Non-adaptive Group Testing: Explicit bounds and novel algorithms

Chun Lam Chan*, Sidharth Jaggi*, Venkatesh Saligrama+, and Samar Agnihotri*
The Chinese University of Hong Kong*, Boston University+

Abstract

We present computationally efficient and provably correct algorithms with near-optimal sample-complexity for noisy non-adaptive group testing. Group testing involves grouping arbitrary subsets of items into pools. Each pool is then tested to identify the defective items, which are usually assumed to be sparse. We consider random non-adaptive pooling where pools are selected randomly and independently of the test outcomes. Our noisy scenario accounts for both false negatives and false positives for the test outcomes. Inspired by compressive sensing algorithms we introduce three computationally efficient algorithms for group testing, namely, Combinatorial Orthogonal Matching Pursuit (COMP), Combinatorial Basis Pursuit (CBP), and CBP via Linear Programming (CBP-LP) decoding. The first and third of these algorithms have several flavours, dealing separately with the noiseless and noisy measurement scenarios. We derive explicit sample-complexity bounds—with all constants made explicit—for these algorithms as a function of the desired error probability; the noise parameters; the number of items; and the size of the defective set (or an upper bound on it). We also derive lower bounds for sample complexity based on Fano's inequality and show that our upper and lower bounds are equal up to a constant factor.

I. Introduction

The goal of *group testing* is to identify a small unknown subset $\mathcal{D}$ of defective items embedded in a much larger set $\mathcal{N}$ (usually in the setting where $d = |\mathcal{D}|$ is much smaller than $n = |\mathcal{N}|$, *i.e.*, d is $o(n)$). This problem was first considered by Dorfman [1] in scenarios where multiple items in a group can be simultaneously tested, with a binary output depending on whether or not a "defective" item is present in the group being tested. In general, the goal of group testing algorithms is to identify the defective set with as few measurements as possible. As demonstrated in [1] and later work (see [2]), with judicious grouping and testing, far fewer than the trivial upper bound of n tests may be required to identify the set of defective items.

We consider *non-adaptive group testing* in this paper. In non-adaptive group testing, the set of items being tested in each test is required to be independent of the outcome of every other test [2]. This restriction is often useful in practice, since this enables parallelization of the testing process. It also allows for an automated testing process. In contrast, the procedures and hardware required for *adaptive* group testing may be significantly more complex.

In this paper we describe computationally efficient algorithms with near-optimal performance for noiseless and noisy non-adaptive group testing problems. We describe the different aspects of the paper in some detail next.

“Noisy” measurements: In addition to the *noiseless* group-testing problem, we consider the “noisy” variant of the problem. In this noisy variant the result of each test may differ from the true result (in an independent and identically distributed manner) with a certain pre-specified probability q. This leads to both false positives and negatives in the test outcomes. Much of the existing work either considers one-sided noise, namely false positives [3] but no false negatives or a “worst-case” noise [4] wherein the number of false positives and negatives are assumed bounded.[1] Since the measurements are noisy, the problem of estimating the set of defective items is more challenging, and is known to require more tests.[2]

Computationally efficient and near-optimal algorithms: Most algorithms in the literature focus on optimizing the number of measurements required – in some cases, this leads to algorithms that may not be computationally efficient to implement (for *e.g.* [3]). In this paper we present algorithms that are not only computationally efficient but are also near-optimal in the number of measurements required.

We derive lower bounds on group-testing algorithms based on information theoretic analysis. For the upper bounds we analyze three different types of algorithms. These algorithms are related to those described in the compressive sensing literature (see Section I-A).

The first two algorithms are based on Basis Pursuit and Orthogonal Matching Pursuit and we call our algorithms Combinatorial Basis Pursuit (CBP), and Combinatorial Orthogonal Matching Pursuit (COMP). Both CBP and COMP are not new and have also been previously considered in the group-testing literature (under different names) for both noiseless and noisy scenarios (see, for instance [7]). Our contribution here is a tighter analysis in comparison to the previous literature for COMP. We present a novel analysis based on the coupon-collector problem for CBP. Our sample-complexity bounds here are explicit and all of the constants involved are made precise.

Our third algorithm is related to linear programming relaxations used in the compressive sensing literature. In compressive sensing the ℓ_0 norm minimization is relaxed to an ℓ_1 norm minimization. In the noise-free case this relaxation results in a linear program since the measurements are linear. In contrast, in group testing, the measurements are non-linear and boolean. In the noise-free case the measurements take the value one if *some* defective item is in the pool and zero if no defective item is part of the pool. Furthermore, noise in the group testing scenario is also boolean unlike additive noises in compressive sensing. For these reasons we also need to relax our boolean measurement equations. We do so by using a novel combination of inequality and positivity constraints. Our LP formulation and analysis is related to error-correction [8], where, one uses a “minimum distance” decoding criteria based on perturbation analysis. The idea is to decode to a vector pair consisting defective items, $\mathbf{x}$, and the

[1] For instance [4] considers group-testing algorithms that are resilient to *all* noise patterns wherein at most a fraction q of the results differ from their true values, rather than the probabilistic guarantee we give against *most* fraction-q errors. This is analogous to the difference between combinatorial coding-theoretic error-correcting codes (for instance Gilbert-Varshamov codes [5]) and probabilistic information-theoretic codes (for instance [6]). In this work we concern ourselves only with the latter, though it is possible that our techniques can also be used to analyzed the former.

[2] **We wish to highlight the difference between *noise* and *errors*. We use the former term to refer to noise in the outcomes of the group-test, regardless of the group-testing algorithm used. The latter term is used to refer to the error due to the estimation process of the group-testing algorithm.**

error vector, η such that the error-vector η is as "small" as possible. We call this algorithm the Noisy Combinatorial Basis Pursuit via LP decoding (NCBP-LP). Using standard concentration results we show that the solution to our LP decoding algorithm recovers the true defective items with high probability. Furthermore, we achieve near-optimal performance in the sense that our sample complexity for NCBP-LP match the lower bounds within a constant factor, where the constant is independent of the number of items n and the defective set size d (but may depend on the noise parameter q, and the error probability ϵ). Based on this analysis, we can directly derive the performance of two other LP-based decoding algorithms. In particular CBP-LP considers the noiseless measurement scenario, and NCBP-SLP considers the noisy measurement scenario, but *only* uses constraints corresponding to positive test outcomes (whether true or false positives). It is perhaps interesting that even though NCBP-SLP only considers a subset of the available information, it nonetheless has performance comparable to NCBP-LP.[3]

"Small-error" probability ϵ: Existing work has considered both deterministic and random pooling designs [2]. In this context both deterministic and probabilistic sample complexity bounds for the number of measurements T that lead to exact identification of the defective items have been derived. There is also existing work on characterizing sample-complexity bounds for the average case scenario (see [3]). These sample complexity bounds are usually asymptotic in nature and describe the scaling of the number of items n with respect to the number of defectives d to ensure that the error probability approaches zero. To gain new insights into the constants involved in the sample-complexity bounds we admit a small but fixed error probability, ϵ. With this new perspective we can derive upper and lower bounds that hold not only in an order-wise sense but also where the constants involved in these order-wise bounds can be made explicit.

Explicit Sample Complexity Bounds: Our sample complexity bounds are of the form $T \geq \beta(D, \epsilon) d \log(n)$. The function $\beta(q, \epsilon)$ is an explicitly computed function of the noise parameter q and admissible error probability ϵ. In the literature, order-optimal upper and lower bounds on the number of tests required are known for the problems we consider (for instance [3], [9]). In both the noiseless and noisy variants, the number of measurements required to identify the set of defective items is known to be $T = \Theta(d \log(n))$ – here $n = |\mathcal{N}|$ is the total number of items and $d = |\mathcal{D}|$ is the size of the defective subset. In fact, if only D, an upper bound on d, is known, then $T = \Theta(D \log(n))$ measurements are also known to be necessary and sufficient. In our algorithms we explicitly demonstrate that we require only a knowledge of D rather than the exact value of d. Furthermore, in the noisy variant, we show that the number of tests required is in general a constant factor larger than in the noiseless case (where this constant β is independent of both n and d, but may depend on the noise parameter q and the allowable *error-probability* ϵ of the algorithm).

This paper is organized as follows. In Section II, we introduce the model and corresponding notation, and describe the algorithms analyzed in this work. In Section III, we describe the main results of this work. Sections IV and V contain the analysis respectively of our information-theoretic lower bounds, and of the group-testing algorithms

[3]In fact, our analysis actually indicates superior performance for NCBP-SLP compared to NCBP-LP, but this is an artifact of the fact that we did not truly optimize over all internal parameters in our proofs.

considered.

A. Compressive Sensing

Compressive sensing has received significant attention over the last decade. We describe the version most related to the topic of this paper [10], [11]. This version considers the following problem. Let $\mathbf{x}$ be an *exactly d-sparse* vector in $\mathbb{R}^n$, *i.e.*, a vector with at most d non-zero components (in general in the situations of interest $d = o(n)$).[4]

Let $\mathbf{z}$ corresponds to a *noise vector* added to the measurement $M\mathbf{x}$. One is given a set of "compressed noisy measurements" of $\mathbf{x}$ as

$$\mathbf{y} = M\mathbf{x} + \mathbf{z} \tag{1}$$

$$||\mathbf{z}||_2 \leq c_2 \tag{2}$$

$$||\mathbf{x}||_0 \leq d \tag{3}$$

Here the constraint (2) corresponds to a guarantee that the noise is not "too large", and the (non-linear) constraint (3) corresponds to the prior knowledge that $\mathbf{x}$ is d-sparse. The $T \times n$ matrix M is designed by choosing each entry i.i.d. from a suitable probability distribution (for instance, the set of zero-mean, $1/n$ variance Gaussian random variables). The decoder must use the resulting *noisy measurement vector* $\mathbf{y} \in \mathbb{R}^T$ and the matrix M to computationally efficiently estimate the underlying vector $\mathbf{x}$. The challenge is to do so with as few measurements as possible, *i.e.*, with the number of rows T of M being as small as possible.

1) Orthogonal Matching Pursuit: We note that it is enough for the decoder to computationally efficiently estimate the *support $\mathcal{D}$ of* $\mathbf{x}$, the set of indices on which $\mathbf{x}$ is non-zero, correctly. This is because the decoder can then estimate $\mathbf{x}$ as $(M_{\mathcal{D}}^t M_{\mathcal{D}})^{-1} M_{\mathcal{D}}^t \mathbf{y}$, which is the minimum mean-square error estimate of $\mathbf{x}$. (Here $M_{\mathcal{D}}$ equals the $T \times d$ sub-matrix of M whose columns numbers correspond to the indices in $\mathcal{D}$, and T is a design parameter chosen to optimize performance.)

One popular method of efficient estimation of $\mathcal{D}$ is that of *Orthogonal matching pursuit* (OMP) [12]. The intuition is that if the columns of the matrix M are "almost orthogonal" (every pair of columns have "relatively small" dot-product) then decoding can proceed in a greedy manner. In particular, the OMP algorithm computes the dot-product between $\mathbf{y}$ and each column $\mathbf{m}_i$ of M, and declares $\mathcal{D}$ to be the set of d indices for which this dot-product has largest absolute value.

One can show [12] that there exists a universal constant c_3 such that if $\mathbf{z} = 0$ then with probability at least $1 - d^{-c_3}$ (over the choice of M, which is assumed to be independent of the vector $\mathbf{x}$) this procedure correctly

[4] As opposed to an *approximately d-sparse* vector, *i.e.*, a vector such that "most" of its energy is confined to d indices of $\mathbf{x}$. One way of characterizing such vectors is to say that $||\mathbf{x} - \mathbf{x}_d||_1 \leq c_1 ||\mathbf{x}||_1$ for some suitably small $0 < c_1 < 1$. Here $\mathbf{x}_d$ is defined as the vector matching $\mathbf{x}$ on the d components that are largest in absolute value, and zero elsewhere. The results of [10], [11] also apply in this more general setting. However, in this work we are primarily concerned with the problem of group testing rather than that of compressive sensing, and present the work of compressive sensing merely by way of analogy. Hence we focus on the simplest scenarios in which we can draw appropriate correspondences between the two problems.

reconstructs $\mathbf{x}$ with $T \leq c_3 d \log(n)$ measurements. Similar results can also be shown with $\mathbf{z} \neq 0$, though the form of the result is more intricate.

