

ONE SKETCH FOR ALL: FAST ALGORITHMS FOR COMPRESSED SENSING

A. C. GILBERT, M. J. STRAUSS, J. A. TROPP, AND R. VERSHYNIN

ABSTRACT. Compressed Sensing is a new paradigm for acquiring the compressible signals that arise in many applications. These signals can be approximated using an amount of information much smaller than the nominal length of the signal. Traditional approaches acquire the entire signal and process it to extract the information. The new approach acquires a small number of nonadaptive linear measurements of the signal and uses sophisticated algorithms to determine its information content. Emerging technologies can compute these general linear measurements of a signal at unit cost per measurement.

This paper exhibits a randomized measurement ensemble and a signal reconstruction algorithm that satisfy four requirements:

- (1) The measurement ensemble succeeds for all signals, with high probability over the random choices in its construction.
- (2) The number of measurements of the signal is optimal, except for a factor polylogarithmic in the signal length.
- (3) The running time of the algorithm is polynomial in the amount of information in the signal and polylogarithmic in the signal length.
- (4) The recovery algorithm offers the strongest possible type of error guarantee. Moreover, it is a fully polynomial approximation scheme with respect to this type of error bound.

Emerging applications demand this level of performance. Yet no other algorithm in the literature simultaneously achieves all four of these desiderata.

Date: November 20, 2006.

Key words and phrases. Approximation, embedding, group testing, sketching, sparse approximation, sublinear algorithms.

Gilbert and Tropp are with the Department of Mathematics, The University of Michigan at Ann Arbor, 2074 East Hall, 530 Church St., Ann Arbor, MI 48109-1043. E-mail: {annacg|jtropp}@umich.edu. ACG is an Alfred P. Sloan Research Fellow. ACG has been supported in part by NSF DMS 0354600. JAT has been supported by NSF DMS 0503299.

Strauss is with the Department of Mathematics and the Department of Electrical Engineering and Computer Science, The University of Michigan at Ann Arbor. E-mail: martinjs@umich.edu. MJS has been supported in part by NSF DMS 0354600.

Vershynin is with the Department of Mathematics, The University of California at Davis, Davis, CA 95616. E-mail: vershynin@math.ucdavis.edu. RV is an Alfred P. Sloan Research Fellow. He was also partially supported by the NSF grant DMS 0401032.

1. INTRODUCTION

Compressed Sensing is a new paradigm for acquiring signals, images, and other types of compressible data. These data have the property that they can be approximated using much less information than their nominal dimension would suggest. At present, the standard approach to signal acquisition is to measure a complete copy of the signal and then process it to extract the important information. For example, one typically measures an image in the pixel basis and then applies JPEG compression to obtain a more efficient representation. Instead, the new approach collects a small number of carefully chosen (but nonadaptive) linear measurements that condense the information in the signal. Sophisticated algorithms are used to approximately reconstruct the signal from these measurements.

Some exciting new technological applications are driving the theoretical work on Compressed Sensing. In these applications, it is possible to compute general linear measurements of the signal with unit cost per measurement. Therefore, the acquisition cost is proportional to the number of signal measurements that we take. (This setting stands in contrast with the digital computation of a dot product component by component.)

In traditional signal acquisition models, measurements of the signal have a straightforward interpretation. On the other hand, Compressed Sensing uses measurements that have no real meaning. In particular, there is no simple map from measurement data back to the signal domain. As a result, we are also very concerned about the time it takes to reconstruct signals from measurements.

Scientists and engineers are developing technologies where the computational model of Compressed Sensing applies. They are building cameras [TLW⁺06, WLD⁺06], analog-to-digital converters [DWB05, LKM⁺06, K LW⁺06], and other sensing devices [ZPB06, ZBSG05] that can obtain a general linear measurement of a signal at unit cost. Compressive imaging cameras use a digital micro-mirror array to optically compute inner products of the image with pseudorandom binary patterns. The image is digitally reconstructed from the projections. Popular media have showcased this application. See *Business Week* (Oct. 16, 2006) and *The Economist* (Oct. 28, 2006).

In fact, certain types of Compressed Sensing devices are already widespread, namely CT and MRI scanners. The detector in a Computed Tomography (CT) scanner takes a number of snapshots or profiles of an attenuated X-ray beam as it passes through a patient. The profiles are used to reconstruct a two-dimensional image. Each snapshot of the X-ray beam is, in essence, the line integral of the X-ray beam through the patient (i.e., an inner product).

1.1. Desiderata for Compressed Sensing. Our premise is that, if one measures a highly compressible signal, it is pointless to reconstruct a full-length copy of the signal because it will include a huge number of small, noisy components that bear no information. Instead, a recovery algorithm should directly identify those few components of the signal that are significant. The algorithm should output this compressed representation directly, and its runtime should be roughly proportional to the size of the representation.

Let us be more formal. We are interested in acquiring signals in \mathbb{R}^d that are well approximated by sparse signals with m nonzero components, where $m \ll d$. The measurement process can be represented by an $n \times d$ matrix Ψ , where n is roughly proportional to m rather than d . Each signal \mathbf{f} yields a sketch $\mathbf{v} = \Psi\mathbf{f}$. The recovery algorithm uses the sketch and a description of the measurement matrix to construct a signal approximation $\hat{\mathbf{f}}$ that has only $O(m)$ nonzero components. We want the measurements and the algorithm to satisfy the following properties:

- (1) One (randomly generated) measurement matrix Ψ is used to measure all signals. With high probability over the random choices in its construction, it must succeed for all signals.
- (2) The number of measurements is nearly optimal, namely $n = m \text{ polylog}(d)$.
- (3) The algorithm must run in time $\text{poly}(m, \log d)$.
- (4) Given the sketch of an arbitrary input signal, the algorithm must return a nearly optimal m -term approximation of that signal.

1.2. Our results. We present a linear measurement procedure that takes a near-optimal number of measurements of a signal. We also present HHS Pursuit,¹ a fully polynomial approximation scheme that uses these measurements to construct a sparse estimate of the signal with an optimal error bound. Moreover, this algorithm is exponentially faster than known recovery algorithms that offer equivalent guarantees.

This section states our major theorem and two important corollaries. We establish these results in Section 4 (and the appendices). We discuss the error bounds in Sections 1.3 and 1.4. Section 1.5 provides a comparison with related work.

Given a signal \mathbf{f} , we write \mathbf{f}_m to denote the signal obtained by zeroing all the components of \mathbf{f} except the m components with largest magnitude. (Break ties lexicographically.) We refer to \mathbf{f}_m as the *head* of the signal; it is the best approximation of the signal using at most m terms with respect to any monotonic norm (such as ℓ_p). The vector $\mathbf{f} - \mathbf{f}_m$ is called the *tail* of the signal since it contains the entries with small magnitude.

Theorem 1. *Fix an integer m and a number $\varepsilon \in (0, 1)$. With probability at least 0.99, the random measurement matrix Ψ has the following property. Suppose that \mathbf{f} is a d -dimensional signal, and let $\mathbf{v} = \Psi \mathbf{f}$ be the signal sketch. Given m , ε , and \mathbf{v} , the HHS Pursuit algorithm produces a signal approximation $\hat{\mathbf{f}}$ with $O(m/\varepsilon^2)$ nonzero entries. The approximation satisfies*

$$\|\mathbf{f} - \hat{\mathbf{f}}\|_2 \leq \frac{\varepsilon}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1.$$

The signal sketch has size $(m/\varepsilon^2) \text{polylog}(d/\varepsilon)$, and HHS Pursuit runs in time $(m^2/\varepsilon^4) \text{polylog}(d/\varepsilon)$. The algorithm uses working space $(m/\varepsilon^2) \text{polylog}(d/\varepsilon)$, including storage of the matrix Ψ .

In particular, note that the algorithm recovers every m -term signal without error.

The first corollary shows that we can construct an m -term signal approximation whose ℓ_2 error is within an additive ℓ_1 term of the optimal ℓ_2 error. One can show that this corollary is equivalent with the theorem.

Corollary 2. *Let $\hat{\mathbf{f}}_m$ be the best m -term approximation to the output $\hat{\mathbf{f}}$ of HHS Pursuit. Then*

$$\|\mathbf{f} - \hat{\mathbf{f}}_m\|_2 \leq \|\mathbf{f} - \mathbf{f}_m\|_2 + \frac{2\varepsilon}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1.$$

This result should be compared with Theorem 2 of [CRT06], which gives an analogous bound for the (superlinear) ℓ_1 minimization algorithm. A second corollary provides an ℓ_1 error estimate.

Corollary 3. *Let $\hat{\mathbf{f}}_m$ be the best m -term approximation to the output $\hat{\mathbf{f}}$ of HHS Pursuit. Then*

$$\|\mathbf{f} - \hat{\mathbf{f}}_m\|_1 \leq (1 + 3\varepsilon) \|\mathbf{f} - \mathbf{f}_m\|_1.$$

The error bound in Corollary 3 is more intuitive but substantially weaker than the bound in Theorem 1. One may check this point by considering a signal whose first component equals $m^{-1/4}$ and whose remaining components equal d^{-1} . The ℓ_1 error bound holds even if an algorithm fails to identify the first signal component, but the mixed-norm error bounds do not.

1.3. Compressible signals. A *compressible signal* has the property that its components decay when sorted by magnitude. These signals arise in numerous applications because one can compress the wavelet and Fourier expansions of certain classes of natural signals [DDD98]. A common measure of compressibility is the weak- ℓ_p norm, which is defined for $0 < p < \infty$ as

$$\|\mathbf{f}\|_{w\ell_p} \stackrel{\text{def}}{=} \inf\{r : |f|_{(k)} \leq r \cdot k^{-1/p} \text{ for } k = 1, 2, \dots, d\}.$$

¹The initials HHS stand for ‘‘Heavy Hitters on Steroids,’’ which reflects the strong demands on the algorithm.

The notation $|f|_{(k)}$ indicates the k th largest magnitude of a signal component. When the weak- ℓ_p norm is small for some $p < 2$, the signal can be approximated efficiently by a sparse signal because

$$\|\mathbf{f} - \mathbf{f}_m\|_1 \leq m^{1-1/p} \|\mathbf{f}\|_{w\ell_p}$$

Theorem 1 shows that the computed approximation $\widehat{\mathbf{f}}$ satisfies the error bound

$$\|\mathbf{f} - \widehat{\mathbf{f}}\|_2 \leq \varepsilon m^{1/2-1/p} \|\mathbf{f}\|_{w\ell_p}.$$

In particular, when $p = 1$, the error decays like $m^{-1/2}$.

1.4. Optimality of error bounds. The error guarantees may look strange at first view. Indeed, one might hope to take (m/ε^2) polylog(d) measurements of a signal \mathbf{f} and produce an m -sparse approximation $\widehat{\mathbf{f}}$ that satisfies the error bound

$$\|\mathbf{f} - \widehat{\mathbf{f}}\|_2 \leq (1 + \varepsilon) \|\mathbf{f} - \mathbf{f}_m\|_2.$$

It has been established [GMS05] that this guarantee is possible if we construct a random measurement matrix *for each signal*. On the other hand, Cohen, Dahmen, and DeVore have shown [CDD06] that it is impossible to obtain this error bound simultaneously *for all* signals unless the number of measurements is $\Omega(d)$.

The same authors also proved a more general lower bound [CDD06]. For each p in the range $[1, 2)$, it requires $\Omega(m(d/m)^{2-2/p})$ measurements to achieve

$$\|\mathbf{f} - \widehat{\mathbf{f}}\|_2 \leq C_p m^{1/2-1/p} \|\mathbf{f} - \mathbf{f}_m\|_p \tag{1.1}$$

simultaneously for all signals. This result holds for all possible recovery algorithms. It becomes vacuous when $p = 1$, which is precisely the case delineated in Theorem 1.

It is not hard to check that the number of measurements required by our algorithm is within a polylogarithmic factor of the lower bound.

Proposition 4. *Fix p in the range $[1, 2)$. With $m(d/m)^{2-2/p}$ polylog(d) measurements, the HHS Pursuit algorithm produces for every signal \mathbf{f} an m -term estimate $\widehat{\mathbf{f}}$ such that (1.1) holds.*

1.5. Related work. The major difference between our work and other algorithms for Compressed Sensing is that we simultaneously provide

- (1) a uniform guarantee for all signals,
- (2) a near-optimal number of measurements,
- (3) a sublinear running time, and
- (4) an optimal error bound.

We discuss these points in turn.

Typically, randomized sketches guarantee that “on each signal, with high probability, the algorithm succeeds.” When the application involves adaptiveness or iteration, it is much better to have a uniform guarantees of the form “with high probability, on all signals, the algorithm succeeds.” Most approaches to Compressed Sensing yield uniform guarantees. Exceptions include work on Orthogonal Matching Pursuit (OMP) due to Tropp–Gilbert [TG06] and the algorithms of Cormode–Muthukrishnan [CM06] that achieve the strongest error bounds.

