

COMPRESSED SENSING WITH CORRUPTED PARTICIPANTS

Meng Wang^{*}, Weiyu Xu[†], Robert Calderbank[‡]

^{*} Dept. of ECSE, Rensselaer Polytechnic Institute, Troy, NY. [†] Dept. of ECE, University of Iowa, Iowa City, IA.

[‡] Dept. of Electrical Engineering, Duke University, Durham, NC.

ABSTRACT

Compressed sensing (CS) theory promises one can recover real-valued sparse signal from a small number of linear measurements. Motivated by network monitoring with link failures, we for the first time consider the problem of recovering signals that contain both real-valued entries and corruptions, where the real entries represent transmission delays on normal links and the corruptions represent failed links. Unlike conventional CS, here a measurement is real-valued only if it does not include a failed link, and it is corrupted otherwise. We prove that $O((d+1)\max(d,k)\log n)$ nonadaptive measurements are enough to recover all n -dimensional signals that contain k nonzero real entries and d corruptions. We provide explicit constructions of measurements and recovery algorithms. We also analyze the performance of signal recovery when the measurements contain errors.

Index Terms—compressed sensing, group testing, fundamental limits, network tomography, corruptions.

1. INTRODUCTION

Compressed sensing (CS) [1–4] indicates that if an n -dimensional signal is k -sparse, i.e., it only has k nonzero entries, then one can efficiently recover the signal from $O(k\log(n/k))$ nonadaptive linear measurements. Network tomography [5–10] attempts to infer system internal characteristics (e.g., link queueing delays) of the Internet from indirect end-to-end (aggregate) measurements (e.g., path delay measurements). Since only a small number of bottleneck links experience large delays, some recent papers like [11–14] have considered the application of CS in network tomography, where the goal is to recover real-valued sparse link delays from a small number of path delay measurements.

In communication networks, a link between two routers may fail either temporarily or permanently. If a link fails, all the packets that travel through it will be lost. Link failure localization has been extensively investigated, e.g., [15–18], where one attempts to locate the failed links from boolean path measurements. A path measurement is a “success” if it does not pass any failed links. Otherwise, it is treated as a “failure”. This is a group testing (GT) problem [19], see [20–23] as some recent examples of a rich literature.

We propose to locate the failed links and recover the transmission delays on normal links *simultaneously* from a set of *nonadaptive* path measurements. A path measurement is a “failure” if it includes at least one failed link, since its packets will be lost. Otherwise, we obtain the real-valued path delay which is the sum of the link delays of links it passes through. We assume that the number of failed links and the number of nonzero link delays are both small. As far as we know, recovering sparse signals that contain failures is a new problem and has not been systematically addressed before.

We for the first time consider the problem of recovering sparse signals that contain corruptions and formulate it into a combined CS and GT problem (Section 2). We provide bounds of the number of measurements needed to recover such signals (Theorem 2) and compare it with CS and GT (Table 1). We provide explicit measurement construction method as well as efficient recovery algorithms (Section 3). When the measurements are erroneous, the number of measurements needed is also characterized (Section 4).

2. PROBLEM FORMULATION

Let $\mathbf{x} \in \bar{\mathbb{R}}^n$ ($\bar{\mathbb{R}} = \mathbb{R} \cup \infty$) denote the unknown signal to recover. ∞ indicates a corruption. Let $\mathbf{y} = M\mathbf{x}$ denote obtained measurements, where $M^{m \times n}$ is the measurement matrix. If $x_j = \infty$, then $y_i = \infty$ for all i such that $M_{ij} \neq 0$. Set $\mathcal{F}(\mathbf{y}) := \{i \in [m] : y_i = \infty\}$ denotes corrupted measurements. Note that in conventional CS, $\mathbf{x}, \mathbf{y}$ are real vectors.

$[q]$ ($q \in \mathbb{N}$) represents the set $\{1, \dots, q\}$. For set $S \subseteq [q]$, $|S|$ denotes its cardinality, and S^c denotes its complimentary set in $[q]$. Given $S \subseteq [n]$ and M , let $\mathcal{N}(S) := \{i \in [m] : \exists j \in S, \text{ s.t. } M_{ij} \neq 0\}$ denote the set of indices of measurements that passes through at least one entry in S , let $\mathcal{N}^c(S)$ denote its complimentary set. For set A and B , $A \cup B$ denotes the union and $A \setminus B$ contains elements that are in A but not in B . Given matrix M , M_{AB} denotes the submatrix with row indices in A and column indices in B .