2) Basis Pursuit: An alternate decoding procedure proceeds by relaxing the compressive sensing problem (in particular the non-linear constraint (3)) into the convex optimization problem called *Basis Pursuit* (BP).

$$\mathbf{x} = \arg\min ||\mathbf{x}||_1 \quad (4)$$

$$\text{subject to } ||\mathbf{y} - M\mathbf{x}||_2 \leq c_2 \quad (5)$$

It can be shown (for instance [10], [11]) that there exist constants c_4, c_5 and c_6 such that with $T = c_4 d \log(n)$, with probability at least $1 - 2^{c_5 n}$, the solution $\mathbf{x}^*$ to BP satisfies $||\mathbf{x}^* - \mathbf{x}||_2 \leq c_6 ||\mathbf{z}||_2$.

II. Background

A. Model and Notation

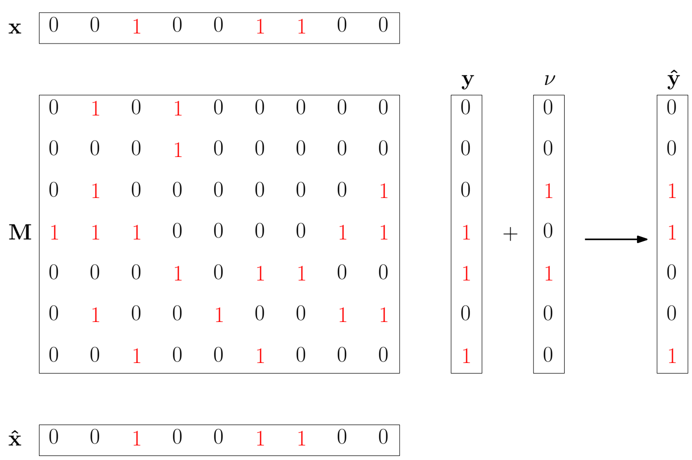

Fig. 1. An example demonstrating a typical non-adaptive group-testing setup. The $T \times n$ binary group-testing matrix represents the items being tested in each test, the length-n binary input vector $\mathbf{x}$ is a weight d vector encoding the locations of the d defective items in $\mathcal{D}$, the length-T binary vector $\mathbf{y}$ denotes the outcomes of the group tests in the absence of noise, the length-T binary noisy result vector $\hat{\mathbf{y}}$ denotes the actually observed noisy outcomes of the group tests, as the result of the noiseless result vector being perturbed by the length-T binary noise vector ν. The length-n binary estimate vector $\hat{\mathbf{x}}$ represents the estimated locations of the defective items.

A set $\mathcal{N}$ contains n items, of which an unknown subset $\mathcal{D}$ are said to be "defective".[5] The goal of group-testing is to correctly identify the set of defective items via a minimal number of "group tests", as defined below (see Figure 1 for a graphical representation).

[5]It is common (see for example [13]) to assume that the number d of defective items in $\mathcal{D}$ is known, or at least a good upper bound D on d, is known *a priori*. If not, other work (for example [14]) considers non-adaptive algorithms with low query complexity that help estimate d. However, in this work we explicitly consider algorithms that do not require such foreknowledge of d – rather, our algorithms have "good" performance with $\mathcal{O}(D \log(n))$ measurements.

Each row of a $T \times n$ binary *group-testing matrix* M corresponds to a distinct test, and each column corresponds to a distinct item. Hence the items that comprise the group being tested in the ith test are exactly those corresponding to columns containing a 1 in the ith location. The method of generating such a matrix M is part of the design of the group test – this and the other part, that of estimating the set $\mathcal{D}$, is described in Section II-B.

The length-n binary *input* vector $\mathbf{x}$ represents the set $\mathcal{N}$, and contains 1s exactly in the locations corresponding to the items of $\mathcal{D}$. The locations with ones/defective items are said to be *positive* – the other locations are said to be *negative*. We use these terms interchangeably.

The outcomes of the *noiseless* tests correspond to the length-T binary *noiseless result* vector $\mathbf{y}$, with a 1 in the i location if and only if the ith test contains at least one defective item.

The observed vector of test outcomes in the *noisy* scenario is denoted by the length-T binary *noisy result* vector $\hat{\mathbf{y}}$ – the probability that each entry y_i of $\mathbf{y}$ differs from the corresponding entry $\hat{y}_i$ in $\hat{\mathbf{y}}$ is q, where q is the *noise parameter*. The locations where the noiseless and the noisy result vectors differ is denoted by the length-T binary *noise vector* ν, with 1s in the locations where they differ.

The estimate of the locations of the defective items is encoded in the length-n binary *estimate vector* $\hat{\mathbf{x}}$, with 1s in the locations where the group-testing algorithms described in Section II-B estimate the defective items to be.

The *probability of error* of any group-testing algorithm is defined as the probability (over the input vector $\mathbf{x}$, group-testing matrix M, and noise vector ν) that the estimated vector differs from the input vector.

B. Algorithms

We now describe the COMP and CBP algorithms in both the noiseless and noisy settings. The algorithms are specified by the choices of encoding matrices and decoding algorithms. Their performance is stated in Section III and the corresponding proofs of the algorithms are presented in Section V.

1) "Column-based" Algorithms: We first consider algorithms that consider columns of the measurement matrix M, and try to correlate these with the observations $\mathbf{y}$. We consider two scenarios – the first when the observations are noiseless, and the second when they are noisy. In both cases we draw the analogy with a corresponding compressive sensing algorithm.

Combinatorial Orthogonal Matching Pursuit (COMP):

The $T \times n$ group-testing matrix M is defined as follows. A *group selection parameter* p is chosen (the exact values of T and p are code-design parameters to be specified later). Then, i.i.d. for each (i, j), $m_{i,j}$ (the (i, j)th element of M) is set to be one with probability p, and zero otherwise.

The decoding algorithm works "column-wise" on M – it attempts to match the columns of M with the result vector $\mathbf{y}$. That is, if a particular column j of M has the property that all locations i where it has ones *also* corresponds to ones in y_i in the result vector, then the jth item (x_j) is declared to be defective (positive). All other

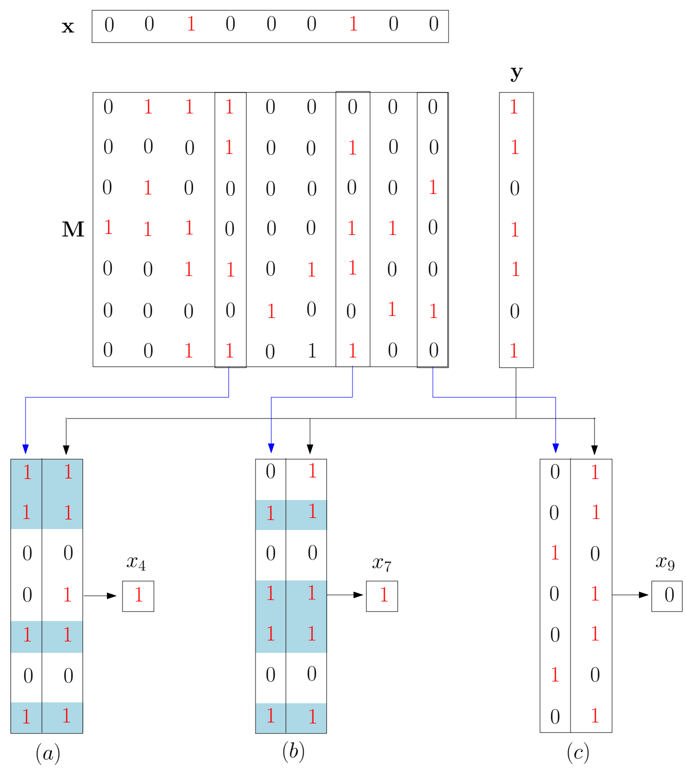

Fig. 2. An example demonstrating the COMP algorithm. The algorithm matches columns of M to the result vector. As in (b) in the figure, since the result vector "contains" the 7th column, then the decoder declares that item to be defective. Conversely, as in (c), since there is no such containment of the last column, then the decoder declares that item to be non-defective. However, sometimes, as in (a), an item that is truly negative, is "hidden" by some other columns corresponding to defective items, leading to a false positive.

items are declared to be non-defective (negative). [6]

Note that this decoding algorithm never has false negatives, only false positives. A false positive occurs when *all* locations with ones in the jth column of M (corresponding to a non-defective item j) are "hidden" by the ones of other columns corresponding to defectives items. That is, let column j and some other columns $j_1, \ldots, j_k$ of matrix M be such that for each i such that $m_{i,j} = 1$, there exists an index j' in $\{j_1, \ldots, j_k\}$ for which $m_{i,j'}$ also equals 1. Then if each of the $\{j_1, \ldots, j_k\}$th items are defective, then the jth item will also always be declared as defective by the COMP decoder, regardless of whether or not it actually is. The probability of this event happening becomes smaller as the number of tests T become larger.

The rough correspondence between this algorithm and Orthogonal Matching Pursuit ([12]) arises from the fact that, as in Orthogonal Matching Pursuit, the decoder attempts to match the columns of the group-testing matrix with the result vector.

[6] Note the similarity between this algorithm and OMP. Another way of phrasing the above decoding rule is to say that the dot-product between $\mathbf{y}$ and each column of M should equal the number of ones in that column. Hence the name Combinatorial OMP, with Combinatorial stressing that the underlying problem (group-testing) is Combinatorial rather than Linear (Compressive sensing).

Noisy Combinatorial Orthogonal Matching Pursuit (NCOMP)

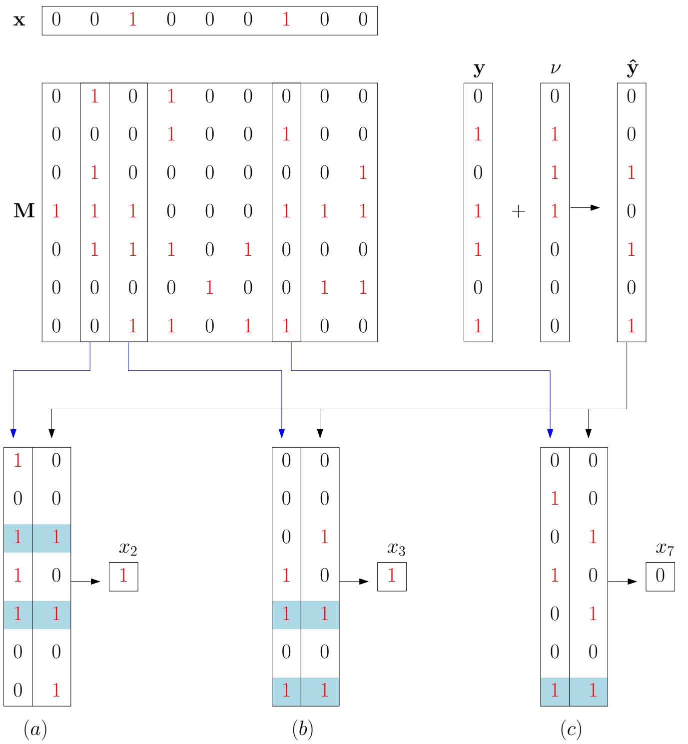

Fig. 3. An example demonstrating the NCOMP algorithm. The algorithm matches columns of M to the result vector *up to a certain number of mismatches* governed by a threshold. In this example, the threshold is set so that the number mismatches be less than the number of matches. For instance, in (b) above, the 1s in the third column of the matrix match the 1s in the result vector in two locations (the 5th and 7th rows), but do not match only in one location in the 4th row (locations wherein there are 0s in the matrix columns but 1s in the result vector do not count as mismatches). Hence the decoder declares that item to be defective, which is the correct decision.
However, consider the false negative generated for the item in (c). This corresponds to the 7th item. The noise in the 2nd, 3rd and 4th rows of ν means that there is only one match (in the 7th row) and two mismatches (2nd and 4th rows) – hence the decoder declares that item to be non-defective.
Also, sometimes, as in (a), an item that is truly negative, has a sufficient number of measurement errors that the number of mismatches is reduced to be below the threshold, leading to a false positive.

In the noisy COMP case, we relax the sharp-threshold requirement in the original COMP algorithm that the set of locations of ones in any column of M corresponding to a positive item be *entirely* contained in the set of locations of ones in the result vector. Instead, we allow for a certain number of "mismatches" – this number of mismatches depends on both the number of ones in each column, and also the noise parameter q.

Let p and τ be design parameters to be specified later. To generate the M for the NCOMP algorithm case, each element of M is selected i.i.d. with probability p to be 1.

The decoder proceeds as follows, For each column j, we define the *indicator set* $\mathcal{T}_j$ as the set of indices i in that column where $m_{i,j} = 1$. We also define the *matching set* $\mathcal{S}_j$ as the set of indices j where both $\hat{y}_i = 1$ (corresponding to the noisy result vector) and $m_{i,j} = 1$.