A major selling point for Compressed Sensing is that it uses only m polylog(d) measurements to recover an entire class of compressible signals. Candès–Romberg–Tao [CRT06] and Donoho [Don06] have shown that a linear programming algorithm achieves this goal. The Chaining Pursuit algorithm of the current authors [GSTV06] also has this property. On the other hand, the algorithms of Cormode–Muthukrishnan that yield a uniform guarantee require $\Omega(m^2)$ measurements [CM06].

The major advantage of our work is that most recovery algorithms for Compressed Sensing have runtimes that are at least linear in the length of the input signal. In particular, the linear

programming technique has cost $\Omega(d^{3/2})$. Cormode–Muthukrishnan have developed some sublinear algorithms whose runtimes are comparable with HHS Pursuit [CM06]. The Chaining Pursuit algorithm has running time $m \text{polylog}(d)$, so it is even faster than HHS Pursuit.

Chaining Pursuit is the only algorithm in the literature that achieves the first three desiderata [GSTV06]. The error bound in Chaining Pursuit, however, is of the form

$$\|\mathbf{f} - \hat{\mathbf{f}}\|_{w\ell_1} \leq C \|\mathbf{f} - \mathbf{f}_m\|_1,$$

where $\hat{\mathbf{f}}$ is the m -term signal estimate produced by the algorithm. Not only is this error bound worse than the HHS error bound, but also Chaining Pursuit is not an approximation scheme.

1.6. Roadmap. The next three sections give an overview of our approach. Section 2 provides a detailed description of the measurement matrix required by HHS Pursuit. Section 3 states the HHS algorithm, along with implementation details and pseudocode. Section 4 shows how to draw the corollaries from the main theorem, and it explains how the analysis of the algorithm breaks into two cases. The bulk of the proof is deferred to the back matter. Appendices B and C analyze the two cases in turn. Appendix D studies the behavior of the identification matrix in detail. Finally, Appendix E studies the properties of the estimation matrix.

2. THE MEASUREMENTS

This section describes a random construction of a measurement matrix Ψ . Afterward, we explain how to store and apply the matrix efficiently. For clarity, we focus on the case $\varepsilon = 1$. To obtain an approximation scheme, we substitute m/ε^2 for m , which increases the costs by $(1/\varepsilon)^{O(1)}$.

The matrix Ψ consists of two pieces: an identification matrix Ω and an estimation matrix Φ . We view the matrix as a linear map that acts on a signal \mathbf{f} in \mathbb{R}^d to produce a two-part sketch.

$$\Psi \mathbf{f} = \begin{bmatrix} \Omega \\ \Phi \end{bmatrix} \mathbf{f} = \begin{bmatrix} \mathbf{v}_{\text{id}} \\ \mathbf{v}_{\text{est}} \end{bmatrix}.$$

The first part of the sketch, $\mathbf{v}_{\text{id}} = \Omega \mathbf{f}$, is used to identify large components of the signal quickly. The second part, $\mathbf{v}_{\text{est}} = \Phi \mathbf{f}$, is used to estimate the size of the identified components. Decoupling the identification and estimation steps allows us to produce strong error guarantees.

2.1. Row tensor products. The identification matrix zeroes out many different subsets of the signal components to isolate large components from each other, and then it computes inner products between these restricted signals and a group testing matrix. We construct this restriction map by applying several different restrictions in sequence. We introduce notation for this operation.

If \mathbf{q} and \mathbf{r} are 0–1 vectors, we can view them as masks that determine which entries of a signal appear and which ones are zeroed out. For example, the signal $\mathbf{q} \circ \mathbf{f}$ is the signal \mathbf{f} restricted to the components in \mathbf{q} that equal one. (The notation \circ indicates the Hadamard, or componentwise, product.) The sequential restriction by \mathbf{q} and \mathbf{r} can be written as $(\mathbf{r} \circ \mathbf{q}) \circ \mathbf{f}$. Given 0–1 matrices \mathbf{Q} and \mathbf{R} , we can form a matrix that encodes sequential restrictions by all pairs of their rows. We express this matrix using the *row tensor product*, as in [GSTV06, CM06].

Definition 5. Let \mathbf{Q} be a $q \times d$ matrix and \mathbf{R} an $r \times d$ matrix with rows $\{\mathbf{q}_i : 0 \leq i < q\}$ and $\{\mathbf{r}_k : 0 \leq k < r\}$, respectively. The row tensor product $\mathbf{A} = \mathbf{Q} \otimes_{\text{r}} \mathbf{R}$ is a $qr \times d$ matrix whose rows are $\{\mathbf{q}_i \circ \mathbf{r}_k : 0 \leq i < q, 0 \leq k < r\}$.

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}^{(1)} \\ \mathbf{A}^{(2)} \\ \vdots \\ \mathbf{A}^{(J)} \end{bmatrix} = \begin{bmatrix} \mathbf{S}^{(1)} \otimes_{\text{r}} \mathbf{R}^{(1)} \\ \mathbf{S}^{(2)} \otimes_{\text{r}} \mathbf{R}^{(2)} \\ \vdots \\ \mathbf{S}^{(J)} \otimes_{\text{r}} \mathbf{R}^{(J)} \end{bmatrix} \quad \text{where} \quad \mathbf{S}^{(j)} = \begin{bmatrix} \mathbf{S}_1^{(j)} \\ \mathbf{S}_2^{(j)} \\ \vdots \\ \mathbf{S}_{T_j}^{(j)} \end{bmatrix} \quad \text{and} \quad \mathbf{R}^{(j)} = \begin{bmatrix} \mathbf{R}_1^{(j)} \\ \mathbf{R}_2^{(j)} \\ \vdots \\ \mathbf{R}_{U_j}^{(j)} \end{bmatrix}$$

FIGURE 1. The structure of the isolation matrix \mathbf{A} . See Section 2.2.2 for details.

2.2. The identification operator. The identification matrix $\mathbf{\Omega}$ is a 0–1 matrix with dimensions $O(m \log^2(m) \log(d/m) \log^2(d)) \times d$. It consists of a combination of a structured deterministic matrix and ensembles of simple random matrices. Formally, $\mathbf{\Omega}$ is the row tensor product $\mathbf{\Omega} = \mathbf{B} \otimes_{\text{r}} \mathbf{A}$. The *bit-test matrix* \mathbf{B} has dimensions $O(\log d) \times d$, and the *isolation matrix* \mathbf{A} has dimensions $O(m \log^2(m) \log(d/m) \log d) \times d$.

2.2.1. The bit-test matrix. The matrix \mathbf{B} is a deterministic matrix that contains a row of 1s appended to a 0–1 matrix \mathbf{B}_0 . The matrix \mathbf{B}_0 has dimensions $\log_2[d] \times d$. Its k th column contains the binary expansion of k . Therefore, the inner product of the i th row of \mathbf{B}_0 with a signal \mathbf{f}^T sums the components of \mathbf{f} that have bit i equal to one. The bit-test matrix with $d = 8$ is

$$\mathbf{B} = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{bmatrix}.$$

In coding theory, \mathbf{B} is called the parity check matrix for the extended Hamming code.

2.2.2. The isolation matrix. The isolation matrix \mathbf{A} is a randomly constructed 0–1 matrix with a hierarchical structure. It consists of $O(\log_2 m)$ blocks $\mathbf{A}^{(j)}$ labeled by $j = 1, 2, 4, 8, \dots, J$, where $J = O(m)$. See Figure 1.

Each block, in turn, has further substructure as a row tensor product of two 0–1 matrices: $\mathbf{A}^{(j)} = \mathbf{R}^{(j)} \otimes_{\text{r}} \mathbf{S}^{(j)}$. The second matrix $\mathbf{S}^{(j)}$ is called the *sifting matrix*, and its dimensions are $O(j \log(d/j)) \times d$. The first matrix $\mathbf{R}^{(j)}$ is called the *noise reduction matrix*, and its dimensions are $O((m/j) \log(m) \log d) \times d$.

2.2.3. The sifting matrix. The purpose of the sifting matrix $\mathbf{S}^{(j)}$ is to isolate about j distinguished signal positions from each other. It is a random 0–1 block matrix, as shown in Figure 1. Each submatrix of $\mathbf{S}^{(j)}$ has dimensions $O(j) \times d$, and the number of submatrices is $T_j = O(\log(d/j))$.

The T_j submatrices are fully independent from each other. Each of the submatrices encodes a $O(j)$ -wise independent random assignment of each signal position to a row. The (i, k) entry of the matrix equals one when the k th signal component is assigned to the i th row. Therefore, with high probability, the componentwise product of the i th row of the matrix with \mathbf{f} generates a copy of the signal with $d/O(j)$ components selected and the others zeroed out. For example,

$$\mathbf{S}_t^{(j)} = \begin{bmatrix} 0 & 1 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

This submatrix can also be viewed as a random linear hash function from the space of d keys onto a set of $O(j)$ buckets.

2.2.4. The noise reduction matrix. The purpose of the noise reduction matrix $\mathbf{R}^{(j)}$ is to attenuate the noise in a signal that has a single large component. It is also a random 0–1 block matrix, as seen in Figure 1. Each submatrix $\mathbf{R}_u^{(j)}$ has dimensions $O(\sqrt{(m/j) \log m}) \times d$, and the total number of submatrices $U_j = O(\sqrt{(m/j) \log m} \log d)$.

The submatrices are fully independent from each other. Each one encodes a pairwise independent assignment of each signal positions to a row. The (i, k) entry of the matrix equals one when the k th signal component is assigned to the i th row, as in the sifting matrix. Each submatrix can be viewed as a random linear hash function from d keys to $O(\sqrt{(m/j) \log m})$ buckets.

2.3. The estimation matrix. The estimation matrix Φ is a randomly constructed matrix with complex entries. Let $\lambda = O(\log^4 d)$ and $L = O(m\sqrt{\log m})$. Choose $q \geq \lambda L$. The estimation matrix consists of q rows drawn independently at random from the $d \times d$ discrete Fourier transform (DFT) matrix. The matrix Φ is scaled by $q^{-1/2}$ so its columns have unit ℓ_2 norm.

2.4. Storage costs. The bit test matrix requires no storage as it is straightforward to generate as needed. The isolation and estimation matrices can be generated from short pseudorandom seeds, as needed.

The total storage for the estimation matrix is $q \log(d) = O(m\sqrt{\log m} \log^5 d)$ because it takes $\log d$ bits to store the index of each of the q rows drawn from the DFT matrix.

The total storage for the isolation matrix \mathbf{A} is $O(m \log^3 d)$, which is negligible compared with the cost of the estimation matrix. To obtain the bound for the isolation matrix, we examine the sifting matrices and the noise reduction matrices separately.

First, observe that each block $\mathbf{S}_t^{(j)}$ of the sifting matrix requires $O(j \log d)$ bits.² Since there are $T_j = O(\log(d/j))$ independent blocks for each j , we have a space bound $O(j \log^2 d)$ for $\mathbf{S}_t^{(j)}$. Summing over $j = 1, 2, 4, \dots, Cm$, we find that the sifting matrices require $O(m \log^2 d)$ bits.

Meanwhile, each block $\mathbf{R}_u^{(j)}$ of the noise reduction matrix requires $O(\log d)$ bits. There are $U_j = O(\sqrt{m/j} \log m \log d)$ blocks, giving a space bound $O(\sqrt{m/j} \log m \log^2 d)$ for $\mathbf{R}^{(j)}$. Summing on j , we see that the noise reduction matrices require space $O(m^{1/2+o(1)} \log^2 d)$.

2.5. Encoding time. The encoding time depends on the technology for computing measurements. If we can compute inner products in constant time, the encoding time is proportional to the number of measurements. This section focuses on the case where the cost is proportional to the minimal sparsity of the vectors that appear in the inner product. This analysis plays a role in determining the runtime of the algorithm.

We show that the time required to measure a signal \mathbf{f} that has exactly one nonzero component is $O(m\sqrt{\log m} \log^4(d))$ word operations. This analysis implies a time bound for measuring a vector with more nonzeros. The time (in word operations) to generate the estimation matrix Φ and to multiply Φ by \mathbf{f} dominates the time (in bit operations) to generate the identification matrix Ω and to multiply Ω by \mathbf{f} .

The time cost for measuring a nonzero component ℓ of a signal with the estimation operator Φ is $q = O(m\sqrt{\log m} \log^4 d)$, assuming column ℓ of Φ has been computed. The (k, ℓ) entry of Φ is simply $q^{-1/2} \exp\{-2\pi i t_k \omega_\ell / d\}$, which is computed from t_k and ω_ℓ in a constant number of word operations.

We now turn to the identification matrix Ω . Each of its columns contains roughly \sqrt{m} nonzeros, so we can ignore the cost to *apply* the matrix. To *generate* a column of Ω , we form the sifting matrices and noise reduction matrices and then compute their row tensor product. Afterward, we form the row tensor product with the bit-test matrix. The total cost is $O(m \log^4 d)$ bit operations, as follows.