Definition 1. $\mathbf{x} \in \bar{\mathbb{R}}^n$ is d -corrupted k -sparse (simplified as (d,k) -type) if $|S| \leq d$ and $|T| \leq k$, where $S = \{j \mid x_j = \infty\}$ and $T = \{j \mid 0 < |x_j| < \infty\}$.

Definition 2. Matrix $M^{m \times n}$ is called (d,k) -type identifiable if and only if for every two (d,k) -type vectors $\mathbf{x}$ and $\mathbf{z}$ such

Table 1. Number of nonadaptive measurements

d corruptions	$O(d^2 \log n)$ [21, 23]
k -sparse real signals	$O(k \log n)$ [1, 2, 4]
(d, k) -type signals	$O((d+1) \max(d, k) \log n)$ (here)

that $\mathbf{x} \neq \mathbf{z}$, it holds that $M\mathbf{x} \neq M\mathbf{z}$.

A (d, k) -type vector indicates that there are at most d failed links and at most k links with nonzero transmission delays. Throughout the paper, we consider the “for all” performance that requires M to identify all (d, k) -type vectors.

For example, consider matrix

$$M = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \end{bmatrix}$$

One can check that M can identify all 2-sparse signals in $\mathcal{R}^5$ when there is no corruption, i.e., M is $(0, 2)$ -type identifiable. However, when there exists one corruption, e.g., $\mathbf{x} = [0, 0.5, \infty, 0, 0]^T$, we have $\mathbf{y} = M\mathbf{x} = [0.5, \infty, \infty, 0]^T$. Although from $\mathbf{y}$ and M , we can infer that $x_3 = \infty$ and locate the corruption, we cannot decide whether $x_1 = 0.5$ or $x_2 = 0.5$. Thus, M is not $(1, 1)$ -type identifiable.

When $d = 0$, the problem reduces to CS problem where one aims to recover k -sparse signals. When $k = 0$, it reduces to the GT problem where one wants to locate d failures from boolean measurements. Here we need to not only locate the d corruptions but also recover the uncorrupted values, among which at most k entries are non-zero. Table 1 compares our result here on the number of measurements needed with existing results in GT and CS. The results in GT and CS can be viewed as special cases for our generalized result.

We remark that the Choir code in [24] is $(1, k)$ -identifiable for some k . But the construction is not directly extendable to general d and does not attempt to reduce m ($m = n$ in [24]). Here for any given n , d , and k , we want to design (d, k) identifiable M with m as small as possible.

We first introduce disjunct matrices in GT and expanders in CS that will be useful for our analysis.

Definition 3 (Disjunct matrices). M is called (d, e) -disjunct if for every $S \subseteq [n]$ and every $i \in [n]$ such that $i \notin S$ and $|S| \leq d$, $|\mathcal{N}(i) \setminus \mathcal{N}(S)| > e$ holds.

A $(d, 0)$ -disjunct matrix is called d -disjunct for simplification. One can locate up to d corruptions with a d -disjunct matrix [21]. The locating algorithm is simple [21]: x_j is identified to be corrupted if and only if $\mathcal{N}(j) \subseteq \mathcal{F}(y)$.

We remark that a $(d + 2k)$ -disjunct matrix is a (d, k) -type identifiable, and one can construct a $(d + 2k)$ -disjunct matrix with $O((d + 2k)^2 \log n)$ measurements [21, 23]. This number is larger than our result in Table 1 when $k \gg d$. We focus on the region that $n \gg k \gg d$ in this paper.

Given 0-1 matrix M , let $\gamma_l = \min_{j \in [n]} |\mathcal{N}(j)|$ and $\gamma_u = \max_{j \in [n]} |\mathcal{N}(j)|$ denote the minimum and maximum number of non-zero entries in a column. Given $T \subseteq [n]$, $E(T) := \sum_{j \in T} |\mathcal{N}(j)|$ measures the total number of nonzero entries in the columns in T .