Then the decoder uses the following "relaxed" thresholding rule. If $|\mathcal{S}_i| \geq |\mathcal{T}_i|(1 - q(1+\tau))$, then the decoder declares the ith item to be defective, else it declares it to be non-defective.[7]

2) "Row-based" algorithms:

We now consider algorithms that consider rows of the measurement matrix M, and try to correlate these with the observations $\mathbf{y}$. We again consider two scenarios – the first when the observations are noiseless, and the second when they are noisy. In both cases we draw the analogy with a corresponding compressive sensing algorithm. This class of algorithms work "row-wise" on M, rather than column-wise as in the OMP algorithms presented above. We first present a combinatorial algorithm for the noiseless case, and then a suite of Linear Programs (LPs) for both the noiseless and noisy cases.

Combinatorial Basis Pursuit (CBP):

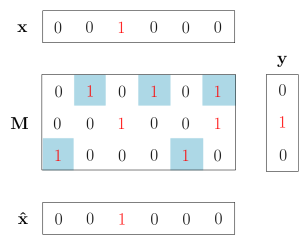

Fig. 4. An example demonstrating the CBP algorithm. Based on only on the outcome of the negative tests (those with output zero), the decoder estimates the set of non-defective items, and "guesses" that the remaining items are defective.

The $T \times n$ group-testing matrix M is defined as follows. A *group sampling parameter* g is chosen (the exact values of T and g are code-design parameters to be specified later). Then, the ith row of M is specified by sampling with replacement[8] from the set $[1, \ldots, n]$ exactly g times, and setting the (i, j) location to be one if j is sampled at least once during this process, and zero otherwise.[9]

The decoding algorithm proceeds by using *only* the tests which have a negative (zero) outcome, to identify all the non-defective items, and declaring all other items to be defective. If M is chosen to have enough rows (tests), each non-defective item should, with significant probability, appear in at least one negative test, and hence will be appropriately accounted for. Errors (false positives) occur when at least one non-defective item is not tested, or only occurs in positive tests (*i.e.,* every test it occurs in has at least one defective item). The analysis of this type of algorithm comprises of estimating the trade-off between the number of tests and the probability of error.

[7]As in Footnote 6, this "relaxed" thresholding rule can be viewed as relaxing the requirement on the dot-product between $\hat{\mathbf{y}}$ and the columns of M.

[8]Sampling without replacement is a more natural way to build tests, but the analysis is trickier. However, the performance of such a group-testing scheme can be shown to be no worse than the one analyzed here [15]. Also see Footnote 9.

[9]Note that this process of sampling each item in each test with replacement results in a slightly different distribution than if the group-size of each test was fixed *a priori* and hence the sampling was "without replacement" in each test. (For instance, in the process we define, each test may, with some probability, test fewer than g items.) The primary advantage of analyzing the "with replacement" sampling is that in the resulting group-testing matrix every entry is then chosen *i.i.d.*.

$$\hat{\mathbf{x}} = \text{ feasible point in} \tag{6}$$

$$\sum_{j:m_{ij}=1} x_j = 0, \text{ if } y_i = 0, \tag{7}$$

$$\sum_{j:m_{ij}=1} x_j \geq 1, \text{ if } y_i = 1, \tag{8}$$

$$\sum_{\forall i} x_j \leq D \tag{9}$$

$$0 \leq x_j \leq 1 \tag{10}$$

Fig. 5. The constraint set of CBP-LP.

More formally, for all tests i whose measurement outcome y_i is a zero, let $\mathbf{m}_i$ denote the corresponding ith row of M. The decoder outputs $\hat{\mathbf{x}}$ as the length-n binary vector which has 0s in exactly those locations where there is a 1 in at least one such $\mathbf{m}_i$.

Combinatorial Basis Pursuit via LP decoding (CBP-LP):

In fact, a linear relaxation of CBP leads naturally to the following set of LPs, which attempts to find a feasible point in the constraint set given by (6–10).

Working backwards through (7)-(10), constraint (10) relaxes the constraint that each $x_j \in \{0, 1\}$ in the usual manner, constraint (9) indicates that there are at most D defective items, constraint (8) indicates that if the ith test outcome is positive ($y_i = 1$) then there must be at least one defective item in the set of items being tested in the ith test (at least one x_j being tested must be 1), and finally constraint (7) indicates that if the ith test outcome is negative ($y_i = 1$) then there cannot be any defective items being tested in test i. The decoder returns $\hat{\mathbf{x}}$ as any feasible solution to this LP.

The next algorithm, NCBP-LP, includes CBP-LP as a special case with no noise. Hence we defer discussion of CBP-LP to the discussion of NCBP-LP.

Noisy Combinatorial Basis Pursuit via LP decoding (NCBP-LP):

In the noisy measurement scenario, the constraints of type (7) and (8) may no longer hold. We hence have to add "slack" variables η_i for all $i \in \{1, \ldots, T\}$. For a particular test i this η_i is defined to be zero if a particular test result is correct, and positive (and at least 1) if the test result is incorrect. Of course, the decoder does not know *a priori* which scenario a particular test outcome falls under, and hence has to also decode η. Nonetheless, as is common in the field of error-correction [8], often using a "minimum distance" decoding criteria (decoding to a vector pair $(\mathbf{x}, \eta)$ such that the error-vector η is as "small" as possible) leads to good decoding performance. Our LP decoder attempts to do so.

To be more precise, we solve the LP in (11–17).

Here the variables η_i are "slack variables" in the equations (12) and (13). For instance, if test i is truly negative, then all the variables in an equation of the form (12) are zero. However, if the test is a false negative, then the

$$(\hat{\mathbf{x}}, \hat{\eta}) = \arg\min_{\mathbf{x},\eta} \sum_{i:\hat{y}_i=1} \eta_i + \frac{1}{e} \sum_{i:\hat{y}_i=0} \eta_i \tag{11}$$

such that

$$-\eta_i + \sum_{j:m_{ij}=1} x_j = 0, \text{ if } \hat{y_i} = 0, \tag{12}$$

$$\eta_i + \sum_{j:m_{ij}=1} x_j \geq 1, \text{ if } \hat{y_i} = 1, \tag{13}$$

$$\sum_{\forall j} x_j \leq D \tag{14}$$

$$0 \leq x_j \leq 1. \tag{15}$$

$$0 \leq \eta_i \leq D, \text{ if } \hat{y_i} = 0, \tag{16}$$

$$0 \leq \eta_i \leq 1, \text{ if } \hat{y_i} = 1, \tag{17}$$

Fig. 6. The LP corresponding to NCBP-LP.

variable η_i is then set to equal the number of defective items in test i. Similarly, if a test i is truly positive, then η_i is zero, and equations of the form (13) are satisfied. However, if the test is a false positive, then η_i is set to equal 1 (and the variables tested in test j are set to zero). Note that η_i is bounded above by 1 in the case of (false) positives, but is only bounded above by D in the case of (false) negatives. This is due to the asymmetry of positive and negative test outcomes – multiple positive items tested simultaneously do not give a different outcome from a single positive item tested. The reason that the objective function is split into two parts, with the slack variables corresponding to negative outcomes being weighted less (by a factor of $1/e$) than the slack variables corresponding to positive outcomes is exactly to compensate for this asymmetry.[10]

In fact, it turns out that an even simpler LP still gives essentially the same performance as NCBP-LP.

Noisy Combinatorial Basis Pursuit via Simpler LP decoding (NCBP-SLP):

Consider the LP given in (18–22). The set of constraints for this LP is a subset of the set of constraints for NCBP-LP (in fact, it only examines the set of tests with positive outcomes). Nonetheless, we can demonstrate that this LP decoding algorithm has comparable performance as NCBP-LP. The intuition is that if NCBP-LP works by finding a η vector with low Hamming weight, then NCBP-SLP does the same by finding a η vector with low Hamming weight restricted just to the set of positive outcomes. Since the noise that converts $\mathbf{y}$ to $\hat{\mathbf{y}}$ is probabilistic,

[10]The reason for this particular choice of the weighting factor equaling $1/e$ comes out of the analysis of our algorithm. We are able to show that for such a value, our algorithm has the claimed performance. However, we have not optimized this weighting factor so as to minimize the number of tests T, since the corresponding calculations quickly become hard to solve exactly in closed form (though they are tractable to computer calculation).

$$(\hat{\mathbf{x}}, \hat{\eta}) = \arg\min_{\mathbf{x},\eta} \sum_{\forall i: \hat{y}_i = 1} \eta_i \tag{18}$$

such that

$$\eta_i + \sum_{j: m_{ij}=1} x_j \geq 1, \text{ if } \hat{y}_i = 1, \tag{19}$$

$$\sum_{\forall j} x_j \leq D \tag{20}$$

$$0 \leq x_j \leq 1. \tag{21}$$

$$0 \leq \eta_i \leq 1, \text{ if } \hat{y}_i = 1, \tag{22}$$

Fig. 7. The LP corresponding to CBP-SLP.

by standard concentration results these two approaches should, with high probability, lead to the same result. [11]

III. Main Results

A. Lower bounds on the number of tests required

We first provide information-theoretic lower bounds on the number of tests required by *any* group-testing algorithm. While we believe these bounds to be "common knowledge" in the field, we have been unable to pinpoint a reference that gives an explicit lower bound on the number of tests in terms of the acceptable probability of error of the group-testing algorithm. For the sake of completeness, so we can benchmark our analysis of the algorithms we present later, we state and prove the lower bounds here. All logarithms in this work are assumed to be binary, except where explicitly identified as otherwise (in some cases we explicitly denote the natural logarithm as $\ln(.)$).

Theorem 1: [Folklore] Any group-testing algorithm with noiseless measurements that has a probability of error of at most ϵ requires at least $(1-\epsilon)D\log(n/D)$ tests.

In fact, the corresponding lower bounds can be extended to the scenario with noisy measurements as well.

Theorem 2: [Folklore] Any group-testing algorithm that has measurements that are noisy i.i.d. with probability q and that has a probability of error of at most ϵ requires at least $[(1-\epsilon)D\log(n/D)]/(1-H(q))$ tests.[12]

Note: Our assumption that $D = o(n)$ implies that the bounds in Theorem 1 and 2 are $\Omega(D\log(n))$.

B. Upper Bounds on the number of tests required

The main contributions of this work are explicit computations of the number of tests required to give a desired probability of error via computationally efficient algorithms. Some of these algorithms are new (LP-based decoding

[11] However, it is interesting to note that the complement of the above statement is not true. In particular, if one removes (13) and (17) from NCBP-LP, to examine only the tests with *negative* outcomes, this approach fails with high probability. This is because the corresponding LP preferentially returns $\hat{\mathbf{x}}$ as the all-zero vector, which indeed would trivially satisfy an LP that examines only negative outcomes. Perhaps this indicates the importance of maintaining a positive attitude in life.

[12] Here $H(.)$ denotes the binary entropy function.

for CBP-LP and NCBP-LP in Theorems). Others have tighter analysis than previously in the literature (COMP and NCOMP) novel analysis (coupon-collector based analysis for CBP, perturbation analysis for CBP-LP and NCBP-LP). In each case, to the best of our knowledge ours is the first work to explicitly compute the tradeoff between the number of tests required to give a desired probability of error, rather than giving order of magnitude estimates of the number of tests required for a "reasonable" probability of success. Also, the novel analysis in the CBP, CBP-LP, and NCBP-LP algorithms may well extend to other information theory problems such as compressive sensing.

Theorem 3: COMP with error probability at most $n^{-\delta}$ requires no more than $eD(1+\delta)\ln(n)$ tests.

Translating COMP into the noisy observation scenario is non-trivial. A more careful analysis for the thresholded scheme in NCOMP leads to the following result.

Theorem 4: NCOMP with error probability at most $n^{-\delta}$ requires no more than $\frac{16(1+\sqrt{\gamma})^2(1+\delta)\ln 2}{(1-e^{-2})(1-2q)^2}D\log n$ tests, where γ is a constant that can be explicitly calculated, and lies in the interval $[\delta, 1)$.

Theorem 5: CBP with error probability at most $n^{-\delta}$ requires no more than $2(1+\delta)eD\ln n$ tests.

The analysis of the constants in the next three theorems are not optimized (doing so is analytically very cumbersome), but are given to demonstrate the functional dependence on δ and q. We define Γ as $\ln(d)/\ln(n)$ (note that in the limit of large n it lies in the interval $[0, 1)$).

For the same noiseless observation scenario as CBP, the LP-based decoding algorithm CBP-LP has the following performance guarantees.

Theorem 6: CBP-LP with error probability at most $n^{-\delta}$ requires no more than $8e\,(\delta+1+\Gamma)\,D\ln n$ tests.

In fact Theorem 6 is implied by the stronger analysis (for the noisy observations scenario) for NCBP-LP in Theorem 7 below.