²We generate the random variables using a polynomial of degree j over a field of size around d . Without loss of generality, we may assume that m and d are powers of two.

To construct $\mathbf{S}_t^{(j)}$, one can use a polynomial of degree $O(j)$ over a field of size approximately d to obtain $O(j)$ -wise independent random variables. This step must be repeated $T_j = O(\log(d/j))$ times, for a total of $O(j \log(d/j))$ field operations over a field of size around d . Therefore, to generate $\mathbf{S}^{(j)}$ costs $O(j \log^3 d)$ bit operations. To generate each $\mathbf{R}_u^{(j)}$ requires $\log d$ operations, so the cost to generate $\mathbf{R}^{(j)}$ is roughly \sqrt{m} , which is negligible. To build $\mathbf{R}^{(j)} \otimes_r \mathbf{S}^{(j)}$ from its factors costs $T_j \cdot U_j = O(\sqrt{m} \log^3 d)$, so it can also be ignored. Summing on j , we find that the bit time to construct \mathbf{A} is $O(m \log^3 d)$ bit operations. The row tensor product with the bit-test matrix gives an additional factor of $O(\log d)$ bit operations.

To summarize, the total time cost to use Ψ to measure a signal with k nonzero entries is $O(km \text{polylog}(d))$. During the HHS Pursuit algorithm, we must encode a list L of $O(m\sqrt{\log m})$ signals with one nonzero component each. The time cost for this encoding procedure is $m^2 \text{polylog}(d)$ if we use the straightforward algorithm for matrix–vector multiplication.³

3. THE HHS ALGORITHM

The HHS algorithm is an iterative procedure, where each iteration has several stages. The first stage is designed to identify a small set of signal positions that carry a constant proportion of the remaining energy. The second stage estimates the values of the signal on these components. The third stage adds the new approximation to the old approximation and prunes it so that it contains only $O(m)$ nonzero components. The fourth stage encodes the new approximation using the measurement matrix and subtracts it from the initial sketch to obtain a sketch of the current residual signal. See Figure 2 for pseudocode. For brevity, we use the term *spike* to refer to the location and size of a single signal component.

The algorithm also employs a preprocessing phase. The preprocessing step encodes the signal with the Chaining Pursuit measurement matrix \mathbf{C} and executes the Chaining Pursuit algorithm to produce a good initial approximation \mathbf{a}^{init} of the input signal. This initial approximation has at most m nonzero components. We encode \mathbf{a}^{init} with the HHS measurement matrix and subtract it from the original chaining sketch $\Psi \mathbf{f}$ to obtain a sketch \mathbf{s} of the initial residual $\mathbf{f} - \mathbf{a}^{\text{init}}$. See Figure 3 for pseudocode.

The ℓ_2 norm of the residual after preprocessing is proportional with m and the optimal ℓ_1 error with m terms. This fact ensures the algorithm recovers sparse signals exactly and that it requires only $O(\log m)$ iterations to reduce the error by a polynomial factor in m . The correctness of the preprocessing phase follows from our previous work [GSTV06].

If we do not run the optional preprocessing step, then the initial residual sketch \mathbf{s} and list L of spike locations and values are $\Psi \mathbf{f}$ and empty, respectively. In that case, one can run the algorithm for $O(\log \Delta)$ iterations, where $\Delta = \|\mathbf{f}\|_1 / \|\mathbf{f} - \mathbf{f}_m\|_1$ is the *dynamic range* of the problem.

3.1. Comparison with other sublinear algorithms. HHS Pursuit resembles a number of sublinear approximation algorithms in the literature [GSTV06, GMS05, CM06]. HHS Pursuit differs from these other algorithms in the way it makes progress toward approximating the signal. The other algorithms search for large spikes in a greedy fashion, and they cannot reject a spike once it has been introduced into the approximation. HHS Pursuit, on the other hand, searches for collections of spikes that carry a lot of energy. This collection may include large spikes, or it may include a relatively large number of small spikes. Spikes can be removed from the list during the course of the algorithm if their values become small. This approach allows us to maintain a small list of spikes that carries more and more of the signal energy.

³Note that encoding requires us to compute a partial discrete Fourier transform with unequally-spaced points on the domain and codomain of the transform. We are not aware of any nontrivial algorithm for this problem, despite the existence of faster algorithms [Bey95] for problems that are superficially similar.

Algorithm: HHS Pursuit

Inputs: The number m of spikes, the HHS measurement matrix Ψ ,
the initial sketch $v = \Psi f$, the initial list L of m spikes,
the initial residual sketch s

Output: A list L of $O(m)$ spikes

For each iteration $k = 0, 1, \dots, O(\log m)$ {
 For each scale $j = 1, 2, 4, \dots, O(m)$ {
 Initialize $L' = \emptyset$.
 For each row of $A^{(j)}$ {
 Use the $O(\log d)$ bit tests to identify one spike location
 }
 Retain a list L'_j of the spike locations
 that appear $\Omega(\sqrt{m/j} \log m \log(d/j) \log d)$ times each
 Update $L' \leftarrow L' \cup L'_j$
 }
 Estimate values for the spikes in L' by forming $\Phi_{L'}^\dagger s_{\text{est}}$
 with Jacobi iteration
 Update L by adding the spikes in L'
 If a spike is duplicated, add the two values together
 Prune L to retain the $O(m)$ largest spikes
 Encode these spikes with measurement matrix Ψ
 Subtract encoded spikes from original sketch v to form
 a new residual sketch s
}

FIGURE 2. Pseudocode for the HHS Pursuit algorithm

3.2. Implementation. The HHS Pursuit algorithm is easily implemented with standard data structures. There are a few steps that require a short discussion. The Jacobi iteration is a standard algorithm from numerical analysis. See Lemma 13 for details of the application here.

The bit tests also require explanation. Each row of the isolation matrix \mathbf{A} effectively generates a copy of the input signal with many locations zeroed out. The bit-test matrix calculates inner products between its rows and the restricted signal. The bit tests attempt to use these numbers to find the location of the largest entry in the restricted signal.

Suppose that the bit tests yield the following $\log_2 \lceil d \rceil + 1$ numbers:

$$c, \quad b(0), \quad b(1), \quad \dots, \quad b(\log_2 \lceil d \rceil - 1).$$

The number c arises from the top row of the bit test matrix, so it is the sum of the components of the restricted signal. We estimate a spike location as follows. If $|b(i)| \geq |c - b(i)|$, then the i th bit of the estimated location is zero. Otherwise, the i th bit of the estimated location is one. It is clear that that the estimated location is correct if the restricted signal contains one large component, and the remaining components have ℓ_1 norm smaller than the magnitude of the large component.

We encode the recovered spikes by accessing the columns of the identification and estimation matrices corresponding to the locations of these spikes and then re-scaling these columns by the spike values. Note that this step requires us to generate arbitrary columns.

3.3. Resource Requirements. In Section 2.4, we showed that we need space $m \text{ polylog}(d)$ to store pseudorandom seeds from which columns of the measurement operator can be generated as needed,

Algorithm: (Optional) Chaining Pursuit Preprocessing

Inputs: The number m of spikes, the Chaining measurement matrix C ,
the Chaining sketch $w = Cf$, the HHS measurement matrix Ψ ,
the HHS sketch $v = \Psi f$

Outputs: A list L of m spikes, the residual sketch s

Run ChainingPursuit(m, w, C) to obtain a list L of m spikes
Encode the spikes in L using the HHS measurement matrix Ψ .
Subtract the encoded spikes from v to form the residual sketch s .

FIGURE 3. Pseudocode for Chaining Pursuit Preprocessing

in time $m \text{polylog}(d)$ each. In Lemma 13, we review how to apply $\Phi_L^\dagger s_{\text{est}}$ via Jacobi iteration in time $m^2 \text{polylog}(d)$. It follows that our algorithm requires $m \text{polylog}(d)$ working space and $m^2 \text{polylog}(d)$ time. Our algorithm becomes an approximation scheme by substituting m/ε^2 for m . This increases the space to $(m/\varepsilon^2) \text{polylog}(d/\varepsilon)$ and the overall time cost to $(m^2/\varepsilon^4) \text{polylog}(d/\varepsilon)$.

4. TOP-LEVEL ANALYSIS OF THE ITERATION

This section describes, at the highest level, why HHS works. We establish the following result.

Theorem 6. *Fix m . Assume that Ψ is a measurement matrix that satisfies the conclusions of Lemmas 19 and 28. Suppose that f is a d -dimensional signal. Given the sketch $v = \Psi f$, the HHS Pursuit algorithm produces a signal \hat{f} with at most $8m$ nonzero entries. This signal estimate satisfies*

$$\|f - \hat{f}\|_2 \leq \frac{20}{\sqrt{m}} \|f - f_m\|_1.$$

We give an overview of the proof in the next subsection. Let m and ε be fixed. Observe that we can apply the theorem with $m' = m/\varepsilon^2$ to obtain a signal estimate \hat{f} with $8m'$ terms that satisfies the error bound

$$\|f - \hat{f}\|_2 \leq \frac{20\varepsilon}{\sqrt{m}} \|f - f_{m'}\|_1.$$

The running time increases by a factor of $(1/\varepsilon^4) \text{polylog}(1/\varepsilon)$. This leads to Theorem 1.

Next, we establish Corollary 2. After running the algorithm, we can prune the final approximation \hat{f} to retain only its m largest terms \hat{f}_m . Calculate that

$$\begin{aligned} \|f - \hat{f}_m\|_2 &\leq \|\hat{f} - \hat{f}_m\|_2 + \|f - \hat{f}\|_2 \\ &\leq \|\hat{f} - f_m\|_2 + \|f - \hat{f}\|_2 \\ &\leq \|f - f_m\|_2 + 2\|f - \hat{f}\|_2 \\ &\leq \|f - f_m\|_2 + \frac{40\varepsilon}{\sqrt{m}} \|f - f_{m'}\|_1 \\ &\leq \|f - f_m\|_2 + \frac{40\varepsilon}{\sqrt{m}} \|f - f_m\|_1. \end{aligned}$$

The first and third inequalities follow from the triangle inequality. The second inequality holds because \hat{f}_m is a best m -term approximation to \hat{f} . The fifth inequality holds because $f_{m'}$ is a best m' -term approximation to f , while f_m is an approximation with fewer terms.

The proof of Corollary 3 is somewhat more involved, so we postpone it to Appendix A.

4.1. The proof. Let \mathbf{f} denote our target signal. The goal of the algorithm is to identify a small set of signal components that carry most of the energy in the signal and to estimate the magnitudes of those components well. We argue that, when our signal estimate is poor, the algorithm makes substantial progress toward this goal. When our estimate is already good, the algorithm does not make it much worse.

The Chaining Preprocessing step is designed to produce a good initial approximation so that fewer iterations of the HHS algorithm are necessary to obtain an approximation of the required quality. The output of the preprocessing step is a signal \mathbf{a}^{init} that satisfies

$$\|\mathbf{a}^{\text{init}}\|_0 \leq m \quad \text{and} \quad \|\mathbf{f} - \mathbf{a}^{\text{init}}\|_2 \leq m \|\mathbf{f} - \mathbf{f}_m\|_1.$$

These claims follow from Theorem 2 in [GSTV06], provided that $m \geq C \log m$ for an absolute constant C .

Let \mathbf{a} denote our approximation to the signal \mathbf{f} at the beginning of an iteration. The pruning step of the algorithm ensures that we have the loop invariant $\|\mathbf{a}\|_0 \leq 8m$. We will abbreviate $p = 8m$ when it makes the argument clearer.

In a given iteration, the performance of the algorithm depends on the size of the residual. In Appendix B, we establish that if the approximation is poor then the algorithm improves it substantially. More precisely, assume that the current approximation \mathbf{a} satisfies

$$\|\mathbf{f} - \mathbf{a}\|_2 > \frac{1}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1. \quad (\text{Case 1})$$

Then one iteration produces a new approximation \mathbf{a}^{new} for which

$$\|\mathbf{f} - \mathbf{a}^{\text{new}}\|_2 \leq \frac{1}{2} \|\mathbf{f} - \mathbf{a}\|_2.$$

On the other hand, when the approximation is good, then the algorithm produces a new approximation that is not too bad. Suppose that \mathbf{a} satisfies

$$\|\mathbf{f} - \mathbf{a}\|_2 \leq \frac{1}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1. \quad (\text{Case 2})$$

Then the next approximation \mathbf{a}^{new} satisfies

$$\|\mathbf{f} - \mathbf{a}^{\text{new}}\|_2 \leq \frac{20}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1.$$

We prove this result in Appendix C.