Definition 4 (Expander). M corresponds to a $(k, \delta, \gamma_l, \gamma_u)$ -expander ($\delta \in (0, 1)$) if $|\mathcal{N}(T)| \geq (1 - \delta)E(T)$ for every $T \subseteq [n]$ with $|T| \leq k$.

We say M corresponds to a (k, δ, γ) -expander if $\gamma_l = \gamma_u = \gamma$. If M corresponds to a $(2k, \delta, \gamma_l, \gamma_u)$ -expander for $\delta\gamma_u/\gamma_l < 1/6$ [25], then one can correctly recover k -sparse signals via ℓ_1 -minimization, which returns the vector with the least ℓ_1 -norm among all the vectors that can produce the obtained measurements. There exist both random and explicit constructions of expanders.

Proposition 1. [26] For any $1 \leq k \leq n/2$, $\epsilon > 0$, one can explicitly construct a (k, ϵ, γ) -expander with $m = k\gamma/\epsilon^{O(1)}$ and $\gamma = 2^{O((\log(\log(n)/\epsilon))^3)}$.

3. RECOVERY OF CORRUPTED SPARSE SIGNALS

In network tomography, $M^{m \times n}$ is naturally a 0-1 matrix since a path delay measurement is an aggregate sum of the corresponding link delays. A lower bound of m for 0-1 matrix M to be (d, k) -type identifiable is stated as follows.

Proposition 2. A 0-1 (d, k) -type identifiable matrix M has at least $[d \log(n/d) + k \log((n - d)/k)] / \log(k + 2)$ rows.

Proof. Consider (d, k) -type vectors that all the non-zero finite values are ‘1’. There are $A := \sum_{i=1}^d \sum_{j=1}^k \binom{n}{i} \binom{n-i}{j}$ such vectors. In this case, each measurement could be an integer from 0 to k , or ∞ . There are at most $B := (k + 2)^m$ possible outcomes. We need $B \geq A$, and the claim follows. $\square$

Next we consider the upper bounds of the number of measurements needed. We start with a sufficient condition for (d, k) -type identifiable matrices.

Definition 5. M is called as $\mathcal{G}(d, 2k, \delta, \gamma_l, \gamma_u)$ if for every $S \subseteq [n]$ with $|S| \leq d$, there exists $G \subseteq \mathcal{N}^c(S)$ such that the submatrix $M' = M_{GS^c}$ is a $(2k, \delta, \gamma_l, \gamma_u)$ -expander.

Theorem 1. A $\mathcal{G}(d, 2k, \delta, \gamma_l, \gamma_u)$ matrix M is (d, k) -type identifiable if $\delta\gamma_u/\gamma_l < 1/6$.

Proof. Since $\gamma_l > 0$, M is d -disjunct, and one can correctly identify up to d corruptions. Since there always exist some uncorrupted measurements that correspond to a $(d, 2k, \gamma_l, \gamma_u)$ -expander, then all the real-valued entries can be correctly recovered via ℓ_1 -minimization. $\square$

One important property for $\mathcal{G}(d, 2k, \delta, \gamma_l, \gamma_u)$ matrices is that we have a polynomial algorithm for recovering $\mathbf{x}$. The

Algorithm 1 Recovery algorithm for error free case**Input:** y, M

- 1 $\forall i, x_i$ is identified as corrupted iff $\mathcal{N}(i) \subseteq \mathcal{F}(y)$.
- 2 Let D be the set of identified corruptions. $R = \mathcal{N}^c(D)$.
- 3 $\mathbf{x}_r = \arg \min_{\mathbf{z}} \|\mathbf{z}\|_1$ s.t. $M_{RD^c} \mathbf{z} = \mathbf{y}_R$.
- 4 **Return:** Corruptions D , uncorrupted values $\mathbf{x}_r$.

recovery algorithm is summarized in Algorithm 1. It first applies the identification algorithm in GT to locates up to d corruptions, which takes time $O(nm)$. Then it recovers the uncorrupted real values with ℓ_1 -minimization, which has running time $O(n^3)$. For comparison, a combinatorial search algorithm to recover $\mathbf{x}$ takes time $O(n^{O(d+k)})$.