Theorem 7: NCBP-LP with error probability at most $n^{-\delta}$ requires no more than $\beta_{LP}D\ln n$ tests, with β_{LP} defined as

$$\max\left\{\frac{4e\,(\delta+1+\Gamma)}{(1-2q)^2}, 8e\,(\delta+1+\Gamma), \frac{4(1-q+2qe)e\,(\delta+1+\Gamma)}{(1-q)^2},\right.$$
$$\left.\frac{8e\,(\delta+1+\Gamma)}{(1-q+2qe)}, \frac{(1-q+qe)\,(\delta+\Gamma)\,(1+e)^2}{e(1-2q)^2}, \frac{8e\,(\delta+\Gamma)}{(1-q+qe)}\right\}.$$

Essentially the same analysis as in Theorem 7 in fact then also implies Theorem 8 below, leading to a simpler LP than NCBP-LP, with essentially the same performance.

Theorem 8: NCBP-LP with error probability at most $n^{-\delta}$ requires no more than $\beta_{SLP}D\ln n$ tests, with β_{SLP} defined as

$$\max\left\{\frac{4e\,(\delta+1+\Gamma)}{(1-2q)^2}, 8e\,(\delta+1+\Gamma), \frac{(1-q+qe)\,(\delta+\Gamma)\,(1+e)^2}{e(1-2q)^2}\right\}.$$

Note: Our achievability schemes in Theorems 3-7 are commensurate (equal up to a constant factor) with the lower bounds in Theorems 1 and 2. For instance, the bound on the number of tests in Theorem 4 differs from the corresponding lower bound in Theorem 2 by a factor that is at most $12.83(1+\sqrt{\gamma})^2(1+\delta)(1-2q)^{-2}$, which is a function only of q and δ. For "small" q and δ this quantity is "small".

IV. Proof of the lower bounds in Theorems 1 and 2

The "usual" proofs of lower bounds for group-testing are combinatorial. To incorporate the allowed probability of error ϵ into our lower bounds, we provide information-theoretic proofs.

We begin by noting that $\mathbf{X} \to \mathbf{Y} \to \hat{\mathbf{Y}} \to \hat{\mathbf{X}}$ (*i.e.* the input vector, noiseless result vector, noisy result vector, and the estimate vector) forms a Markov chain. By standard information-theoretic definitions we have

$$H(\mathbf{X}) = H(\mathbf{X}|\hat{\mathbf{X}}) + I(\mathbf{X};\hat{\mathbf{X}})$$

Since $\mathbf{X}$ is uniformly distributed over all length-n and D-sparse data vectors (since d could be as large as D), $H(\mathbf{X}) = \log|\mathcal{X}| = \log\binom{n}{D}$. By Fano's inequality, $H(\mathbf{X}|\hat{\mathbf{X}}) \leq 1 + \epsilon \log\binom{n}{D}$. Also, we have $I(\mathbf{X};\hat{\mathbf{X}}) \leq I(\mathbf{Y};\hat{\mathbf{Y}})$ by the data-processing inequality. Finally, note that

$$I(\hat{\mathbf{Y}};\hat{\mathbf{Y}}) \leq \sum_{i=1}^{T} \left[H(\hat{Y}_i) - H(\hat{Y}_i|Y_i)\right]$$

since the first term is maximized when each of the $\hat{Y}_i$ are independent, and because the measurement noise is memoryless. For the BSC(q) noise we consider in this work, this summation is at most $T(1-H(q))$ by standard arguments.[13]

Combining the above inequalities, we obtain

$$(1-\epsilon)\log\binom{n}{D} \leq 1 + T(1-H(q))$$

Also, by standard arguments via Stirling's approximation [16], $\log\binom{n}{D}$ is at least $D\log(n/D)$. Substituting this gives us the desired result

$$\begin{aligned} T &\geq \frac{1-\epsilon}{1-H(q)} \log\binom{n}{D} \\ &\geq \frac{1-\epsilon}{1-H(q)} D \log\left(\frac{n}{D}\right). \end{aligned}$$

■

V. Proof of the performance of algorithms in Theorems 3-7

A. Column-based algorithms

We first consider column-based algorithms. The COMP and NCOMP algorithms respectively deal with the noiseless and noisy observation scenarios.

Proof of Theorem 3:

As noted in the discussion on COMP in Section II-B, the error-events for the algorithm correspond to false positives, when a column of M corresponding to a non-defective item is "hidden" by other columns corresponding to defective items. To calculate this probability, recall that each entry of M equals one with probability p, i.i.d. Let j index

[13] This technique also holds for more general types of discrete memoryless noise – for ease of presentation, in this work we focus on the simple case of the Binary Symmetric Channel.

a column of M corresponding to a non-defective item, and let $j_1, \ldots, j_d$ index the columns of M corresponding to defective items. Then the probability that $m_{i,j}$ equals one, and at least one of $m_{i,j_1}, \ldots, m_{i,j_d}$ *also* equals one is $p(1-(1-p)^d)$. Hence the probability that the jth column is hidden by a column corresponding to a defective item is $\left(1-p(1-p)^d\right)^T$. Taking the union bound over all $n-d$ non-defective items gives us that the probability of false positives is bounded from above by

$$P_e = P_e^+ \leq (n-d)\left(1-p(1-p)^d\right)^T. \tag{23}$$

By differentiation, optimizing (23) with respect to p suggests choosing p as $1/d$. However, the precise value of d may not be known, only D, an upper bound on it, might be. Substituting the value $p = 1/D$ back into (23), and setting T as $\beta D \ln n$ gives us

$$\begin{aligned} P_e &\leq (n-d)\left(1-\frac{1}{D}\left(1-\frac{1}{D}\right)^d\right)^{\beta D \ln n} \\ &\leq (n-d)\left(1-\frac{1}{De}\right)^{\beta D \ln n} && (24) \\ &\leq (n-d)e^{-\beta e^{-1}\ln n} \\ &\leq n^{1-\beta e^{-1}}. && (25) \end{aligned}$$

Inequality (24) follows from the previous since $d \leq D$ by definition, and since $(1-1/x)^x \geq e^{-1}$. Choosing $\beta = (1+\delta)e$ thus ensures the required decay in the probability of error. Hence choosing T to be at least $(1+\delta)eD\ln n$ suffices to prove the theorem. ■.

Proof of Theorem 4:

Due to the presence of noise, both false positives and false negatives may occur in the noisy COMP algorithm – the overall probability of error is the sum of the probability of false positives and that of false negatives. As in the previous algorithm, we set $p = 1/D$ and $T = \beta D \log n$. We first calculate the probability of false negatives by computing the probability that more than the expected number of ones get flipped to zero in the result vector in locations corresponding to ones in the column indexing the defective item. This can be computed as

$$
\begin{aligned}
P_e^- &= \bigcup_{i=1}^{d} P\left(|\mathcal{T}_i| = t\right) \Pr\left(|\mathcal{S}_i| < |\mathcal{T}_i|(1-q(1+\tau))\right) \\
&\leq d \sum_{t=0}^{T} \binom{T}{t} p^t (1-p)^{T-t} \qquad (26) \\
&\qquad \sum_{r=t-t(1-q(1+\tau))}^{t} \binom{t}{r} q^r (1-q)^{t-r} \\
&\leq d \sum_{t=0}^{T} \binom{T}{t} p^t (1-p)^{T-t} e^{-2t(q\tau)^2} \qquad (27) \\
&= d\left(1 - p + pe^{-2(q\tau)^2}\right)^T \qquad (28) \\
&= d\left(1 - \frac{1}{D} + \frac{1}{D} e^{-2(q\tau)^2}\right)^{\beta D \log n} \qquad (29) \\
&\leq d \exp\left[-\beta \log n \left(1 - e^{-2(q\tau)^2}\right)\right] \qquad (30) \\
&\leq d \exp\left[-\beta \log n (1 - e^{-2})(q\tau)^2\right] \qquad (31)
\end{aligned}
$$

Here, as in Section II-B, $\mathcal{T}_i$ denotes the locations of ones in the ith column of M. Inequality (26) follows from the union bound over the possible errors for each of the defective items, with the first summation accounting for the different possible sizes of $\mathcal{T}_i$, and the second summation accounting for the error events corresponding to the number of one-to-zero flips exceeding the threshold chosen by the algorithm. Inequality (27) follows from the Chernoff bound. Equality (28) comes from the binomial theorem. Equality (29) comes from substituting in the values of p and T. Inequality (30) follows from the leading terms of the Taylor series of the exponential function. Inequality (31) follows from bounding the concave function $1 - e^{-2x}$ by the linear function $(1 - e^{-2})x$ for $x > 0$.

The requirement that the probability of false negatives P_e^- to be at most $n^{-\delta}$ implies that β^- (the bound on β due to this restriction) satisfies

$$
\begin{aligned}
& \ln\left(d \exp\left[-\beta(1 - e^{-2})(q\tau)^2 \log n\right]\right) < -\delta \ln n \\
\Rightarrow \quad & \ln d - \tfrac{\beta(1-e^{-2})(q\tau)^2}{\ln 2} \ln n < -\delta \ln n \\
\Rightarrow \quad & \beta^- > \tfrac{\left(\frac{\ln d}{\ln n} + \delta\right) \ln 2}{(1-e^{-2})(q\tau)^2} \qquad (32)
\end{aligned}
$$

We now focus on the probability of false positives. In the noiseless COMP algorithm, the only way a false positive could occur was if all the ones in a column are hidden by ones in columns corresponding to defective items. In the noisy COMP algorithm this still happens, but in addition noise could also lead to a similar masking effect. That is, even in the 1 locations of a non-defective column not hidden by other defective columns, measurement noise may flip enough zeroes to ones so that the decoding threshold is exceeded, and the decoder hence incorrectly declares that particular item to be defective. See Figure 3(a) for an example.

Hence we define a new quantity a, which denotes the probability for any (i, j)th location in M that a 1 in that location is "hidden by other columns *or* by noise". It equals

$$\begin{aligned} a &= 1-[(1-q)(1-p)^d+q(1-(1-p)^d)] \\ &= \left(1-q-\left(1-\frac{1}{D}\right)^d(1-2q)\right) \end{aligned} \tag{33}$$

We set $D \geq 2$ (the case $D=1$ can be handled separately by the same analysis, but setting $p=1/(D+1)=1/2$ rather than $1/D=1$), and note that by definition $d \leq D$. We then bound a from above as

$$\begin{aligned} \max_{D\geq 2, d\leq D} a &= \max_{D\geq 2, d\leq D}\left(1-q-\left(1-\frac{1}{D}\right)^d(1-2q)\right) \\ &= 1-q-(1-2q)\min_{D\geq 2, d\leq D}\left(\left(1-\frac{1}{D}\right)^d\right) \qquad (34) \\ &= 1-q-(1-2q)\min_{D\geq 2}\left(\left(1-\frac{1}{D}\right)^D\right) \qquad (35) \\ &= (1-q)-(1-2q)/4. \qquad (36) \end{aligned}$$

Equation (36) follows from the observations that (34) is optimized when $d=D$ and (35) is optimized when $D=2$. The probability of false positives is then computed in a similar manner to that of false negatives as in (26)–(31).

$$\begin{aligned} P_e^+ &= \bigcup_{i=1}^{n-d} P\left(|\mathcal{T}_i|=t\right) P\left(|\mathcal{S}_i| \geq |\mathcal{T}_i|(1-q(1+\tau))\right) \\ &\leq (n-d)\sum_{t=0}^{T}\binom{T}{t}p^t(1-p)^{T-t} \\ &\qquad \sum_{r=t(1-q(1+\tau))}^{t}\binom{t}{r}a^r(1-a)^{t-r} \\ &\leq (n-d)\left(1-p+pe^{-2((1-q(1+\tau))-a)^2}\right)^T \qquad (37) \\ &\leq (n-d)\left(1-p+pe^{-2((1-2q)/4-q\tau)^2}\right)^T \qquad (38) \\ &\leq (n-d)\exp\left[-(1-e^{-2((1-2q)/4-q\tau)^2})\beta\log n\right] \qquad (39) \\ &\leq (n-d)\exp\left[-\beta\log n(1-e^{-2})\left((1-2q)/4-q\tau)^2\right)\right] \qquad (40) \end{aligned}$$

Note that for the Chernoff bound to be applicable in (37), $1-q(1+\tau)>a$, which implies that $\tau<(1-2q)/(4q)$. Equation (38) follows from substituting the bound derived on a in (36) into (37), and (39) follows by substituting $p=1/D$ into the previous equation. Inequality (40) follows from bounding the concave function $1-e^{-2x}$ by the linear function $(1-e^{-2})x$ for $x>0$.