After $1.5 \log_2 m$ iterations, we claim that the algorithm produces a final approximation \mathbf{a}^{fin} that satisfies

$$\|\mathbf{f} - \mathbf{a}^{\text{fin}}\|_2 \leq \frac{20}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1.$$

First, note that the condition (Case 1) must fail during some iteration or else

$$\|\mathbf{f} - \mathbf{a}^{\text{fin}}\|_2 \leq \left(\frac{1}{2}\right)^{1.5 \log_2 m} \|\mathbf{f} - \mathbf{a}^{\text{init}}\|_1 \leq \frac{1}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1,$$

which is a contradiction. In the subsequent iterations, one of two events must occur. Either

$$\|\mathbf{f} - \mathbf{a}\|_2 \leq \frac{1}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1.$$

or else

$$\frac{1}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1 < \|\mathbf{f} - \mathbf{a}\|_2 \leq \frac{20}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1.$$

In both cases, the arguments above show that

$$\|\mathbf{f} - \mathbf{a}^{\text{new}}\|_2 \leq \frac{20}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1.$$

ACKNOWLEDGMENTS

We wish to thank Mark Rudelson for showing us this lovely proof of Lemma 29. We thank Howard Karloff and Piotr Indyk for many insightful discussions.

REFERENCES

- [Bey95] G. Beylkin. On the Fast Fourier Transform of functions with singularities. *Appl. Comp. Harmonic Anal.*, 2:363–381, 1995.
- [CDD06] A. Cohen, W. Dahmen, and R. DeVore. Compressed Sensing and best k -term approximation. Submitted for publication, July 2006.
- [CM06] G. Cormode and S. Muthukrishnan. Combinatorial algorithms for Compressed Sensing. In *Proc. 40th Ann. Conf. Information Sciences and Systems*, Princeton, Mar. 2006.
- [CRT06] E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59(8):1208–1223, 2006.
- [DDD98] D. L. Donoho, I. Daubechies, R. DeVore, and M. Vetterli. Data compression and harmonic analysis. *IEEE Trans. Info. Theory*, 44(6):2435–2476, 1998.
- [Don06] D. L. Donoho. Compressed Sensing. *IEEE Trans. Info. Theory*, 52(4):1289–1306, Apr. 2006.
- [DWB05] M. F. Duarte, M. B. Wakin, and R. G. Baraniuk. Fast reconstruction of piecewise smooth signals from random projections. In *Proc. SPARS05*, Rennes, France, Nov. 2005.
- [GMS05] A. C. Gilbert, S. Muthukrishnan, and M. J. Strauss. Improved time bounds for near-optimal sparse Fourier representation via sampling. In *Proc. SPIE Wavelets XI*, San Diego, 2005.
- [GSTV06] A. C. Gilbert, M. J. Strauss, J. A. Tropp, and R. Vershynin. Algorithmic linear dimension reduction in the ℓ_1 norm for sparse vectors. Submitted for publication, 2006.
- [Jan94] S. Janson. Large deviation inequalities for sums of indicator variables. Technical report, Mathematics Dept., Uppsala Univ., 1994.
- [KLW⁺06] Sami Kirolos, Jason Laska, Michael Wakin, Marco Duarte, Dror Baron, Tamer Ragheb, Yehia Massoud, and Richard Baraniuk. Analog-to-information conversion via random demodulation. In *IEEE Dallas Circuits and Systems Workshop (DCAS)*, Dallas, Texas, Oct. 2006.
- [LKM⁺06] Jason Laska, Sami Kirolos, Yehia Massoud, Richard Baraniuk, Anna Gilbert, Mark Iwen, and Martin Strauss. Random sampling for analog-to-information conversion of wideband signals. In *IEEE Dallas Circuits and Systems Workshop (DCAS)*, Dallas, Texas, Oct. 2006.
- [RV06] M. Rudelson and R. Vershynin. Sparse reconstruction by convex relaxation: Fourier and Gaussian measurements. In *Proc. 40th Ann. Conf. Information Sciences and Systems*, Princeton, Mar. 2006.
- [TG06] J. A. Tropp and A. C. Gilbert. Signal recovery from random measurements via Orthogonal Matching Pursuit. Submitted for publication, Apr. 2005. Revised, Nov. 2006, 2006.
- [TLW⁺06] Dharmpal Takhar, Jason Laska, Michael B. Wakin, Marco F. Duarte, Dror Baron, Shriram Sarvotham, Kevin Kelly, and Richard G. Baraniuk. A new compressive imaging camera architecture using optical-domain compression. In *Proc. IS&T/SPIE Symposium on Electronic Imaging*, 2006.
- [VM82] V. A. Vatutin and V. G. Mikhailov. Limit theorems for the number of empty cells in an equiprobable scheme for group allocation of particles. *Theor. Probab. Appl.*, 27:734–743, 1982.
- [WLD⁺06] Michael Wakin, Jason Laska, Marco Duarte, Dror Baron, Shriram Sarvotham, Dharmpal Takhar, Kevin Kelly, and Richard Baraniuk. Compressive imaging for video representation and coding. In *Proc. Picture Coding Symposium 2006*, Beijing, China, Apr. 2006.
- [ZBSG05] Y.H. Zheng, D. J. Brady, M. E. Sullivan, and B. D. Guenther. Fiber-optic localization by geometric space coding with a two-dimensional gray code. *Applied Optics*, 44(20):4306–4314, 2005.
- [ZPB06] Y. H. Zheng, N. P. Pitsianis, and D. J. Brady. Nonadaptive group testing based fiber sensor deployment for multiperson tracking. *IEEE Sensors Journal*, 6(2):490–494, 2006.

APPENDIX A. PROOF OF THE ℓ_1 ERROR BOUND

In this appendix, we establish Corollary 3, the fact that the pruned output of HHS Pursuit also satisfies an ℓ_1 error bound.

Corollary 3. *Let $\widehat{\mathbf{f}}_m$ be the best m -term approximation to the output $\widehat{\mathbf{f}}$ of HHS Pursuit. Then*

$$\|\mathbf{f} - \widehat{\mathbf{f}}_m\|_1 \leq (1 + 3\varepsilon) \|\mathbf{f} - \mathbf{f}_m\|_1.$$

Proof. On input \mathbf{f} , the output $\widehat{\mathbf{f}}$ of the HHS algorithm satisfies

$$\|\mathbf{f} - \widehat{\mathbf{f}}\|_2 \leq \frac{\varepsilon}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1. \quad (\text{A.1})$$

We partition $\{1, 2, 3, \dots, d\}$ into two pieces as follows. Let $H = \text{supp}(\mathbf{f}_m) \cup \text{supp}(\widehat{\mathbf{f}}_m)$ and let H^c denote the complement of H . We write $\mathbf{f}|_A$ to denote the signal restricted to its components in a set A . Now make a calculation:

$$\begin{aligned} \|\mathbf{f} - \widehat{\mathbf{f}}_m\|_1 &= \|(\mathbf{f} - \widehat{\mathbf{f}}_m)|_H\|_1 + \|(\mathbf{f} - \widehat{\mathbf{f}}_m)|_{H^c}\|_1 \\ &\leq \|(\widehat{\mathbf{f}} - \widehat{\mathbf{f}}_m)|_H\|_1 + \|(\mathbf{f} - \widehat{\mathbf{f}})|_H\|_1 + \|(\mathbf{f} - \widehat{\mathbf{f}}_m)|_{H^c}\|_1 \\ &\leq \|(\widehat{\mathbf{f}} - \mathbf{f}_m)|_H\|_1 + \|(\mathbf{f} - \widehat{\mathbf{f}})|_H\|_1 + \|(\mathbf{f} - \widehat{\mathbf{f}}_m)|_{H^c}\|_1 \\ &\leq \|(\mathbf{f} - \mathbf{f}_m)|_H\|_1 + 2\|(\mathbf{f} - \widehat{\mathbf{f}})|_H\|_1 + \|(\mathbf{f} - \widehat{\mathbf{f}}_m)|_{H^c}\|_1 \\ &= \|\mathbf{f} - \mathbf{f}_m\|_1 + 2\|(\mathbf{f} - \widehat{\mathbf{f}})|_H\|_1 \\ &\leq \|\mathbf{f} - \mathbf{f}_m\|_1 + 2\sqrt{2m}\|(\mathbf{f} - \widehat{\mathbf{f}})|_H\|_2 \\ &\leq \|\mathbf{f} - \mathbf{f}_m\|_1 + 2\sqrt{2}\varepsilon\|\mathbf{f} - \mathbf{f}_m\|_1 \\ &< (1 + 3\varepsilon)\|\mathbf{f} - \mathbf{f}_m\|_1. \end{aligned}$$

The first line follows from the fact that H and H^c are disjoint. We use the triangle inequality in the second and fourth inequalities. In the third inequality, we use the fact that $\widehat{\mathbf{f}}_m$ bests \mathbf{f}_m as an m -term representation of $\widehat{\mathbf{f}}$, along with the fact that both $\widehat{\mathbf{f}}_m$ and \mathbf{f}_m are supported on H . The fifth line relies on the observation that $\mathbf{f} - \widehat{\mathbf{f}}_m = \mathbf{f} = \mathbf{f} - \mathbf{f}_m$ on H^c and the disjointness of H and H^c . The sixth inequality depends on $|H| \leq 2m$ and Cauchy–Schwarz. The seventh inequality removes the restriction to H and invokes the mixed-norm approximation bound (A.1). \square

APPENDIX B. CASE 1: LARGE RESIDUAL

Suppose that (Case 1) is in force at the beginning of an iteration. This section shows that the iteration reduces the error in the approximation by a factor of two.

B.1. Signal Properties. This section describes some generic properties of signals that are important in the analysis. First, we show that the ℓ_2 norm of the tail of a signal is much smaller than the ℓ_1 norm of the entire signal.

Lemma 7. *For any signal \mathbf{g} , it holds that*

$$\|\mathbf{g} - \mathbf{g}_t\|_2 \leq \frac{1}{2\sqrt{t}} \|\mathbf{g}\|_1.$$

In particular, for $\mathbf{g} = \mathbf{f} - \mathbf{f}_m$, we have

$$\|\mathbf{f} - \mathbf{f}_p\|_2 \leq \frac{1}{2\sqrt{7m}} \|\mathbf{f} - \mathbf{f}_m\|_1$$

since $p - m = 7m$.

Proof. Let μ denote the magnitude of the $(t + 1)$ th largest entry of \mathbf{g} . If $\mu = 0$, the result clearly holds. Otherwise, we can make the following calculation:

$$\begin{aligned} \frac{\|\mathbf{g} - \mathbf{g}_t\|_2}{\|\mathbf{g}\|_1} &\leq \frac{\sqrt{\|\mathbf{g} - \mathbf{g}_t\|_1 \|\mathbf{g} - \mathbf{g}_t\|_\infty}}{\|\mathbf{g}_t\|_1 + \|\mathbf{g} - \mathbf{g}_t\|_1} \\ &\leq \frac{\sqrt{\|\mathbf{g} - \mathbf{g}_t\|_1 \cdot \mu}}{t\mu + \|\mathbf{g} - \mathbf{g}_t\|_1} \\ &\leq \frac{\sqrt{t\mu^2}}{t\mu + t\mu} \\ &= \frac{1}{2\sqrt{t}}. \end{aligned}$$

The last inequality holds because the function $x \mapsto \sqrt{bx}/(c+x)$ achieves its maximum at $x = c$. \square

When the condition (Case 1) holds, most of the energy in the signal is concentrated in its largest components. Let $\mathbf{r} = \mathbf{f} - \mathbf{a}$ denote the residual signal. The number α is this lemma in a constant that will be fixed later.

Lemma 8 (Heads and Tails). *Suppose that (Case 1) is in force. Fix a number $\alpha \geq 1$, and let M be the smallest power of two that exceeds*

$$16\alpha^2 m + 9m.$$

Then the following bounds hold.

$$\|\mathbf{r}\|_2 \geq \alpha \|\mathbf{r} - \mathbf{r}_M\|_2 \tag{B.1}$$

$$\|\mathbf{r}\|_2 \geq \frac{\alpha}{\sqrt{M}} \|\mathbf{r}\|_1. \tag{B.2}$$

Proof. The loop invariant $\|\mathbf{a}\|_0 \leq 8m$ implies that the total number of terms in the signal $(\mathbf{a} - \mathbf{f}_m)$ is not more than $9m$. Therefore, relation (Case 1) implies that

$$\begin{aligned} \|\mathbf{r}\|_2 = \|\mathbf{f} - \mathbf{a}\|_2 &> \frac{1}{\sqrt{m}} \|(\mathbf{f} - \mathbf{a}) + (\mathbf{a} - \mathbf{f}_m)\|_1 \\ &\geq \frac{1}{\sqrt{m}} \|\mathbf{r} - \mathbf{r}_{9m}\|_1 \\ &= \frac{3}{\sqrt{9m}} \|\mathbf{r} - \mathbf{r}_{9m}\|_1. \end{aligned} \tag{B.3}$$

because \mathbf{r}_{9m} is the best approximation to $\mathbf{r} = \mathbf{f} - \mathbf{a}$ using $9m$ terms. Next, apply Lemma 7 to the residual tail $\mathbf{g} = \mathbf{r} - \mathbf{r}_{9m}$ with $t = 16\alpha^2 m$ to obtain

$$\|\mathbf{r} - \mathbf{r}_{9m}\|_1 \geq 2\sqrt{16\alpha^2 m} \|\mathbf{r} - \mathbf{r}_M\|_2. \tag{B.4}$$

Introduce this relation into (B.3) to reach

$$\|\mathbf{r}\|_2 \geq 8\alpha \|\mathbf{r} - \mathbf{r}_M\|_2.$$

This estimate implies (B.1).