The requirement that the probability of false positives P_e^+ be at most $n^{-\delta}$ implies that β^+ (the bound on β due to this restriction) be at least

$$\beta^+ > \frac{\left(\frac{\ln(n-d)}{\ln n}+\delta\right)\ln 2}{(1-e^{-2})((1-2q)/4-q\tau)^2}. \tag{41}$$

Note that β must be at least as large as $\max\{\beta^-,\beta^+\}$ so that both (32) and (41) are satisfied.

When the threshold in the NCOMP algorithm is high (*i.e.,* τ is small) then the probability of false negatives increases; conversely, the threshold being low (τ being large) increases the probability of false positives. Algebraically, this expresses as the condition that $\tau > 0$ (else the probability of false negatives is significant), and conversely to the condition that $1-q(1+\tau) > a$ (so that the Chernoff bound can be used in (37)) – combined with (36) this implies that $\tau \leq (1-2q)/4q$. Each of (32) and (41) as a function of τ is a reciprocal of a parabola, with a pole at the corresponding extremal value of τ. Furthermore, β^- is strictly increasing and β^+ is strictly decreasing in the region of valid τ in $(0,(1-2q)/(4q))$. Hence the corresponding curves on the right-hand sides of (32) and (41) intersect within the region of valid τ, and a good choice for β is at the τ where these two curves intersect. To find this β, we make another simplifying substitution. Let γ be defined as

$$\gamma = \lim_{n,d\to\infty} \frac{\ln d + \delta \ln n}{\ln(n-d) + \delta \ln n}. \tag{42}$$

and Γ as

$$\Gamma = \lim_{n,d\to\infty} \frac{\ln d}{\ln n}.$$

Hence

$$\gamma = \frac{\Gamma + \delta}{1+\delta}. \tag{43}$$

(Note that since $d = o(n)$, Γ lies in the interval $[0,1)$, and hence for large n, γ approaches a constant in the interval $[\delta, 1)$.) Then equating the RHS of (32) and (41) implies that the optimal τ^* satisfies

$$\frac{\ln 2}{(1-e^{-2})((1-2q)/4 - q\tau^*)^2} = \frac{\gamma \ln 2}{(1-e^{-2})(q\tau^*)^2} \tag{44}$$

Simplifying (44) gives us that

$$\tau^* = \frac{1-2q}{4q(1+\gamma^{-1/2})}. \tag{45}$$

Substituting these values of γ and τ into (32) gives us the explicit bound for large n

$$\beta^* = \frac{16(1+\gamma^{-0.5})^2(\Gamma+\delta)\ln 2}{(1-e^{-2})(1-2q)^2}. \tag{46}$$

Using (43) to simplify (46) gives

$$\beta^* = \frac{16(1+\sqrt{\gamma})^2(1+\delta)\ln 2}{(1-e^{-2})(1-2q)^2} \approx \frac{12.83(1+\sqrt{\gamma})^2(1+\delta)}{(1-2q)^2}.$$

∎

B. *Row-based algorithms*

We now consider row-based algorithms. We first consider CBP whose analysis is based on a novel use of the Coupon Collector Problem [17]. We then consider CBP-LP and NCBP-LP (respectively for noiseless and noisy observation scenarios), whose analyses correspond to a novel "perturbation method" that has potential applications in other estimation problems (such as compressive sensing).

Proof of Theorem 5:

The Coupon Collector's Problem (CCP) is a classical problem that considers the following scenario. Suppose there are n types of coupons, each of which is equiprobable. A collector selects coupons (with replacement) until he has a coupon of each type. What is the distribution on his stopping time? It is well-known ([17]) that the expected stopping time is $n \ln n + \Theta(n)$. Also, reasonable bounds on the tail of the distribution are also known – for instance, it is known that the probability that the stopping time is more than $\chi n \ln n$ is at most $n^{-\chi+1}$.

Analogously to the above, we view the group-testing procedure of CBP as a Coupon Collector Problem. Consider the following thought experiment. Suppose we consider any test as a length-g *test-vector*[14] whose entries index the items being tested in that test (repeated entries are allowed in this vector, hence there might be less than g distinct items in this vector). Due to the design of our group-testing procedure in CBP, the probability that any item occurs in any location of such a vector is uniform and independent. In fact this property (uniformity and independence of the value of each entry of each test) also holds *across* tests. Hence, the items in any subsequence of k tests may be viewed as the outcome of a process of selecting a single chain of gk coupons. This is still true even if we restrict ourselves solely to the tests that have a negative outcome. The goal of CBP may now be viewed as the task of collecting *all* the *non-defective* items. This can be summarized in the following equation

$$Tg\left(\frac{n-d}{n}\right)^g \geq (n-d)\ln(n-d). \tag{47}$$

The left-hand side of Equation (47) refers to the expected number of (possibly repeated) items in negative tests (since there are a total of T tests, each containing g (possibly repeated) items, and the probability of a test being negative equals $((n-d)/n)^g$). The right-hand side of (47) refers to the expected stopping-time of the underlying CCP. We thus optimize (47) w.r.t. g to obtain an optimal value of g equaling $1/\ln(n/(n-d))$. However, since the exact value of d is not known, but rather only D, an upper bound on it, we set g to equal $1/\ln(n/(n-D))$. Taking the appropriate limit of n going to infinity, and noting $D = o(n)$, enables us to determine that, in expectation over the testing process and the location of the defective items, (47) implies that $T \geq eD\ln n$.

However, (47) only holds in expectation. For us to design a testing procedure for which we can demonstrate that the number of tests decays to zero as $n^{-\delta}$, we need to modify (47) to obtain the corresponding tail bound on T. This takes a bit more work.

The right-hand side is then modified to $\chi(n-d)\ln(n-d)$. This corresponds to the event that all types of coupons have not been collected if $\chi(n-d)\ln(n-d)$ total coupons have been collected. The probability of this event is at most $(n-d)^{-\chi+1}$).

The left-hand side is multiplied with $(1-\rho)$, where ρ is a design parameter to be specified by Chernoff's bound on the probability that the actual number of items in the negative tests is smaller than $(1-\rho)$ times the expected number. By Chernoff's bound this is at most $\exp\left(-\rho^2 T\left(\frac{n-d}{n}\right)^g\right)$. Taking the union bound over these

[14]Note that this test-vector is different from the binary length-n vectors that specify tests in the group testing-matrix, though there is indeed a natural bijection between them.

two low-probability events gives us that the probability that

$$(1-\rho)Tg\left(\frac{n-d}{n}\right)^{g} \geq \chi(n-d)\ln(n-d) \tag{48}$$

does *not* hold is at most

$$\exp\left(-\rho^2 T\left(\frac{n-d}{n}\right)^{g}\right) + (n-d)^{-\chi+1}. \tag{49}$$

So, we again optimize for g in (47) and substitute $g^* = 1/\ln\left(\frac{n}{n-D}\right)$ into (48). We note that since both D and d are $o(n)$, $\left(\frac{n-d}{n}\right)^{g^*}$ converges to e^{-1}. Hence we have, for large n,

$$\begin{aligned} T &\geq \frac{\chi}{1-\rho}\frac{(n-d)\ln(n-d)}{g^*\left(\frac{n-d}{n}\right)^{g^*}} \\ &\approx \frac{\chi}{1-\rho}\frac{(n-d)\ln(n-d)}{\frac{1}{\ln\left(\frac{n}{n-D}\right)}e^{-1}} \\ &= \frac{\chi}{1-\rho}\frac{(n-d)\ln(n-d)\ln\left(\frac{n}{n-D}\right)}{e^{-1}}. \end{aligned} \tag{50}$$

Using the inequality $\ln(1+x) \geq x - x^2/2$ with x as $D/(n-D)$ simplifies the RHS of (50) to

$$T \geq \frac{\chi}{1-\rho}e\left(D - \frac{D^2}{2(n-d)}\right)\ln(n-d). \tag{51}$$

Choosing T to be greater than the bound in (51) can only reduce the probability of error, hence choosing

$$T \geq \frac{\chi}{1-\rho}eD\ln(n-d)$$

still implies a probability of error at most as large as in (49).

Choosing $\rho = \frac{1}{2}$, noting that $D \geq d$, and substituting (52) into (49) implies, for large enough d, the probability of error P_e satisfies

$$\begin{aligned} P_e &\leq e^{-\frac{\delta^2\chi}{1-\delta}d\ln(n-d)} + (n-d)^{-\chi+1} \\ &= (n-d)^{-\frac{\delta^2}{1-\delta}\chi d} + (n-d)^{-\chi+1} \\ &\leq 2(n-d)^{-\chi+1}. \end{aligned}$$

Taking $2(n-d)^{-\chi+1} = n^{-\delta}$, we have $\chi = \delta\frac{\log n}{\log(n-d)} + \frac{1}{\log(n-d)} + 1$. For large n, χ approaches $\delta + 1$.

Therefore, the probability of error is at most $n^{-\delta}$, with sufficiently large n, the following number of tests suffice to satisfy the probability of error condition stated in the theorem.

$$T \geq 2(1+\delta)eD\ln n.$$

■

We first prove Theorem 7, and then derive Theorems 6 and 8 as direct corollaries.

Proof of Theorem 7:

At a high level, our proof proceeds as follows. First, we define two *finite* sets $\mathbf{\Phi}'$ and $\mathbf{\Phi}''$ containing so-called "perturbation vectors" (these vectors, defined below, depend only on $\mathbf{x}$). We demonstrate in Claim 9 that any $\bar{\mathbf{x}}$ in the feasible set of the constraint set of NCBP-LP can be written as the true $\mathbf{x}$ plus a *non-negative linear combination* of perturbation vectors from one or both of these sets. The linear combination property is important, since this enables us to characterize the directions in which a vector can be perturbed from $\mathbf{x}$ to another vector that satisfies the constraints of NCBP-LP, in a "finite" manner (instead of having to consider the uncountably infinite number of directions that $\mathbf{x}$ could be perturbed to). The non-negativity of the linear combination is also crucial since, as we explain below, this property ensures that the objective function of the LP can only increase when perturbed in a direction corresponding to any vector in $\mathbf{\Phi}'$ or $\mathbf{\Phi}''$.

In Claims 10 and 11 (which form the heart of our argument) we then characterize the expected change (over randomness in the matrix M and noise ν) in the value of each slack variable η_i when $\mathbf{x}$ is perturbed to some $\bar{\mathbf{x}}$ by a vector in $\mathbf{\Phi}'$ or $\mathbf{\Phi}''$. In particular, we demonstrate that for each such *individual* perturbation vector, the expected change in the value of each slack variable η_i is *strictly* positive with high probability.

In Claim 13 with slightly careful use of standard concentration inequalities (specifically, we need to use both the additive and multiplicative forms of the Chernoff bound, reprised in Claim 12) we show that the probability distributions derived in Claim 11. We then take the union bound over all vectors in $\mathbf{\Phi}'$ and $\mathbf{\Phi}''$ (in fact, there are a total of $d(n-d)+d$ such vectors in $\mathbf{\Phi}'$ and $\mathbf{\Phi}''$ together) and show that with high probability the expected change in the value of the objective function (which equals the weighted sum of the changes in the values of the slack variables η_i) for *each* perturbation vector in $\mathbf{\Phi}' \cup \mathbf{\Phi}''$ is also strictly positive. Finally, we note that the set of feasible $(\bar{\mathbf{x}}, \eta)$ satisfying NCBP-LP forms a convex set. Hence for any $\bar{\mathbf{x}} \neq \mathbf{x}$ in the feasible set of NCBP-LP, the value of NCBP-LP's objective function corresponding to $\bar{\mathbf{x}}$ must be strictly greater than the value of the objective function corresponding to $\mathbf{x}$, which implies that the LP decodes correctly as $\hat{\mathbf{x}} = \mathbf{x}$.

We proceed by proving a sequence of claims that when strung together formalize the above argument.

Without loss of generality, let $\mathbf{x}$ be the vector with 1s in the first d locations, and 0s in the last $n-d$ locations.[15] Choose $\mathbf{\Phi}' = \{\phi'\}_{k'=1}^{d(n-d)}$ as the set of $d(n-d)$ vectors with a single -1 in the support of $\mathbf{x}$, a single 1 outside the support of $\mathbf{x}$, and zeroes everywhere else. For instance, the first ϕ' in the set equals $(-1,0,\ldots,0,1,0,\ldots,0)$, where the 1 is in the $(d+1)$th location. Analogously, choose $\mathbf{\Phi}'' = \{\phi''\}_{k''=1}^{d}$ as the set of d vectors with a single -1 in the support of $\mathbf{x}$, and zeroes everywhere else. For instance, the first ϕ'' in the set equals $(-1,0,\ldots,0)$.

Then, Claim 9 below gives a "nice" characterization of the set of $\bar{\mathbf{x}}$ in the feasible set of NCBP-LP.