To develop the second bound, use Cauchy–Schwarz to see that

$$3 \|\mathbf{r}\|_2 \geq 3 \|\mathbf{r}_{9m}\|_2 \geq \frac{3}{\sqrt{9m}} \|\mathbf{r}_{9m}\|_1.$$

Add this inequality to (B.3) and combine the ℓ_1 norms to obtain

$$4 \|\mathbf{r}\|_2 \geq \frac{3}{\sqrt{9m}} \|\mathbf{r}\|_1.$$

On account that $\sqrt{M/9m} > (4/3)\alpha$, we conclude that (B.2) holds. \square

B.2. Identification. In Section D, we prove that the identification step generates a list L of signal positions that satisfies

$$|L| = O(M\sqrt{\log M}).$$

Now we argue that the signal positions in the list carry most of the energy in the residual signal. This fact ensures that the iteration is making progress toward finding significant signal positions. Moreover, it guarantees that the estimation step can accurately predict the values of the signal positions listed in L .

Lemma 9 (Identification). *Suppose that the condition (Case 1) holds. Then the identification step creates a list L of signal positions such that*

$$\begin{aligned} \|\mathbf{r}|_{L^c}\|_2 &\leq 4\alpha^{-1} \|\mathbf{r}\|_2 \\ \frac{1}{\sqrt{M}} \|\mathbf{r}|_{L^c}\|_1 &\leq 4\alpha^{-1} \|\mathbf{r}\|_2. \end{aligned}$$

Proof. The parts of the residual that we miss fall into four categories. In the first three cases, we show that we can simply disregard parts of the residual that have little energy. The fourth case requires some serious work, and we defer the details until Section D.

Let H index the M components of the residual that are largest in magnitude, breaking any ties lexicographically. We first consider the residual tail, which consists of the components in H^c . Lemma 8 shows that

$$\begin{aligned} \|\mathbf{r}|_{H^c}\|_2 &\leq \alpha^{-1} \|\mathbf{r}\|_2 \\ \frac{1}{\sqrt{M}} \|\mathbf{r}|_{H^c}\|_1 &\leq \alpha^{-1} \|\mathbf{r}\|_2. \end{aligned}$$

We decrease the left-hand sides by restricting the residual to the smaller set of components $H^c \cap L^c$.

Next, look at the small components of the residual indexed in H . Let H_1 be the subset of components in H that have magnitude no greater than $M^{-1} \|\mathbf{r}\|_1$. Since H_1 contains no more than M components,

$$\|\mathbf{r}|_{H_1}\|_2 \leq \left(M \cdot \frac{1}{M^2} \|\mathbf{r}\|_1^2 \right)^{1/2}.$$

Turning to Lemma 8, we see that

$$\|\mathbf{r}|_{H_1}\|_2 \leq \alpha^{-1} \|\mathbf{r}\|_2.$$

For the remaining two parts, we require a few definitions. Let s be a power of two between 1 and $M/2$, and note that s takes $\log_2 M$ values. We define the s th band of the residual to be the set

$$B_s = \left\{ i : \frac{1}{2s} \|\mathbf{r}\|_1 < |r_i| \leq \frac{1}{s} \|\mathbf{r}\|_1 \right\}.$$

We say that the s th band is *insignificant* if

$$\|\mathbf{r}|_{B_s}\|_2 \leq \frac{\alpha^{-1}}{\sqrt{\log_2 M}} \|\mathbf{r}\|_2. \tag{B.5}$$

Otherwise, the s th band is *significant*.

Let H_2 be the subset of H that lists components in the insignificant bands. Since that are at most $\log_2 M$ insignificant bands,

$$\|\mathbf{r}|_{H_2}\|_2 \leq \alpha^{-1} \|\mathbf{r}\|_2.$$

Fourth, define H_3 to be the subset of H that lists components in the significant bands that do not appear in the list L . In Section D, we will prove that

$$\left\| \mathbf{r}|_{H_3} \right\|_2 \leq \alpha^{-1} \|\mathbf{r}\|_2.$$

Adding up the contributions to the residual from H_1 , H_2 , and H_3 , we have

$$\left\| \mathbf{r}|_{H_1 \cup H_2 \cup H_3} \right\|_2 \leq 3\alpha^{-1} \|\mathbf{r}\|_2.$$

To obtain the ℓ_1 norm estimate, note that $|H_1 \cup H_2 \cup H_3| \leq |H| = M$. Hence

$$\frac{1}{\sqrt{M}} \left\| \mathbf{r}|_{H_1 \cup H_2 \cup H_3} \right\|_1 \leq \left\| \mathbf{r}|_{H_1 \cup H_2 \cup H_3} \right\|_2 \leq 3\alpha^{-1} \|\mathbf{r}\|_2.$$

Add the contribution from $L^c \cap H^c$ to each bound to complete the proof. \square

B.3. Estimation. This section establishes that the estimation step produces an approximation \mathbf{b} to the residual \mathbf{r} that lies relatively close to the residual, even though it contains $O(M\sqrt{\log M})$ nonzero entries.

First, we state two technical results that are essential to the proof. We discuss these results in more detail in Section E.

Lemma 10 (Restricted Isometry). *The estimation matrix Φ has the property that every $|L|$ -column submatrix \mathbf{A} satisfies*

$$\frac{1}{2} \|\mathbf{x}\|_2 \leq \|\mathbf{A}\mathbf{x}\|_2 \leq \frac{3}{2} \|\mathbf{x}\|_2.$$

for every vector \mathbf{x} .

Lemma 11. *The estimation matrix Φ has the property that*

$$\|\Phi\mathbf{x}\|_2 \leq \frac{3}{2} \left[\|\mathbf{x}\|_2 + \frac{1}{\sqrt{M}} \|\mathbf{x}\|_1 \right].$$

for every vector \mathbf{x} .

Let L be the list of signal components generated by the identification step, and define Φ_L to be the estimation matrix restricted to these L columns. The estimation step uses an iterative algorithm to apply the pseudoinverse Φ_L^\dagger to the estimation sketch $\mathbf{s}_{\text{est}} = \Phi\mathbf{r}$. This yields an approximation to the values of the signal on the set L .

Lemma 12. *Suppose that the list L of coefficient positions satisfies the conclusions of Lemma 9. Then the estimation step produces an approximation \mathbf{b} of the residual \mathbf{r} that satisfies*

$$\|\mathbf{r} - \mathbf{b}\|_2 \leq 0.15 \|\mathbf{r}\|_2.$$

Proof. The ideal approximation $\mathbf{b}^{\text{ideal}}$ is calculated by restricting the estimation matrix to L and applying its pseudoinverse to the current data:

$$\mathbf{b}^{\text{ideal}}|_L = \Phi_L^\dagger \Phi\mathbf{r} = \mathbf{r}|_L + \Phi_L^\dagger \Phi(\mathbf{r}|_{L^c}).$$

The vector $\mathbf{b}^{\text{ideal}}|_{L^c}$ is identically zero. It follows that

$$\|\mathbf{r} - \mathbf{b}^{\text{ideal}}\|_2 \leq \|\mathbf{r}|_{L^c}\|_2 + \left\| \Phi_L^\dagger \Phi(\mathbf{r}|_{L^c}) \right\|_2. \quad (\text{B.6})$$

Lemma 9 establishes

$$\|\mathbf{r}|_{L^c}\|_2 \leq 4\alpha^{-1} \|\mathbf{r}\|_2. \quad (\text{B.7})$$

It remains to estimate the second norm on the right-hand side of (B.6).

Lemma 10, the restricted isometry bound, shows that $\|\Phi_L^\dagger\|_{2,2} \leq 2$. Applying this bound, we obtain

$$\left\| \Phi_L^\dagger \Phi(\mathbf{r}|_{L^c}) \right\|_2 \leq 2 \|\Phi(\mathbf{r}|_{L^c})\|_2.$$

Next, bound the right-hand side via Lemma 11.

$$\left\| \Phi_L^\dagger \Phi(\mathbf{r}|_{L^c}) \right\|_2 \leq 3 \left[\|\mathbf{r}|_{L^c}\|_2 + \frac{1}{\sqrt{M}} \|\mathbf{r}|_{L^c}\|_1 \right].$$

Introduce the estimates from Lemma 9 to reach

$$\left\| \Phi_L^\dagger \Phi(\mathbf{r}|_{L^c}) \right\|_2 \leq 24\alpha^{-1} \|\mathbf{r}\|_2. \quad (\text{B.8})$$

Substitute (B.7) and (B.8) into (B.6) to find

$$\|\mathbf{r} - \mathbf{b}^{\text{ideal}}\|_2 \leq 28\alpha^{-1} \|\mathbf{r}\|_2.$$

There is an additional wrinkle that occurs because the algorithm uses an iterative procedure to apply the pseudoinverse. As a result, the computed approximation \mathbf{b} may be slightly worse than the ideal approximation $\mathbf{b}^{\text{ideal}}$. By performing a sufficient number of iterations, we can ensure that

$$\|\mathbf{b} - \mathbf{b}^{\text{ideal}}\|_2 \leq 2\alpha^{-1} \|\mathbf{r}\|_2, \quad (\text{B.9})$$

whence

$$\|\mathbf{r} - \mathbf{b}\|_2 \leq 30\alpha^{-1} \|\mathbf{r}\|_2.$$

By choosing α to be a large enough constant, such as 200, we obtain the required estimate. \square

Finally, we prove that we need only perform a small number of iterations of Jacobi's algorithm to apply the pseudoinverse $\Phi_L^\dagger = (\Phi_L^* \Phi_L)^{-1} \Phi_L^*$ to a vector and still achieve sufficient precision.

Lemma 13. *To compute $\Phi_L^\dagger(\Phi \mathbf{r})$ with precision sufficient for Inequality (B.9) requires a constant number of iterations of Jacobi's algorithm. The cost of each iteration is equivalent to the cost of applying $\Phi_L^* \Phi_L$ to a vector.*

Proof. Define $\mathbf{x}^0 = \mathbf{0}$. Jacobi's algorithm defines a sequence by

$$\mathbf{x}^k = (\mathbf{I} - \Phi_L^* \Phi_L) \mathbf{x}^{k-1} + \Phi_L^* (\Phi \mathbf{r}).$$

The standard analysis of Jacobi's algorithm shows that

$$\|\mathbf{x}^k - \mathbf{b}^{\text{ideal}}\|_2 \leq \|\mathbf{I} - \Phi_L^* \Phi_L\|^k \|\mathbf{x}^0 - \mathbf{b}^{\text{ideal}}\|_2.$$

According to Lemma 10, the norm of the iteration matrix is at most 0.75. It is easy to see from the last proof that $\|\mathbf{b}^{\text{ideal}}\|_2 \leq 2 \|\mathbf{r}\|_2$. Therefore,

$$\|\mathbf{x}^k - \mathbf{b}^{\text{ideal}}\|_2 \leq 2\alpha^{-1} \|\mathbf{r}\|_2.$$

after at most $k = -\log \alpha / \log(0.75)$ iterations. \square

B.4. Pruning. The estimation step produces an approximation \mathbf{b} that is relatively close to \mathbf{r} and has a small number of nonzero components. The pruning step combines this approximation with the previous approximation and extracts the most significant components to obtain a very sparse new approximation to the input signal. We need to make sure that the new approximation reduces the residual substantially.

More formally, the current residual is $\mathbf{r} = \mathbf{f} - \mathbf{a}$ where \mathbf{a} is the current approximation to the signal. We add the vector \mathbf{b} to \mathbf{a} to obtain an overall approximation $\mathbf{c} = \mathbf{a} + \mathbf{b}$. The pruning step computes a new signal approximation $\mathbf{a}^{\text{new}} = \mathbf{c}_p$, the best p -term approximation to \mathbf{c} .

Lemma 14. *Suppose that the condition (Case 1) holds. The error in the pruned approximation \mathbf{c}_p satisfies*

$$\|\mathbf{f} - \mathbf{c}_p\|_2 \leq \frac{1}{2} \|\mathbf{f} - \mathbf{a}\|_2.$$

Proof. We make the following calculation.

$$\|\mathbf{f} - \mathbf{c}_p\|_2 \leq \|\mathbf{f} - \mathbf{c}\|_2 + \|\mathbf{c} - \mathbf{c}_p\|_2.$$

Since \mathbf{c}_p is the best approximation to \mathbf{c} using p terms, we increase the error by replacing it with the p -sparse vector \mathbf{f}_p . So

$$\begin{aligned} \|\mathbf{f} - \mathbf{c}_p\|_2 &\leq \|\mathbf{f} - \mathbf{c}\|_2 + \|\mathbf{c} - \mathbf{f}_p\|_2 \\ &\leq 2\|\mathbf{f} - \mathbf{c}\|_2 + \|\mathbf{f} - \mathbf{f}_p\|_2 \\ &= 2\|\mathbf{r} - \mathbf{b}\|_2 + \|\mathbf{f} - \mathbf{f}_p\|_2. \end{aligned} \tag{B.10}$$

Lemma 12 bounds the first term and Lemma 7 bounds the second term, hence

$$\|\mathbf{f} - \mathbf{c}_p\|_2 \leq 0.3\|\mathbf{r}\|_2 + \frac{1}{2\sqrt{7m}}\|\mathbf{f} - \mathbf{f}_m\|_1.$$

Since condition (Case 1) is in force and $(2\sqrt{7})^{-1} < 0.2$, we arrive at

$$\|\mathbf{f} - \mathbf{c}_p\|_2 \leq (0.3 + 0.2)\|\mathbf{f} - \mathbf{a}\|_2.$$

□

In summary, when condition (Case 1) obtains, an iteration of the algorithm reduces the ℓ_2 norm of the residual by a factor of two.

APPENDIX C. CASE 2: SMALL RESIDUAL

Suppose that (Case 2) is in force at the beginning of an iteration. This section shows that the iteration produces a new approximation to the signal that is not much worse than the current approximation.

C.1. Signal Properties. First, we establish a norm property of the residual $\mathbf{r} = \mathbf{f} - \mathbf{a}$ that is important in the subsequent analysis.

Lemma 15. *When condition (Case 2) is in force, then*

$$\|\mathbf{r}\|_1 \leq 6\|\mathbf{f} - \mathbf{f}_m\|_1.$$

Proof. First, note that

$$\begin{aligned} \|\mathbf{f} - \mathbf{a}\|_1 &\leq \|\mathbf{f} - \mathbf{f}_p\|_1 + \|\mathbf{f}_p - \mathbf{a}\|_1 \\ &\leq \|\mathbf{f} - \mathbf{f}_m\|_1 + \|\mathbf{f}_p - \mathbf{a}\|_1. \end{aligned} \tag{C.1}$$

Since the vector $(\mathbf{f}_p - \mathbf{a})$ has at most $2p$ nonzero components and $p = 8m$, it follows that

$$\begin{aligned} \|\mathbf{f}_p - \mathbf{a}\|_1 &\leq \sqrt{2p}\|\mathbf{f}_p - \mathbf{a}\|_2 \\ &\leq 4\sqrt{m}[\|\mathbf{f} - \mathbf{f}_p\|_2 + \|\mathbf{f} - \mathbf{a}\|_2]. \end{aligned}$$

According to Lemma 7,

$$\|\mathbf{f} - \mathbf{f}_p\|_2 \leq \frac{1}{2\sqrt{7m}}\|\mathbf{f} - \mathbf{f}_m\|_1.$$

Since (Case 2) holds,

$$\|\mathbf{f} - \mathbf{a}\|_2 \leq \frac{1}{\sqrt{m}}\|\mathbf{f} - \mathbf{f}_m\|_1.$$

Introducing the last three bounds into (C.1), we obtain

$$\|\mathbf{f} - \mathbf{a}\|_2 \leq \left[1 + \frac{4\sqrt{m}}{2\sqrt{7m}} + \frac{4\sqrt{m}}{\sqrt{m}} \right] \|\mathbf{f} - \mathbf{f}_m\|_1.$$

A numerical calculation shows that the bracket is no greater than 6, which completes the proof. \square

C.2. Identification. Under (Case 2), the identification step may fail utterly in its quest to find the largest components remaining in the residual. As we check in Section D, it still produces a list L of signal positions that satisfies

$$|L| = O(M\sqrt{\log M}).$$

This is the only property we require.

C.3. Estimation. We must ensure that the estimation step does not perform too badly. That is, we must show that it does not yield such inaccurate estimates for the signal positions in L that the subsequent approximation suffers. We can establish this fact using only the hypothesis that the residual is small, the bound on the list size, and the properties of the estimation matrix.

Lemma 16. *Suppose that condition (Case 2) is in effect. Then the estimation step produces an approximation \mathbf{b} of the residual \mathbf{r} that satisfies*

$$\|\mathbf{r} - \mathbf{b}\|_2 \leq 8.5 \|\mathbf{r}\|_2 + \frac{1}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1.$$

Proof. As in Lemma 12, the ideal approximation $\mathbf{b}^{\text{ideal}}$ satisfies

$$\|\mathbf{r} - \mathbf{b}^{\text{ideal}}\|_2 \leq \|\mathbf{r}|_{L^c}\|_2 + \|\Phi_L^\dagger \Phi(\mathbf{r}|_{L^c})\|_2. \quad (\text{C.2})$$

Since the list L is not too large, Lemma 10 yields $\|\Phi_L^\dagger\|_{2,2} \leq 2$. Apply this bound and use Lemma 11 to reach

$$\|\mathbf{r} - \mathbf{b}^{\text{ideal}}\|_2 \leq \|\mathbf{r}|_{L^c}\|_2 + 3 \left[\|\mathbf{r}|_{L^c}\|_2 + \frac{1}{\sqrt{M}} \|\mathbf{r}|_{L^c}\|_1 \right].$$

Drop the restriction to L^c , which increases all the norms on the right-hand side:

$$\|\mathbf{r} - \mathbf{b}^{\text{ideal}}\|_2 \leq \|\mathbf{r}\|_2 + 3 \left[\|\mathbf{r}\|_2 + \frac{1}{\sqrt{M}} \|\mathbf{r}\|_1 \right].$$

Since $\sqrt{M} > 4\sqrt{m}$, Lemma 15 and condition (Case 2) allow that

$$\frac{1}{\sqrt{M}} \|\mathbf{r}\|_1 \leq \frac{6}{4\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1 \leq 1.5 \|\mathbf{r}\|_2.$$

We conclude that

$$\|\mathbf{r} - \mathbf{b}^{\text{ideal}}\|_2 \leq 8.5 \|\mathbf{r}\|_2.$$

We are unable to compute $\mathbf{b}^{\text{ideal}}$ exactly with an iterative algorithm, but we may perform sufficient iterations so that the calculated approximation \mathbf{b} satisfies

$$\|\mathbf{b} - \mathbf{b}^{\text{ideal}}\|_2 \leq \frac{1}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1.$$

Apply the triangle inequality to complete the argument. \square

Using an argument very similar to Lemma 13, we can show that a constant number of iterations of Jacobi suffice to determine $\mathbf{b}^{\text{ideal}}$ within the required tolerance.

C.4. Pruning. As before, the pruning step takes the overall approximation $\mathbf{c} = (\mathbf{a} + \mathbf{b})$ and prunes it to its largest p terms \mathbf{c}_p . This vector \mathbf{c}_p will become the new approximation \mathbf{a}^{new} . We must check that this new approximation is not too poor.

Lemma 17. *Suppose that condition (Case 2) is in force. The error in the pruned approximation \mathbf{c}_p satisfies*

$$\|\mathbf{f} - \mathbf{c}_p\|_2 \leq \frac{20}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1.$$

Proof. The argument is similar to the proof of Lemma 14. In particular, we have the analog of (B.10):

$$\|\mathbf{f} - \mathbf{c}_p\|_2 \leq 2 \|\mathbf{r} - \mathbf{b}\|_2 + \|\mathbf{f} - \mathbf{f}_p\|_2.$$

We use Lemma 16 to bound the first term and Lemma 7 to bound the second term, whence

$$\|\mathbf{f} - \mathbf{c}_p\|_2 \leq 17 \|\mathbf{r}\|_2 + \frac{2}{\sqrt{m}} \|\mathbf{f} - \mathbf{f}_m\|_1 + \frac{1}{2\sqrt{7m}} \|\mathbf{f} - \mathbf{f}_m\|_1.$$

Since condition (Case 2) is in effect,

$$\|\mathbf{f} - \mathbf{c}_p\|_2 \leq \frac{1}{\sqrt{m}} \left[19 + \frac{1}{2\sqrt{7}} \right] \|\mathbf{f} - \mathbf{f}_m\|_1.$$

Conclude by noting that the bracket is less than 20. \square

APPENDIX D. IDENTIFICATION, IN DETAIL

The identification matrix $\mathbf{\Omega} = \mathbf{B} \otimes_{\mathbf{r}} \mathbf{A}$ is a complicated thing. We can best understand its behavior by studying its pieces separately.

The isolation matrix \mathbf{A} consists of $\log_2 M$ blocks $\mathbf{A}^{(j)}$, where $j = 1, 2, 4, 8, \dots, M/2$:

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}^{(1)} \\ \mathbf{A}^{(2)} \\ \vdots \\ \mathbf{A}^{(M/2)} \end{bmatrix}$$

The j th block is designed to identify coefficients from those significant bands that contain between j and $2j$ signal components. The band can be considered as a set of distinguished components that we want to find, and we need to ask how the identification matrix treats them.

Each block $\mathbf{A}^{(j)} = \mathbf{R}^{(j)} \otimes_{\mathbf{r}} \mathbf{S}^{(j)}$ where $\mathbf{S}^{(j)}$ is the sifting matrix and $\mathbf{R}^{(j)}$ is the noise reduction matrix. It is best to think about the action of the isolation matrix $\mathbf{A}^{(j)}$ in two phases.

- (1) First $\mathbf{S}^{(j)}$ takes an input signal and generates a collection of output signals of the same length by zeroing out different collections of components. The idea is that most of the distinguished components will appear in an output signal that contains no other distinguished component.
- (2) Then $\mathbf{R}^{(j)}$ takes each of these signals and generates a further collection of output signals by zeroing out additional subsets of components. The idea is that, in many of the output signals, a distinguished component will survive, but the ℓ_1 norm of the other components will be substantially reduced.

Afterward, the bit-test matrix \mathbf{B} forms the inner product between each of its rows and each of the numerous output signals. Whenever an output signal contains a distinguished component and a small amount of noise, the $\log d$ bit tests allow us to determine the location of the distinguished component correctly.

D.1. High-Level Analysis of Identification. The bit test process can always identify the largest component of a signal, provided that the ℓ_1 norm of the remaining components is not too large. The following lemma is trivial.

Lemma 18 (Bit Tests). *Suppose that $\mathbf{g} = \boldsymbol{\delta} + \boldsymbol{\nu}$, where $\|\boldsymbol{\delta}\|_0 = 1$ and $\|\boldsymbol{\nu}\|_1 \leq \frac{1}{2}\|\boldsymbol{\delta}\|_1$. If we apply bit tests to the signal \mathbf{g} , then we correctly identify the location of the spike $\boldsymbol{\delta}$.*

The ℓ_0 quasi-norm $\|\cdot\|_0$ counts the number of nonzero components in its argument.

Let us consider a fixed collection I of components in a signal \mathbf{g} , where $j \leq |I| < 2j$. The next result shows that $\mathbf{A}^{(j)}$ succeeds in generating a lot of output signals where a large proportion of the components in I are isolated from each other. Moreover, the ℓ_1 norm of the other components in these signals is small in comparison with the total norm of the signal. The number ρ in this lemma is a constant (depending only on α) that will be determined shortly.

Lemma 19. *Except with probability $O(d^{-1} \log m)$, the random isolation operator $\mathbf{A}^{(j)}$ satisfies the following property. Let \mathbf{g} be a signal, and let I be an arbitrary subset of $\{1, 2, \dots, d\}$ with $j \leq |I| < 2j$. For at least $(1 - \rho)|I|$ of the components $i \in I$, then operator $\mathbf{A}^{(j)}$ generates at least*

$$O(\sqrt{(M/j) \log M \log(d/j) \log d})$$

signals of the form $g_i \mathbf{e}_i + \boldsymbol{\nu}$ where

$$\|\boldsymbol{\nu}\|_1 \leq \frac{1}{4|I|} \sqrt{\frac{j}{M \log_2 M}} \|\mathbf{g}\|_1.$$

The proof of this result takes several long steps, so we postpone it to the sequel.

Lemma 20 (Significant Bands). *Suppose that B_s is a significant band. Then the size $|B_s|$ of the band satisfies*

$$\frac{s^2}{M \log_2 M} \leq |B_s| < 2s.$$

Proof. Fix a significant band B_s . Since the entries in B_s exceed $(2s)^{-1} \|\mathbf{r}\|_1$, Markov's inequality shows that the size of B_s is less than $2s$.