Claim 9: Any vector $\bar{\mathbf{x}}$ that satisfies the constraints (12–17) in NCBP-LP can be written as

$$\bar{\mathbf{x}} = \mathbf{x} + \sum_{k'=1}^{d(n-d)} c_{k'}\phi'_{k'} + \sum_{k''=1}^{d} c_{k''}\phi''_{k''}, \text{ with all } c_{k'}, c_{k''} \geq 0. \tag{52}$$

Proof of Claim 9: The intuition behind this claim is a type of "conservation law", so to speak. A good analogy is the following two-stage physical process.

[15] As can be verified, our analysis is agnostic to the actual choice of $\mathbf{x}$, as long as it is a vector in $\{0,1\}^n$ of weight any $d \leq D$.

Imagine that $\mathbf{x}$ corresponds to n bottles of water each with capacity one litre, with the first d bottles full, and the remaining empty. Imagine $\bar{\mathbf{x}}$ as another state of these bottles with $\bar{d} \leq d$ litres of water.

In the first stage, we throw away $d - \bar{d}$ litres of water (taking care not to undershoot – *i.e.,* to ensure that we always leave at least $\bar{x}_j$ litres in the jth bottle, for each of the first d bottles). This stage corresponds to using non-negative linear combinations of perturbation vectors from the set $\mathbf{\Phi}''$ (non-negativity arises from the fact that we took care not to undershoot).

Then in the second stage, for each bottle j among the first d bottles that still has more water remaining than in the corresponding bottle in the final state $\bar{x}_i$, we use its water to increase the water level of bottles among the last $n - d$ bottles (taking care not to overshoot, *i.e.,* not to put more than the desired water level $\bar{x}_i$ in any such bottle). The fact that this is doable follows from conservation of mass. This stage corresponds to using non-negative linear combinations of perturbation vectors from the set $\mathbf{\Phi}'$ (non-negativity arises from the fact the we took care not to overshoot).[16] $\square$

Next, Claim 11 below computes the expected change in the value of the slack variable η_i as $\mathbf{x}$ is perturbed by ϕ' or ϕ''. A small example in Table V-B (with $n = 3$, $d = 2$) demonstrates the calculations in the proof of Claim 11 explicitly.

For any fixed $\mathbf{x} \in \{0,1\}^n$ of weight $d \leq D$, let $\mathbf{x}' = \mathbf{x} + \phi'$, and $\mathbf{x}'' = \mathbf{x} + \phi''$. Over the randomness in $\mathbf{m}_i$ and the noise in the test outcome ν_i, we define the *cost perturbation* random variables

$$\begin{aligned} \Delta'_{0,i} = (\eta_i(\mathbf{x}') - \eta_i(\mathbf{x})), \quad \Delta''_{0,i} = (\eta_i(\mathbf{x}'') - \eta_i(\mathbf{x})) \text{ conditioned on } \hat{y}_i = 0, \\ \Delta'_{1,i} = (\eta_i(\mathbf{x}') - \eta_i(\mathbf{x})), \quad \Delta''_{1,i} = (\eta_i(\mathbf{x}'') - \eta_i(\mathbf{x})) \text{ conditioned on } \hat{y}_i = 1. \end{aligned} \tag{53}$$

Claim 10: The cost perturbation random variables defined in (53) all take values only in $\{-1, 0, 1\}$.

Proof of Claim 10: We first analyze the case when if $\hat{y}_i = 0$. By (12), $\eta_i(\mathbf{x}) = \mathbf{m}_i.\mathbf{x}$. Hence $\Delta''_{0,i} = \eta_i(\mathbf{x}'') - \eta_i(\mathbf{x}) = \mathbf{m}_i.(\mathbf{x}'' - \mathbf{x}) = \mathbf{m}_i.\phi''$. Similarly $\Delta'_{0,i} = \mathbf{m}_i.\phi'$. But ϕ'' has exactly one non-zero component (equaling -1), and ϕ' has exactly one component equaling -1 and one equaling 1. By definition, $\mathbf{m}_i$, is a $0/1$ vector. Hence both $\Delta''_{0,i}$ and $\Delta'_{0,i}$ take values in $\{-1, 0, 1\}$.

Similarly, if $\hat{y}_i = 1$, by (13), the minimum value of $\eta_i(\mathbf{x})$ occurs at $(1 - \mathbf{m}_i.\mathbf{x})^+$ (*i.e.*, equals $1 - \mathbf{m}_i.\mathbf{x}$ if $1 - \mathbf{m}_i.\mathbf{x}$ is positive, and 0 otherwise). Since $\mathbf{m}_i$, $\mathbf{x}$, $\mathbf{x}'$and $\mathbf{x}''$ are all $0/1$ vectors. Hence $\eta_i(\mathbf{x})$, $\eta_i(\mathbf{x}')$ and $\eta_i(\mathbf{x}'')$ are all $0/1$ vectors, and thus their pairwise differences (and in particular $\Delta''_{1,i}$ and $\Delta'_{1,i}$) take values in $\{-1, 0, 1\}$. $\square$.

The next claim forms the heart of our proof. It does an exhaustive[17] case analysis that computes the probabilities that the cost perturbation random variables take values 1 or -1 (the case that they equal zero can be derived from these calculations in a straightforward manner too, but since these values turn out not to matter for our analysis, we omit these details).

[16]Formalizing this intuition to explicitly obtain the non-negative linear combination is tedious simply due to notation complexity rather than any intrinsic hardness in the concept, so we leave it as an exercise for HAROLD (Hypothetical Alert Reader Of Limitless Dedication).

[17]And exhausting!

Claim 11:

$$P(\Delta''_{1,i} = 1) = p(1-p)^{d-1}(1-q), \qquad P(\Delta''_{1,i} = -1) = 0, \tag{54}$$

$$P(\Delta''_{0,i} = 1) = 0, \qquad P(\Delta''_{0,i} = -1) = pq, \tag{55}$$

$$P(\Delta'_{1,i} = 1) = p(1-p)^{d}(1-q), \qquad P(\Delta'_{1,i} = -1) = p(1-p)^{d}q, \tag{56}$$

$$P(\Delta'_{0,i} = 1) = p(1-p)\left[(1-q)(1-p)^{d-1} + q\right], \qquad P(\Delta'_{0,i} = -1) = p(1-p)q. \tag{57}$$

Proof of Claim 11: The proof follows from a case-analysis similar to the one performed in the example in Table V-B. The reader is strongly encouraged to read that example before looking at the following case analysis, which can appear quite intricate.

Without loss of generality (due to the symmetry of the distribution on $\mathbf{m}_i$), we assume that $\mathbf{x}$ has 1s in the first d locations and 0s elsewhere.

To ease the reader into the analysis, we begin by considering the easier cases first.

- Equation (55) analyzes $\Delta''_{0,i}$, the expected change in η_i if $\mathbf{x}$ is perturbed by a vector from $\mathbf{\Phi}''$, and $\hat{y}_i = 0$. In this case, note that by (12), $\eta_i(\mathbf{x}) = \mathbf{m}_i.\mathbf{x}$ and $\eta_i(\mathbf{x}'') = \mathbf{m}_i.\mathbf{x}''$. Hence $\Delta''_{0,i} = \eta_i(\mathbf{x}'') - \eta_i(\mathbf{x}) = \mathbf{m}_i.(\mathbf{x}'' - \mathbf{x}) = \mathbf{m}_i.\phi''$. But ϕ'' has only a single negative component and no positive components, hence $\Delta''_{0,i}$ is never positive. Conversely, $\Delta''_{0,i}$ equals -1 whenever the support of $\mathbf{m}_i$ intersects the support of $\mathbf{x}$ at the location where $\phi'' = -1$ (which happens with probability pq since the first term corresponds to the probability of a 1 in that location in $\mathbf{m}_i$, and since y_i equals 1 in this case but we assumed that $\hat{y}_i = 0$), giving the desired result.
- Analogously, Equation (54) analyzes $\Delta''_{1,i}$, the expected change in η_i if $\mathbf{x}$ is perturbed by a vector from $\mathbf{\Phi}''$, and $\hat{y}_i = 1$. In this case, note that by (13), the minimum value of $\eta_i(\mathbf{x})$ occurs at $(1 - \mathbf{m}_i.\mathbf{x})^+$, (*i.e.*, equals $1 - \mathbf{m}_i.\mathbf{x}$ if $1 - \mathbf{m}_i.\mathbf{x}$ is positive, and 0 otherwise). Note that since $\mathbf{m}_i$, $\mathbf{x}$, and $\mathbf{x}''$ are all binary vectors, therefore this function must also have integral values, and in fact must equal either 0 or 1. But both $\mathbf{x}$ and $\mathbf{x}''$ are $0/1$-vectors, and the support of $\mathbf{x}''$ is a strict subset of the support of $\mathbf{x}$ with exactly one less positive component, hence $\eta_i(\mathbf{x}'') \geq \eta_i(\mathbf{x})$, which means that $\Delta''_{1,i} \geq 0$. The strict inequality holds if and only if $\eta_i(\mathbf{x}'') = 1$ and $\eta_i(\mathbf{x}) = 0$, or equivalently if $\mathbf{m}_i.\mathbf{x}'' = 1$ and $\mathbf{m}_i.\mathbf{x} = 0$. The only scenarios when this happens is when the support of $\mathbf{m}_i$ intersects the support of $\mathbf{x}$ at exactly one location (hence y_i equals 1), and $\phi'' = -1$ at exactly this location. Thus, only d indices of $\mathbf{m}_i$ matter for this scenario, and of these indices, $d-1$ must equal 0 and one must be 1. These scenarios occurs with probability $(1-q)(1-p)^{d-1}p$ (the term q indicates probability of the event that $y_i = 1$ and $\hat{y} = 1$, and the remaining terms indicate the probability of $\mathbf{m}_i$ being as specified). Computing the expectation of Δ''_1 due to this scenario gives the desired result.

The analysis of Equations (56) and (57) is more intricate.

- Equation (56) analyzes $\Delta'_{1,i}$ the expected change in η_i if $\mathbf{x}$ is perturbed by a vector from $\mathbf{\Phi}'$, and $\hat{y}_i = 1$. As in (54), the minimum value of $\eta_i(\mathbf{x})$ occurs at $(1 - \mathbf{m}_i.\mathbf{x})^+$ and must equal either 0 or 1. We now consider the cases $\Delta'_{1,i} = -1$ and $\Delta'_{1,i} = 1$ separately.

Using similar arguments as before, the only scenarios when $\Delta'_{1,i} = -1$ occur when $y_i = 0$, the support of $\mathbf{m}_i$ is entirely outside the support of $\mathbf{x}$, and further that $\mathbf{m}_i$ equals 1 in the location where $\phi' = 1$. Thus, only $d+1$ indices of $\mathbf{m}_i$ matter for this scenario – of these indices, d must equal 0 and one must be 1. These scenarios occur with probability $q(1-p)^d p$.

Analogously, the only scenarios when $\Delta'_{1,i} = 1$ occur when $y_i = 1$, the support of $\mathbf{m}_i$ intersects the support of $\mathbf{x}$ exactly at the location where $\phi' = -1$, and further that $\mathbf{m}_i$ equals 0 in the location where $\phi' = 1$. Thus, only $d+1$ indices of $\mathbf{m}_i$ matter for this scenario. Of these indices, the d must equal 0 and one must be 1. These scenarios occurs with probability $(1-q)(1-p)^d p$ (the $(1-q)$ term indicates probability of the event that $y_i = 1$ and $\hat{y} = 1$, and the remaining terms indicate the probability of $\mathbf{m}_i$ being as specified).

- Equation (57) analyzes $\Delta'_{0,i} = -1$, the expected change in η_i if $\mathbf{x}$ is perturbed by a vector from $\mathbf{\Phi}'$, and $\hat{y}_i = 0$. As in Equation (55), $\eta_i(\mathbf{x}) = \mathbf{m}_i.\mathbf{x}$, hence $\Delta'_0 = \eta_i(\mathbf{x}') - \eta_i(\mathbf{x}) = \mathbf{m}_i.\phi'$.

  We first analyze the case when $\Delta'_{0,i} = -1$. This occurs only when $y_i = 1$, $\mathbf{m}_i = 1$ where $\phi' = -1$, and further that $\mathbf{m}_i$ equals 0 in the location where $\phi' = 1$. Thus, only 2 indices of $\mathbf{m}_i$ matter for this scenario. These scenarios occurs with probability $q(1-p)p$.

  Analogously, the only scenarios when $\Delta'_{0,i} = 1$ occur when $y_i = 1$, $\mathbf{m}_i$ equals 1 in the location where $\phi' = 1$ and $\mathbf{m}_i$ equals 0 in the location where $\phi' = -1$. This can happen in two ways. It could be that if the support of $\mathbf{m}_i$ is entirely outside the support of $\mathbf{x}$ (hence $y_i = 0$), and further that $\mathbf{m}_i$ equals 1 in the location where $\phi' = 1$ (this happens with probability $(1-q)(1-p)^d p$). Or, it could be that $\mathbf{m}_i$ equals 0 where $\phi' = -1$, $\mathbf{m}_i$ equals 1 in at least one of the other $(d-1)$ locations in the support of $\mathbf{x}$ (which ensures that $y_i = 1$), and further that $\mathbf{m}_i$ equals 1 where $\phi' = 1$ (the probability of such a scenario is $q(1-p)(1-(1-p)^{d-1})p$). Adding together the probabilities corresponding to these two scenarios gives the desired result.

□

We now recall specific forms we need of standard concentration inequalities, that we use to make statements about the expected value of objective value of NCBP-LP.

Claim 12 (Chernoff bound (multiplicative form) [17]): Let $\{W_i\}_{i=1}^T$ be a sequence i.i.d. binary random variables with probability distribution $P(W_i = 1) = \theta$.

$$P\left(\left|\sum_{i=1}^{T}\frac{W_i}{T} - \theta\right| > \sigma\right) \leq \exp\left(-2T\sigma^2\right), \text{ (additive form)}, \tag{58}$$

$$P\left(\left|\sum_{i=1}^{T}\frac{W_i}{T} - \theta\right| > \sigma\theta\right) \leq \exp\left(-\frac{T\sigma^2\theta}{2}\right), \text{ (multiplicative form)}. \tag{59}$$

Finally, Claim 13 below demonstrates that regardless of which direction $\mathbf{x}$ is perturbed in, as long as it remains within the feasible set for NCBP-LP, with high probability over M and the noise vector ν, the value of the objective function of NCBP-LP increases. This implies that the minimization performed in NCBP-LP returns $\hat{\mathbf{x}} = \mathbf{x}$ correctly, and completes the proof.

1. $\hat{\mathbf{y}}_i$	2. $\eta(\mathbf{x})$	3. $\mathbf{y}_i$	4. $\mathbf{m}_i$	5. $P(\hat{\mathbf{y}}_i, \mathbf{m}_i \vert \mathbf{x})$	6. $\eta_i(\mathbf{m}_i, \mathbf{x})$	$\mathbf{x}$ $(\mathbf{1},\mathbf{1},\mathbf{0})$ 7. $\eta_i(\mathbf{m}_i, \mathbf{x})$	$\mathbf{x}' = \mathbf{x} + \phi'$ $(\mathbf{0},\mathbf{1},\mathbf{1})$ 8(a). $\eta_i(\mathbf{m}_i, \mathbf{x}')$	8(b). $\mathrm{E}(\mathbf{m}_i, \Delta_i')$	$\mathbf{x}'' = \mathbf{x} + \phi''$ $(\mathbf{0},\mathbf{1},\mathbf{0})$ 9(a). $\eta_i(\mathbf{m}_i, \mathbf{x}'')$	9(b). $\mathrm{E}(\mathbf{m}_i, \Delta_i'')$
1	$(1 - \mathbf{m}_i.\mathbf{x})^+$	0	$(\mathbf{0},\mathbf{0},\mathbf{0})$	$q(1-p)^3$	1	1	1	0	1	0
			$(\mathbf{0},\mathbf{0},\mathbf{1})$	$q(1-p)^2p$	$(1-x_3)^+$	1	0	$-q(1-p)^2p$	1	0
		1	$(\mathbf{0},\mathbf{1},\mathbf{0})$	$(1-q)(1-p)^2p$	$(1-x_2)^+$	0	0	0	0	0
			$(\mathbf{0},\mathbf{1},\mathbf{1})$	$(1-q)(1-p)p^2$	$(1-x_2-x_3)^+$	0	0	0	0	0
			$(\mathbf{1},\mathbf{0},\mathbf{0})$	$(1-q)(1-p)^2p$	$(1-x_1)^+$	0	1	$(1-q)(1-p)^2p$	1	$(1-q)(1-p)^2p$
			$(\mathbf{1},\mathbf{0},\mathbf{1})$	$(1-q)(1-p)p^2$	$(1-x_1-x_3)^+$	0	0	0	1	$(1-q)(1-p)p^2$
			$(\mathbf{1},\mathbf{1},\mathbf{0})$	$(1-q)(1-p)p^2$	$(1-x_1-x_2)^+$	0	0	0	0	0
			$(\mathbf{1},\mathbf{1},\mathbf{1})$	$(1-q)p^3$	$(1-x_1-x_2-x_3)^+$	0	0	0	0	0
								$(1-2q)(1-p)^2p$		$(1-q)(1-p)p$
0	$\mathbf{m}_i.\mathbf{x}$	0	$(\mathbf{0},\mathbf{0},\mathbf{0})$	$(1-q)(1-p)^3$	0	0	0	0	0	0
			$(\mathbf{0},\mathbf{0},\mathbf{1})$	$(1-q)(1-p)^2p$	x_3	0	1	$(1-q)(1-p)^2p$	0	0
		1	$(\mathbf{0},\mathbf{1},\mathbf{0})$	$q(1-p)^2p$	x_2	1	1	0	1	0
			$(\mathbf{0},\mathbf{1},\mathbf{1})$	$q(1-p)p^2$	x_2+x_3	1	2	$q(1-p)p^2$	1	0
			$(\mathbf{1},\mathbf{0},\mathbf{0})$	$q(1-p)^2p$	x_1	1	0	$-q(1-p)^2p$	0	$-q(1-p)^2p$
			$(\mathbf{1},\mathbf{0},\mathbf{1})$	$q(1-p)p^2$	x_1+x_3	1	0	0	0	$-q(1-p)p^2$
			$(\mathbf{1},\mathbf{1},\mathbf{0})$	$q(1-p)p^2$	x_1+x_2	2	1	$-q(1-p)p^2$	1	$-q(1-p)p^2$
			$(\mathbf{1},\mathbf{1},\mathbf{1})$	qp^3	$x_1+x_2+x_3$	2	2	0	1	$-qp^3$
								$(1-2q)(1-p)^2p$		$-qp$

TABLE I. Suppose $\mathbf{x} = (1,1,0)$. Choose some $\mathbf{x}' \neq \mathbf{x}$ (in this example, $\mathbf{x}' = \mathbf{x} + \phi'$ and $\mathbf{x}'' = \mathbf{x} + \phi''$, where $\phi' = (-1,0,1)$ and $\phi'' = (-1,0,0)$ are the two types of *perturbation vectors*). This example analyzes the expectation (over the randomness in the particular row $\mathbf{m}_i$ of the measurement matrix M) of the difference in value of the corresponding slack variables $\eta_i(\mathbf{x})$ and $\eta_i(\mathbf{x}')$ in column $8(b)$, and also between $\eta_i(\mathbf{x})$ and $\eta_i(\mathbf{x}'')$ in column $9(b)$. We distinguish between the cases when $\mathbf{y}_i$ equals zero, and when it equals one, and compute the corresponding quantities separately. To compute these, we consider the columns of the table above sequentially from left to right. Column 1 considers the two possible values of the observed vector $\hat{\mathbf{y}}_i$. Column 2 gives the corresponding values of the slack variables corresponding to the ith test, as returned by the constraints (12) and (13) of NCBP-LP – here $(f(\mathbf{x}))^+$ denotes the function $\max\{f(\mathbf{x}), 0\}$. Column 3 indexes the possibilities of the (noiseless) test outcomes $\mathbf{y}_i$, and column 4 enumerates possible values for $\mathbf{m}_i$, the i-th row of M, that could have generated the values of $\mathbf{y}_i$ in column 3, given that $\mathbf{x} = (1,1,0)$. Column 5 computes the probability of a particular observation $\hat{\mathbf{y}}_i$ and a row $\mathbf{m}_i$, given that the noiseless output $\mathbf{y}$ equaled a particular value. Column 6 computes the function in column 2, given that $\mathbf{m}_i$ equals the value given in Column 4. Columns 7 and $8(a)$ respectively explicitly compute the value of the function in column 6 for the vectors $\mathbf{x}$ and $\mathbf{x}'$ – the red entries in column $8(a)$ index those locations where $\eta(\mathbf{x}')$, the slack variable for the perturbed vector, is less than $\eta(\mathbf{x})$, and the green cells indicate those locations where the situation is reverse. Column $8(b)$ then computes the product of column 5 with the difference of the entries in column 7 from those of column $8(a)$, *i.e.*, the expected change in the value of the slack variable $\eta_i()$. The value $(1-2q)(1-p)^2p$ in blue at the bottom represents the expected change (averaged over all possible tuples $(\mathbf{y}_i, \mathbf{m}_i, \hat{\mathbf{y}}_i)$). Columns $9(a)$ and $9(b)$ compute values analogous to columns $8(a)$ and $8(b)$, for the perturbation vector $\phi'' = (-1,0,0)$.

Claim 13: Choose T as $\beta_{LP} D \log(n)$, with β_{LP} as given in Theorem 7. Then NCBP-LP fails with a probability of error of at most $n^{-\delta}$.

Proof of Claim 13: Since each $\mathbf{m}_i$ is chosen independently, the random variables corresponding to each type of cost perturbation variable in (53) are distributed i.i.d. according to (54)-(57).

We define the sets $\mathcal{S}'_1 = \{i|\Delta'_{1,i} \neq 0\}$, $\mathcal{S}'_0 = \{i|\Delta'_{0,i} \neq 0\}$, and $\mathcal{S}'' = \{i|\Delta''_{1,i}, \Delta''_{0,i} \neq 0\}$. By (54)-(57), the expected sizes of $\mathcal{S}'_1$, $\mathcal{S}'_0$ and $\mathcal{S}''$ respectively equal

$$E(|\mathcal{S}'_1|) = Tp(1-p)^d > \frac{\beta_{LP}\log(n)}{e}, \tag{60}$$

$$E(|\mathcal{S}'_0|) = Tp(1-p)((1-p)^{d-1}(1-q)+2q) > \frac{\beta_{LP}(1-q+2qe)\log(n)}{e}. \tag{61}$$

$$E(|\mathcal{S}''|) = Tp((1-p)^{d-1}(1-q)+q) > \frac{\beta_{LP}(1-q+qe)\log(n)}{e}, \tag{62}$$

The above inequalities follow by using the definition of $p = 1/D$, T as defined in this claim, and $(1-1/D)^{d-1} > (1-1/D)^D > 1/e$. Hence by the concentration inequality (59) above, for some positive constants σ'_1, σ'_0 and σ'',

$$P\left(|\mathcal{S}'_1| < E(|\mathcal{S}'_1|)(1-\sigma'_1)\right) < n^{-\frac{\beta_{LP}(\sigma'_1)^2}{2e}}, \tag{63}$$

$$P\left(|\mathcal{S}'_0| < E(|\mathcal{S}'_0|)(1-\sigma'_0)\right) < n^{-\beta_{LP}\left(\frac{1-q+2qe}{2e}\right)(\sigma'_0)^2}, \tag{64}$$

$$P\left(|\mathcal{S}''| < E(|\mathcal{S}''|)(1-\sigma'')\right) < n^{-\beta_{LP}\left(\frac{1-q+qe}{2e}\right)\sigma''^2}. \tag{65}$$

Next, we compute conditional probabilities, conditioning on the event that i belongs to a set (in fact it may belong to several, but we compute each of these conditional probabilities individually),

$$\begin{aligned} &P\left(\Delta'_{1,i} = 1|i \in \mathcal{S}'_1\right) = 1-q, && P\left(\Delta'_{1,i} = -1|i \in \mathcal{S}'_1\right) = q,\\ &P\left(\Delta'_{0,i} = 1|i \in \mathcal{S}'_0\right) = \frac{(1-p)^{d-1}(1-q)+q}{(1-p)^{d-1}(1-q)+2q}, && P\left(\Delta'_{0,i} = -1|i \in \mathcal{S}'_0\right) = \frac{q}{(1-p)^{d-1}(1-q)+2q},\\ &P\left(\Delta''_{1,i} = 1|i \in \mathcal{S}''\right) = \frac{(1-p)^{d-1}(1-q)}{(1-p)^{d-1}(1-q)+q}, && P\left(\Delta''_{0,i} = -1|i \in \mathcal{S}''\right) = \frac{q}{(1-p)^{d-1}(1-q)+q}. \end{aligned} \tag{66}$$