By definition of ‘‘significant,’’ we have a lower bound on the energy in the band, relation (B.5). It is also clear that this energy is achieved with the fewest coefficients provided that each coefficient has the maximal magnitude $\|\mathbf{r}\|_1/s$. That is,

$$|B_s| \frac{\|\mathbf{r}\|_1^2}{s^2} \geq \frac{\alpha^{-2}}{\log M} \|\mathbf{r}\|_2^2.$$

Rearranging this inequality, we have

$$|B_s| \geq \frac{\alpha^{-2} s^2}{\log M} \left(\frac{\|\mathbf{r}\|_2}{\|\mathbf{r}\|_1} \right)^2.$$

Lemma 8 provides a lower bound on the parenthesis. We conclude that

$$|B_s| \geq \frac{\alpha^{-2} s^2}{\log M} \left(\frac{\alpha}{\sqrt{M}} \right)^2,$$

as advertised. □

Lemma 21 (Identifying Significant Bands). *Let B_s be a significant band for which $j \leq |B_s| < 2j$. The identification process finds at least $(1 - \rho)|B_s|$ of the components in the band at least $O(\sqrt{(M/j) \log M \log(d/j) \log d})$ times each.*

Proof. Suppose that $j \leq |B_s| < 2j$. Let i be one of the components in B_s for which the identification operator $\mathbf{A}^{(j)}$ generates a signal of the form $r_i \mathbf{e}_i + \boldsymbol{\nu}$ with

$$\|\boldsymbol{\nu}\|_1 \leq \frac{1}{4|B_s|} \sqrt{\frac{j}{M \log_2 M}} \|\mathbf{r}\|_1.$$

Using the fact $j \leq |B_s|$ and the lower bound on the band size from Lemma 20, we have

$$\begin{aligned} \|\boldsymbol{\nu}\|_1 &\leq \frac{1}{4\sqrt{|B_s|} M \log_2 M} \|\mathbf{r}\|_1 \\ &\leq \frac{1}{4s} \|\mathbf{r}\|_1. \end{aligned}$$

By definition, the components in B_s have magnitude at least $(2s)^{-1} \|\mathbf{r}\|_1$. Lemma 18 shows that the bit tests succeed in identifying the location i from the signal $r_i \mathbf{e}_i + \boldsymbol{\nu}$.

We conclude that the identification process succeeds in finding a $(1 - \rho)$ -fraction of the indices in B_s , and it lists each of these at least $O(\sqrt{(M/j) \log M} \log(d/j) \log d)$ times. \square

Lemma 22 (Energy Loss). *Select $\rho = (4\alpha^2 - 7)^{-1}$. If the identification process misses no more than a ρ -proportion of the components in each significant band, the total ℓ_2 norm of the lost components is at most $\alpha^{-1} \|\mathbf{r}\|_2$.*

Proof. Fix a band B_s . Using basic calculus, one can check that the worst-case scenario occurs when (i) we miss fully $\rho|B_s|$ components from this band, (ii) the missing components have the maximum magnitude $s^{-1} \|\mathbf{r}\|_1$, and (iii) the identified components have the minimum magnitude $(2s)^{-1} \|\mathbf{r}\|_1$. In this case, the ratio of the energy (i.e., sums of squares) in the missing components to the total energy in the band satisfies

$$\frac{\rho|B_s| s^{-2} \|\mathbf{r}\|_1^2}{\rho|B_s| s^{-2} \|\mathbf{r}\|_1^2 + (1 - \rho)|B_s| (2s)^{-2} \|\mathbf{r}\|_1^2} = \frac{4\rho}{3\rho + 1}.$$

Altogether, the energy in the significant bands is at least $(1 - \alpha^{-2}) \|\mathbf{r}\|_2^2$. Therefore, the total energy loss over all significant bands is at most

$$\frac{4\rho(1 - \alpha^{-2})}{3\rho + 1} \|\mathbf{r}\|_2^2.$$

If we select $\rho = (4\alpha^2 - 7)^{-1}$, then it follows that the total energy loss is at most $\alpha^{-2} \|\mathbf{r}\|_2^2$. \square

Lemma 23 (List Reduction). *The final list of identified components contains $O(M\sqrt{\log M})$ items.*

Note that this result does not depend on the “success” of the identification process in finding energy in the residual.

Proof. The matrix $\mathbf{A}^{(j)}$ has a total of $O(M \log(M) \log(d/j) \log d)$ rows. The bit tests identify one signal location per row of the matrix, so the identification process makes an initial list of $O(M \log(M) \log(d/j) \log d)$ signal locations.

We only retain list items that appear $\Omega(\sqrt{(M/j) \log M} \log(d/j) \log d)$ times. Therefore, the total number of items that remain is $O(\sqrt{jM \log M})$. Sum over $j = 1, 2, 4, \dots, M/2$ to complete the proof. \square

D.2. Sifting. Given a signal \mathbf{g} and a set I of distinguished components with $j \leq |I| < 2j$, we establish that the sifting matrix $\mathbf{S}^{(j)}$ succeeds in isolating most of these components from each other, while reducing the norm of the other components somewhat.

Recall that the sifting matrix $\mathbf{S}^{(j)}$ can be viewed as performing T_j independent trials, where $T_j = O(\log(d/j))$. The matrix has a hierarchical structure:

$$\mathbf{S}^{(j)} = \begin{bmatrix} \mathbf{S}_1^{(j)} \\ \mathbf{S}_2^{(j)} \\ \vdots \\ \mathbf{S}_{T_j}^{(j)} \end{bmatrix}.$$

Each submatrix $\mathbf{S}_t^{(j)}$ consists of $N = O(j)$ rows, so the whole sifting matrix consists of $O(j \log(d/j))$ rows.

We can think about the action of one submatrix $\mathbf{S}_t^{(j)}$ as

$$\mathbf{S}_t^{(j)} : \mathbf{g} \mapsto [\mathbf{h}^1 \quad \mathbf{h}^2 \quad \dots \quad \mathbf{h}^N].$$

That is, each input signal is mapped to a collection of output signals. Each component of the input signal is randomly assigned to one output signal. The random variables that describe where the components are sent only need to be $O(j)$ -wise independent. This important point follows from the fact that we only consider $O(j)$ components at a time.

The first result shows that one trial of sifting is very likely to isolate all but a constant proportion of the distinguished indices.

Lemma 24 (Sifting: One Trial). *Let \mathbf{g} be a signal, and let $I \subset \{1, 2, \dots, d\}$ with $j \leq |I| < 2j$. Write $k = |I|$. Suppose we apply the random operator $\mathbf{S}_t^{(j)}$ to \mathbf{g} . Except with probability $e^{-1.7 - \rho k/5}$, for at least $(1 - \rho)k$ of the indices $i \in I$, there is an output signal \mathbf{h} of the form*

- (1) $\mathbf{h} = g_i \mathbf{e}_i + \boldsymbol{\nu}$
- (2) $\|\boldsymbol{\nu}\|_1 \leq \frac{2}{\rho k} \|\mathbf{g}\|_1$.

Proof. We can think of the sifting operator as assigning each of the k distinguished positions (balls) to one of the N output signals (bins) uniformly at random. We hope that the balls are isolated from one another. We will see that if the number of bins satisfies $N \geq \max\{10k\rho^{-1}, 850\rho^{-1}\}$, then the result holds.

For $n = 1, 2, \dots, N$, let X_n be the indicator variable for the event that the n th bin is empty, and write $X = \sum X_n$ for the total number of empty bins. The symbols μ and σ^2 will denote the expectation and variance of X . To understand large deviations of X requires some effort because the set of indicators $\{X_n\}$ is not stochastically independent. Nevertheless, X satisfies a rather strong tail bound.

Fact 25 (Theorem 6, [Jan94]).

$$\mathbb{P}\{X \geq \mathbb{E}X + a\} \leq \exp\left\{-\sigma^2 \left(\left(1 + \frac{a}{\sigma^2}\right) \log\left(1 + \frac{a}{\sigma^2}\right) - \frac{a}{\sigma^2}\right)\right\}.$$

This result is based on the surprising fact, due to Vatutin and Mikhailov [VM82], that X can be expressed as a sum of independent indicators.

The content of our argument is to develop explicit bounds on the expectation and variance of X , which will allow us to apply Janson's result. By calculating the means and covariances of the variables X_n , we determine that

$$\mu = N \left(1 - \frac{1}{N}\right)^k$$

and that

$$\sigma^2 = N \left(1 - \frac{1}{N}\right)^k + N(N-1) \left(1 - \frac{2}{N}\right)^k - N^2 \left(1 - \frac{1}{N}\right)^{2k}.$$

It is straightforward to develop an upper bound on the mean by deploying Taylor series:

$$\mu < N \left(1 - \frac{k}{N} + \frac{k^2}{2N^2}\right) = (N - k) + \frac{k^2}{2N}.$$

The variance takes more work. First, regroup terms and factor to obtain

$$\sigma^2 = N \left(1 - \frac{1}{N}\right)^k \left[1 - \left(\frac{1 - 2/N}{1 - 1/N}\right)^k\right] + N^2 \left(1 - \frac{2}{N}\right)^k \left[1 - \left(\frac{(1 - 1/N)^2}{1 - 2/N}\right)^k\right].$$

We apply Bernoulli's inequality $(1+x)^k \geq 1+kx$, which is valid for $x \geq -1$, to obtain the bounds

$$\begin{aligned} 1 - \left(\frac{1 - 2/N}{1 - 1/N}\right)^k &\leq \frac{k}{N(1 - 1/N)} \\ 1 - \left(\frac{(1 - 1/N)^2}{1 - 2/N}\right)^k &\leq \frac{-k}{N^2(1 - 2/N)}. \end{aligned}$$

Combining these estimates, we reach

$$\sigma^2 \leq k \left[\left(1 - \frac{1}{N}\right)^{k-1} - \left(1 - \frac{2}{N}\right)^{k-1} \right].$$

The function $h(x) = (1-x)^{k-1} - (1-2x)^{k-1}$ is concave for $x \in [0, 0.5]$, so the tangent line at $x = 0$ yields an overestimate. Since $h(0) = 0$ and $h'(0) = k - 1$, we have

$$\sigma^2 \leq k \cdot h(1/N) \leq \frac{k(k-1)}{N} < \frac{k^2}{N}$$

provided that $N \geq 2$.

Depending on the size of k , we need to choose a different number N of bins to obtain the required probabilities. First, assume that $0.2\rho k > 1.7$. In this case, we select $N \geq 10k/\rho$, which yields the following estimates on the mean and variance of X :

$$\begin{aligned} \mu &\leq (N - k) + 0.05\rho k \\ \sigma^2 &< 0.1\rho k. \end{aligned}$$

We invoke Fact 25 with the value $a = 0.2\rho k$ to reach

$$\mathbb{P}\{X > (N - k) + 0.25\rho k\} < e^{-0.4\rho k} < e^{-1.7 - 0.2\rho k} \quad (\text{D.1})$$

using $0.2\rho k > 1.7$.

This estimate allows us to bound the number Y of balls that fail to be isolated. It takes at least $(N - X)$ balls to fill the nonempty bins. The remaining $(k - (N - X))$ balls can be placed in no more than $(k - (N - X))$ bins, where they will result in no more than $Y = 2(k - (N - X)) = 2(X - (N - k))$ collisions. Using the deviation bound (D.1), we conclude that

$$\mathbb{P}\left\{Y > \frac{\rho k}{2}\right\} < e^{-1.7 - 0.2\rho k}.$$

Second, we assume that k is small. Precisely, consider the case where $0.2\rho k \leq 1.7$. Now, select $N \geq 100k^2$. The mean and variance of X satisfy

$$\begin{aligned} \mu &\leq (N - k) + 0.005 \\ \sigma^2 &\leq 0.01 \end{aligned}$$

Apply Fact 25 with the value $a = 0.995$ to reach

$$\mathbb{P}\{X \geq (N - k) + 1\} < e^{-3.6} < e^{-1.7 - 0.2\rho k}$$

using $0.2\rho k \leq 1.7$. When $X < (N - k) + 1$, the maximum number of bins are empty, and so all k of the balls are isolated. Furthermore, we observe that the number N of bins required here satisfies $N \leq 850\rho^{-1}$.

Finally, we need to argue that few of the output signals have a lot of noise. Let u_n denote the ℓ_1 norm of the n th output signal. Since each position in the input signal \mathbf{g} is assigned to exactly one output signal, $\sum_n u_n = \|\mathbf{g}\|_1$. By Markov's inequality,

$$\#\left\{n : u_n \geq \frac{2}{\rho k} \|\mathbf{g}\|_1\right\} \leq \frac{\rho k}{2\|\mathbf{g}\|_1} \sum_n u_n = \frac{\rho k}{2}.$$

In particular, no more than $\rho k/2$ of the isolated balls can appear in a bin whose ℓ_1 norm exceeds $(2/\rho k) \|\mathbf{g}\|_1$. Therefore, the total number of positions in I lost to collisions or noise is at most ρk except with probability $e^{-1.7 - 0.2\rho k}$. \square

The failure probability for one trial is not small enough to take a union bound over all possible sets I . We perform $T_j = O(\log(d/j))$ repeated trials to drive down the failure probability.