In the limit $D \to \infty$, (66) reduces to

$$P\left(\Delta'_{1,i} = 1|i \in \mathcal{S}'_1\right) = 1-q, \qquad P\left(\Delta'_{1,i} = -1|i \in \mathcal{S}'_1\right) = q, \tag{67}$$

$$P\left(\Delta'_{0,i} = 1|i \in \mathcal{S}'_0\right) = \frac{1-q+qe}{1-q+2qe}, \qquad P\left(\Delta'_{0,i} = -1|i \in \mathcal{S}'_0\right) = \frac{qe}{1-q+2qe}, \tag{68}$$

$$P\left(\Delta''_{1,i} = 1|i \in \mathcal{S}''\right) = \frac{1-q}{1-q+qe}, \qquad P\left(\Delta''_{0,i} = -1|i \in \mathcal{S}''\right) = \frac{qe}{1-q+qe}. \tag{69}$$

By the definition of the cost perturbation variables and (11) the objective value of NCBP-LP for $\mathbf{x}' = \mathbf{x} + \phi'$ equals the objective value of NCBP-LP for $\mathbf{x}$ plus the *objective value perturbation*

$$\sum_{i \in \mathcal{S}'_1} (1(\Delta'_{1,i} = 1) - 1(\Delta'_{1,i} = -1)) + \frac{1}{e}\sum_{i \in \mathcal{S}'_0} (1(\Delta'_{0,i} = 1) - 1(\Delta'_{0,i} = -1)). \tag{70}$$

The first term[18] in (70) is non-positive if and only if $\sum_{i\in\mathcal{S}'_1} 1(\Delta'_{1,i}=1)$ equals $\sum_{i\in\mathcal{S}'_1} 1(\Delta'_{1,i}=-1)$. By (67) and the concentration inequality (58) this happens with probability at most $\exp\left(-\frac{|\mathcal{S}'_1|(1-2q)^2}{2}\right)$. But by (63) and (60), the probability that $|\mathcal{S}'_1|$ is less than $\frac{\beta_{LP}\log(n)(1-\sigma'_1)}{e}$ is at most $n^{-\frac{\beta_{LP}(\sigma'_1)^2}{2e}}$. Hence the probability that the first term in (70) is non-positive is at most

$$n^{-\frac{\beta_{LP}(1-\sigma'_1)(1-2q)^2}{2e}}+n^{-\frac{\beta_{LP}(\sigma'_1)^2}{2e}}. \tag{71}$$

The second term in (70) is non-positive if and only if $\sum_{i\in\mathcal{S}'_0} 1(\Delta'_{0,i}=1)$ equals $\sum_{i\in\mathcal{S}'_0} 1(\Delta'_{0,i}=-1)$. By (68) and the concentration inequality (58) this happens with probability at most $\exp\left(-\frac{|\mathcal{S}'_2|(1-q)^2}{2(1-q+2qe)^2}\right)$. But by (64) and (61), the probability that $|\mathcal{S}'_0|$ is less than $\frac{\beta_{LP}(1-q+2qe)\log(n)(1-\sigma'_0)}{e}$ is at most $n^{-\frac{\beta_{LP}(1-q+2qe)(\sigma'_0)^2}{2e}}$. Hence the probability that the second term in (70) is non-positive is at most

$$n^{-\frac{\beta_{LP}(1-\sigma'_0)(1-q)^2}{2(1-q+2qe)e}}+n^{-\frac{\beta_{LP}(1-q+2qe)(\sigma'_0)^2}{2e}}. \tag{72}$$

Similarly, by the definition of the cost perturbation variables and (11) the objective value of NCBP-LP for $\mathbf{x}''=\mathbf{x}+\phi''$ equals the objective value of NCBP-LP for $\mathbf{x}$ plus the *objective value perturbation*

$$\sum_{i\in\mathcal{S}''}(1(\Delta''_{1,i}=1)-1(\Delta''_{1,i}=-1))+\frac{1}{e}\sum_{i\in\mathcal{S}''}(1(\Delta''_{0,i}=1)-1(\Delta''_{0,i}=-1)). \tag{73}$$

But by (54) and (55), $1(\Delta''_{0,i}=1)$ and $1(\Delta''_{1,i}=-1)$ are always zero. Hence (73) is non-negative if and only if $\sum_{i\in\mathcal{S}''} 1(\Delta''_{1,i}=1)$ equals $e\sum_{i\in\mathcal{S}''} 1(\Delta''_{0,i}=-1)$. By (69) and the concentration inequality (58) this happens with probability at most $\exp\left(-\frac{|\mathcal{S}''|2(1-2q)^2e^2}{(1-q+qe)^2(1+e)^2}\right)$. But by (65) and (62), the probability that $|\mathcal{S}''|$ is less than $\frac{\beta_{LP}(1-q+qe)\log(n)(1-\sigma'')}{e}$ is at most $n^{-\frac{\beta_{LP}(1-q+qe)(\sigma'')^2}{2e}}$. Hence the probability that (73) is non-positive is at most

$$n^{-\frac{2\beta_{LP}e(1-\sigma'')(1-2q)^2}{(1-q+qe)(1+e)^2}}+n^{-\frac{\beta_{LP}(1-q+qe)(\sigma'')^2}{2e}}. \tag{74}$$

We now note that (71) and (72) give bounds on the probability that a single perturbation vector in $\mathbf{\Phi}'$ causes a non-positive perturbation in optimal value of the objective function of NCBP-LP, and similarly (74) does the same for a vector in $\mathbf{\Phi}''$. But there are $d(n-d)$ vectors in $\mathbf{\Phi}'$, and d vectors in $\mathbf{\Phi}''$. We take a union bound over all of these vectors by multiplying the terms in (71) and (72) by $d(n-d)$, and the terms in (74) by d. Hence the overall bound on the probability that even a single vector from $\mathbf{\Phi}'\cup\mathbf{\Phi}''$ causes a non-positive perturbation in optimal value of the objective function of NCBP-LP is given by

$$d(n-d)\left(n^{-\frac{\beta_{LP}(1-\sigma'_1)(1-2q)^2}{2e}}+n^{-\frac{\beta_{LP}(\sigma'_1)^2}{2e}}+n^{-\frac{\beta_{LP}(1-\sigma'_0)(1-q)^2}{2(1-q+2qe)e}}+n^{-\frac{\beta_{LP}(1-q+2qe)(\sigma'_0)^2}{2e}}\right) \tag{75}$$

$$+d\left(n^{-\frac{2\beta_{LP}e(1-\sigma'')(1-2q)^2}{(1-q+qe)(1+e)^2}}+n^{-\frac{\beta_{LP}(1-q+qe)(\sigma'')^2}{2e}}\right). \tag{76}$$

[18]One can in fact improve the bounds somewhat by simultaneously considering both the first and the second term simultaneously, but the corresponding calculations are very messy and not very insightful, hence we instead separately optimize these two terms leading to equations that are somewhat more tractable to symbolic manipulations.

One can further optimize by choosing σ_0', σ_1' and σ'' so that each of the terms in (76) is as small as possible, but the corresponding calculations have a messy closed form solution and are not very insightful, hence we simply choose $\sigma_0' = \sigma_1' = \sigma'' = 1/2$ and obtain

$$d(n-d)\left(n^{-\frac{\beta_{LP}(1-2q)^2}{4e}} + n^{-\frac{\beta_{LP}}{8e}} + n^{-\frac{\beta_{LP}(1-q)^2}{4(1-q+2qe)e}} + n^{-\frac{\beta_{LP}(1-q+2qe)}{8e}}\right) + d\left(n^{-\frac{2\beta_{LP}e(1-2q)^2}{2(1-q+qe)(1+e)^2}} + n^{-\frac{\beta_{LP}(1-q+qe)}{8e}}\right). \tag{77}$$

We then choose β_{LP} as

$$\max\left\{\frac{4e\,(\delta+1+\Gamma)}{(1-2q)^2}, 8e\,(\delta+1+\Gamma), \frac{4(1-q+2qe)e\,(\delta+1+\Gamma)}{(1-q)^2}, \frac{8e\,(\delta+1+\Gamma)}{(1-q+2qe)}, \frac{(1-q+qe)\,(\delta+\Gamma)\,(1+e)^2}{e(1-2q)^2}, \frac{8e\,(\delta+\Gamma)}{(1-q+qe)}\right\} \tag{78}$$

so that the maximum of the six terms in (77) is less than $n^{-\delta}$ (here $\Gamma = (\ln d)/(\ln n)$ is a constant in $[0,1)$). This shows that for large d, n and D, with probability at least $1-n^{-\delta}$ *all* vectors in $\mathbf{\Phi}'$ and $\mathbf{\Phi}''$ cause a *strictly positive* change in the optimal value of the objective function of NCBP-LP.

Finally, we note that the set of feasible $(\bar{\mathbf{x}}, \eta)$ of NCBP-LP forms a convex set. Hence if η strictly increases along every direction in $\mathbf{\Phi}'$ and $\mathbf{\Phi}''$, then in fact η strictly increases when the true $\mathbf{x}$ is perturbed in *any* direction. Hence the true $\mathbf{x}$ must be the solution to NCBP-LP. □ ■

Proof of Theorem 6:

We substitute $q = 0$ into (78) and choose the largest term to obtain the corresponding β as

$$8e\,(\delta+1+\Gamma), \tag{79}$$

■

Proof of Theorem 8:

The proof is essentially the same as in the case of Theorem 7. The only different lies in the fact that η_i now depends only on the tests with positive outcomes. Hence (57) and (55) are not required, and the difference in expectation implies that the claim analogous to Claim 13 has a slightly different concentration result. In particular only the first, second, and fifth terms in (78) survive, hence the corresponding β_{SLP} is

$$\max\left\{\frac{4e\,(\delta+1+\Gamma)}{(1-2q)^2}, 8e\,(\delta+1+\Gamma)\,, \frac{(1-q+qe)\,(\delta+\Gamma)\,(1+e)^2}{e(1-2q)^2}\right\} \tag{80}$$

■

VI. Conclusion

In this work we consider the problem of non-adaptive group-testing with possibly noisy measurements. We provide information-theoretic lower bounds on the number of tests necessary to identify the set of defective items with high probability. We also present a suite of algorithms that match these lower bounds up to "small" factors. In particular, we consider three types of decoding algorithms – "OMP-type", "BP-type", and "LP-type". While the OMP-type and BP-type algorithms are not new, their analysis is tighter than before and results in an explicit characterization of their performance. In the case of the BP-type algorithm, the analysis is novel, and depends on a novel use of the well-studied Coupon Collector Problem. The LP-type algorithms are entirely new, and the "perturbation analysis" used in these algorithms is possibly interesting in its own right for a larger class of sparse recovery problems.

REFERENCES

[1] R. Dorfman, "The detection of defective members of large populations," *Annals of Mathematical Statistics*, vol. 14, pp. 436–411, 1943.

[2] D.-Z. Du and F. K. Hwang, *Combinatorial Group Testing and Its Applications*, 2nd ed. World Scientific Publishing Company, 2000.

[3] G. Atia and V. Saligrama, "Boolean compressed sensing and noisy group testing," *CoRR*, vol. abs/0907.1061, 2009.

[4] A. J. Macula, "Error-correcting nonadaptive group testing with de-disjunct matrices," *Discrete Applied Mathematics*, vol. 80, no. 2-3, pp. 217 – 222, 1997.

[5] E. N. Gilbert, "A comparison of signaling alphabets," *Bell System Technical Journal*, vol. 31, pp. 504–522, 1952.

[6] C. E. Shannon, "A mathematical theory of communication," *The Bell system technical journal*, vol. 27, pp. 379–423, Jul. 1948.

[7] D.-Z. Du and F. K. Hwang, *Pooling designs and nonadaptive group testing: important tools for DNA sequencing*. World Scientific Publishing Company, 2006.

[8] F. J. McWilliams and N. J. A. Sloane, *The Theory of Error-Correcting Codes*. North-Holland, 1977.

[9] D. Sejdinovic and O. Johnson, "Note on noisy group testing: Asymptotic bounds and belief propagation reconstruction," in *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, Oct. 2010, pp. 998 –1003.

[10] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," *Comptes Rendus Mathematique*, vol. 346, no. 9-10, pp. 589–592, 2008.

[11] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, no. 3, pp. 253–263, December 2008.

[12] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Transactions on Information Theory*, vol. 53, pp. 4655–4666, 2007.

[13] A. Macula, "Probabilistic nonadaptive group testing in the presence of errors and dna library screening," *Annals of Combinatorics*, vol. 3, pp. 61–69, 1999.

[14] M. Sobel and R. M. Elashoff, "Group testing with a new goal, estimation," *Biometrika*, vol. 62, no. 1, pp. 181–193, 1975.

[15] J. A. Rice, *Mathematical Statistics and Data Analysis*, 3rd ed. Duxbury Advanced Series, 2006.

[16] T. Cover and J. Thomas, *Elements of Information Theory*. John Wiley and Sons, 1991.

[17] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.