Lemma 26 (Sifting: All Trials). *Except with probability $\exp(-j \log d)$, the sifting operator $\mathcal{S}^{(j)}$ has the following property. Let \mathbf{g} be an arbitrary signal, and let $I \subset \{1, 2, \dots, d\}$ satisfy $j \leq |I| < 2j$. For some set of $(1 - \rho)|I|$ indices $i \in I$, there are at least $0.5T_j$ output signals \mathbf{h} of the form*

- (1) $\mathbf{h} = g_i \mathbf{e}_i + \boldsymbol{\nu}$
- (2) $\|\boldsymbol{\nu}\|_1 \leq \frac{2}{\rho k} \|\mathbf{g}\|_1$.

Proof. Let \mathbf{g} be a signal, and fix a set $I \subset \{1, 2, \dots, d\}$ that contains k or more indices. Lemma 24 shows that each isolation trial succeeds for at least $(1 - \rho)|I|$ of the distinguished indices, except with probability $p = e^{-1.7 - 0.2\rho k}$. We repeat this experiment T_j times. Let X_t be the indicator variable for the event that trial t fails, so $\mathbb{E}X_t \leq p$. Then the random variable $X = \sum X_t$ counts the total number of trials in which ρk balls fail to be isolated, and its mean satisfies $\mu \leq pT_j$. Chernoff's bound shows that

$$\mathbb{P}\{X > 0.5T_j\} < \left[\frac{e}{0.5T_j/(pT_j)}\right]^{T_j} < e^{-0.2\rho k T_j}$$

Choose $T_j = 15\rho^{-1} \log(ed/j)$, and use the fact that $k \geq j$ to obtain

$$\mathbb{P}\{X > T_j\} < e^{-3j \log(ed/j)}.$$

Next, we must count the total number of subsets of $\{1, 2, \dots, d\}$ whose size is between j and $(2j - 1)$. It is well known that

$$\binom{d}{r} \leq e^{r \log(ed/r)}.$$

Therefore,

$$\sum_{r=j}^{2j-1} \binom{d}{r} \leq \sum_{r=j}^{2j-1} e^{r \log(ed/r)} \leq \int_j^{2j} e^{x \log(ed/x)} dx,$$

which uses the fact that $x \mapsto x \log(ed/x)$ is monotonically increasing. A crude estimate of the integral suffices:

$$\int_j^{2j} e^{x \log(ed/x)} dx \leq e^{\log(ed/j)} \int_j^{2j} e^x dx \leq e^{2j \log(ed/j)}.$$

Finally, we take a union bound over all sets I with size between j and $(2j - 1)$ to obtain a failure probability of

$$e^{2j \log(ed/j)} \cdot e^{-3j \log(ed/j)} = e^{-j \log(ed/j)}.$$

In other words, for every such I , the sifting matrix $\mathbf{S}^{(j)}$ isolates at least $(1 - \rho) |I|$ of the distinguished indices in at least half the trials. \square

D.3. Noise Reduction. The noise reduction matrices $\mathbf{R}^{(j)}$ are designed to take the output signals containing isolated components and further attenuate the noise in those signals. Once again, the matrix has a hierarchical structure:

$$\mathbf{R}^{(j)} = \begin{bmatrix} \mathbf{R}_1^{(j)} \\ \mathbf{R}_2^{(j)} \\ \vdots \\ \mathbf{R}_{U_j}^{(j)} \end{bmatrix}$$

Let us define $r = \sqrt{M/j}$ and $C = 80\rho^{-1}\sqrt{\log_2 M}$. Each of the submatrices consists of at least Cr rows, and the number U_j of submatrices is at least $Cr \log d$. Therefore, the total number of rows in the noise reduction matrix $\mathbf{R}^{(j)}$ is $O((M/j) \log(M) \log d)$.

It is most convenient to think of the action of the submatrix $\mathbf{R}_t^{(j)}$ as

$$\mathbf{R}_t^{(j)} : \mathbf{g} \mapsto [\mathbf{h}^1 \quad \mathbf{h}^2 \quad \dots \quad \mathbf{h}^N].$$

where $N = O(Cr)$ is the number of rows in the submatrix. Each submatrix assigns the positions of the input signal randomly to one output signal. The proof shows that this assignment only needs to be pairwise independent. Each of the submatrices, however, must be fully independent from the others.

We establish that, if \mathbf{g} contains a single distinguished component (a spike) plus noise, then a large number of the output signals contain that spike along with a reduced amount of noise.

Lemma 27 (Noise Reduction). *Except with probability d^{-1} , the noise reduction matrix $\mathbf{R}^{(j)}$ has the following property. Let \mathbf{g} be an input signal that satisfies*

- (1) $\mathbf{g} = \boldsymbol{\delta} + \boldsymbol{\nu}$
- (2) $\|\boldsymbol{\delta}\|_0 = 1$
- (3) $\text{supp}(\boldsymbol{\delta}) \cap \text{supp}(\boldsymbol{\nu}) = \emptyset$

Then there are at least $0.5Cr \log d$ output signals \mathbf{h} of the form $\mathbf{h} = \boldsymbol{\delta} + \boldsymbol{\mu}$ where

$$\|\boldsymbol{\mu}\|_1 \leq \frac{10}{Cr} \|\boldsymbol{\nu}\|_1 = \frac{\sqrt{j}}{8\rho^{-1}\sqrt{M \log_2 M}} \|\boldsymbol{\nu}\|_1.$$

Proof. Let \mathbf{g} be a signal of the form $\mathbf{g} = \boldsymbol{\delta} + \boldsymbol{\nu}$, where $\boldsymbol{\delta}$ is a spike at position i . Consider the submatrix \mathbf{X} of $\mathbf{R}^{(j)}$ constructed by extracting the rows of $\mathbf{R}^{(j)}$ where the index i appears and then removing the i th column. This submatrix contains exactly $Cr \log d$ rows and $(d - 1)$ columns. Let $\boldsymbol{\nu}'$ be the vector $\boldsymbol{\nu}$ without its i th component (which equals zero by hypothesis). Note that $\boldsymbol{\nu}'$ has the same ℓ_1 norm as $\boldsymbol{\nu}$.

Let \mathbf{x} be a column of \mathbf{X} . The entries of \mathbf{x} are independent binary random variables with mutual expectation $(Cr)^{-1}$ because each one comes from a different submatrix. Therefore,

$$\mathbb{E} \|\mathbf{x}\|_1 = \frac{Cr \log d}{Cr} = \log d.$$

Chernoff's bound shows that

$$\mathbb{P} \{ \|\mathbf{x}\|_1 > 5 \log d \} \leq \left[\frac{e^4}{5^5} \right]^{\log d} < d^{-3}.$$

Applying the union bound over all $(d-1)$ columns of \mathbf{X} ,

$$\mathbb{P} \left\{ \|\mathbf{X}\|_{1,1} > 5 \log d \right\} \leq d^{-2}.$$

Therefore, the number of output signals in which position i appears and where the noise is large satisfies

$$\begin{aligned} \#\{n : |(\mathbf{X}\boldsymbol{\nu}')_n| > \frac{10\|\boldsymbol{\nu}'\|_1}{Cr}\} &\leq 0.1Cr \frac{\|\mathbf{X}\boldsymbol{\nu}'\|_1}{\|\boldsymbol{\nu}'\|_1} \\ &\leq 0.5Cr \log d. \end{aligned}$$

Since position i appears in exactly $Cr \log d$ output signals, we discover that the number of output signals where position i appears and the noise is less than $(10/Cr)\|\boldsymbol{\nu}'\|_1$ is at least $0.5Cr \log d$.

So far, we have only established the result for a single spike location i . To complete the proof, we perform a union bound over the d possible locations for the spike to obtain a final failure probability of d^{-1} . \square

Combine Lemma 26 and Lemma 27 to obtain the announced Lemma 19.

APPENDIX E. THE ESTIMATION MATRIX

In this section, we show that the estimation matrix Φ has the properties described in Lemmas 10 and 11. The first result appears in the work of Rudelson–Vershynin [RV06]. They shows that a random row-submatrix of the discrete Fourier transform (DFT) matrix has the Restricted Isometry Property (RIP).

Lemma 28. *Let $\lambda = O(\log^4 d)$. Choose $q \geq \lambda L$, and draw q rows independently at random from the $d \times d$ discrete Fourier transform matrix, and rescale them by $q^{-1/2}$. With fixed constant probability (e.g., 0.99), the resulting matrix Φ has the property that each L -column submatrix \mathbf{A} satisfies*

$$\frac{1}{2} \|\mathbf{x}\|_2 \leq \|\mathbf{A}\mathbf{x}\|_2 \leq \frac{3}{2} \|\mathbf{x}\|_2.$$

The upper estimate in the RIP is essentially equivalent to the following result, which plays an essential role in our analysis.

Lemma 29. *Suppose that Φ is a matrix whose L -column submatrices satisfy the upper estimate of Lemma 28. Then the following bound holds for every vector \mathbf{x} .*

$$\|\Phi\mathbf{x}\|_2 \leq \frac{3}{2} \left[\|\mathbf{x}\|_2 + \frac{1}{\sqrt{L}} \|\mathbf{x}\|_1 \right].$$

Since $M \leq L$, we also have the weaker estimate

$$\|\Phi\mathbf{x}\|_2 \leq \frac{3}{2} \left[\|\mathbf{x}\|_2 + \frac{1}{\sqrt{M}} \|\mathbf{x}\|_1 \right].$$

Proof. Define the matrix norm

$$\|\Phi\| = \max_{\substack{\|\mathbf{x}\|_2 \leq 1 \\ \|\mathbf{x}\|_0 \leq L}} \|\Phi\mathbf{x}\|_2$$

and define the vector norm

$$\|\mathbf{x}\|_K = \|\mathbf{x}\|_2 + \frac{1}{\sqrt{L}} \|\mathbf{x}\|_1.$$

We establish the result by proving that

$$\|\mathbf{A}\|_{K \rightarrow 2} \leq \|\Phi\| \quad (\text{E.1})$$

where the operator norm

$$\|\Phi\|_{K \rightarrow 2} = \max_{\|\mathbf{x}\|_K \leq 1} \|\Phi \mathbf{x}\|_2.$$

Afterward, we invoke Lemma 28 to see that $\|\Phi\| \leq 3/2$.

Let B_q denote the unit ball in ℓ_q^d , and let B_q^I be the unit ball in $\ell_q(I)$ where $I \subset \{1, 2, \dots, d\}$. Define the convex body

$$S = \text{conv} \left\{ \bigcup_{|I| \leq L} B_2^I \right\}.$$

Observe that the matrix norm $\|\cdot\| = \|\cdot\|_{S \rightarrow 2}$. Therefore, to verify (E.1), it suffices to show that $K \subset S$. We do so by establishing the reverse inclusion for the polars, $S^\circ \subset K^\circ$.

The norm with unit ball S° is

$$\|\mathbf{y}\|_{S^\circ} = \max_{|I| \leq L} \|\mathbf{y}|_I\|_2.$$

Choose a vector \mathbf{y} in the unit ball S° , and let I be a set of L coordinates of \mathbf{y} with greatest magnitudes. Then

$$|y_i| \leq \frac{1}{\sqrt{L}} \quad \text{for each } i \in I^c.$$

Otherwise, $|y_i| > \frac{1}{\sqrt{L}}$ for each $i \in I$, which would imply $\|\mathbf{y}\|_{S^\circ} \geq \|\mathbf{y}|_I\|_2 > 1$. To proceed, we decompose

$$\mathbf{y} = \mathbf{y}|_I + \mathbf{y}|_{I^c}.$$

It follows from this discussion that $\|\mathbf{y}|_I\|_2 \leq 1$ and that $\|\mathbf{y}|_{I^c}\|_\infty \leq \frac{1}{\sqrt{L}}$. In other terms,

$$\mathbf{y} \in B_2 + \frac{1}{\sqrt{L}} B_\infty.$$

But this set is the unit ball of K° because

$$\begin{aligned} \|\mathbf{x}\|_K &= \|\mathbf{x}\|_2 + \frac{1}{\sqrt{L}} \|\mathbf{x}\|_1 \\ &= \|\mathbf{x}\|_{B_2^\circ} + \|\mathbf{x}\|_{\sqrt{L} B_\infty^\circ} \\ &= \|\mathbf{x}\|_{(B_2 + \frac{1}{\sqrt{L}} B_\infty)^\circ} \end{aligned}$$

and because polarity is an involution on norm balls. \square