Now we present one main result regarding the number of measurements needed for M to be $\mathcal{G}(d, 2k, \delta, \gamma_l, \gamma_u)$.

Theorem 2. $M^{m \times n}$ is a 0-1 matrix with i.i.d. entries M_{ij} , and $P(M_{ij} = 1) = p = \delta / \max(d, 2k)$, where constant $\delta \in (0, (\sqrt{73} - 7)/12)$. If $m = O(d \max(d, k) \log n)$, then with probability $1 - o(1)$, M is $\mathcal{G}(d, 2k, \delta, \gamma_l, \gamma_u)$ with $\delta \gamma_u / \gamma_l < 1/6$, and is thus (d, k) -type identifiable.

Proof. Pick any $\epsilon \in (0, 1)$. Let I_m denote the event that for every set $S \subseteq [n]$ with $|S| \leq d$, it holds that (1) $|\mathcal{N}^c(S)| \geq (1 - \epsilon)(1 - p)^d m$, and (2) for a fixed $G \subseteq \mathcal{N}^c(S)$ with $|G| = (1 - \epsilon)(1 - p)^d m$, M_{GS^c} corresponds to a $(2k, \delta, (1 - \delta)p|G|, (1 + \delta)p|G|)$ -expander.

Since $\delta(1 + \delta)/(1 - \delta) < 1/6$ from the assumption, clearly if I_m happens, the claim holds. We will prove that when m is as stated, $\Pr[I_m]$ goes to 1 as n goes to infinity.

Given S with $|S| = s$, let F_s denote the event that $|\mathcal{N}^c(S)| \geq (1 - \epsilon)(1 - p)^d m$ holds. Given $G \subseteq \mathcal{N}^c(S)$ with $|G| = r$, let E_r denote the event that M_{GS^c} corresponds to a $(2k, \delta, (1 - \delta)p|G|, (1 + \delta)p|G|)$ -expander. Since M has i.i.d. entries, once s and r are fixed, $\Pr[F_s]$ and $\Pr[E_r]$ do not depend on S and G . From the union bound,

$$\Pr[I_m^c] \leq \sum_{s=1}^d \binom{n}{s} (Pr[F_s^c] + Pr[E_{(1-\epsilon)(1-p)^d m}^c]). \quad (1)$$

We will next calculate $\Pr[I_m^c]$. The following form of Chernoff bound [27] is applied in our analysis.

Lemma 1. Let X be the sum of n independent random variables $x_i \in \{0, 1\}$, and let μ be its expectation. $\forall \delta \in (0, 1)$,

$$\Pr[X > (1 + \delta)\mu] \leq e^{-\frac{\delta^2 \mu}{3}}, \text{ and } \Pr[X < (1 - \delta)\mu] \leq e^{-\frac{\delta^2 \mu}{2}}.$$

Given S with $|S| = s \leq d$, from Lemma 1, we have

$$\Pr[F_s^c] \leq \Pr[F_d^c] \leq e^{-\epsilon^2(1-p)^d m/2}. \quad (2)$$

Given $G \subseteq \mathcal{N}^c(S)$ with $|G| = r$, let D_r denote the event that the number of nonzero entries in every column of M_{GS^c} is in $[(1 - \delta)pr, (1 + \delta)pr]$. From Lemma 1 and the Union bound,

$$\Pr[D_r^c] \leq ne^{-\delta^2 pr/3} + ne^{-\delta^2 pr/2} \leq 2ne^{-\delta^2 pr/3}. \quad (3)$$

Given $T \subseteq S^c$ with $|T| = t$, from Lemma 1, we have

$$\Pr[|\mathcal{N}(T)| \leq (1 - \delta/8)[1 - (1 - p)^t]r] \leq e^{-\frac{-\delta^2[1-(1-p)^t]r}{128}}, \quad (4)$$

$$\Pr[E(T) \geq (1 + \delta/8)ptr] \leq e^{-\delta^2 ptr/192}, \quad (5)$$

where $\mathcal{N}(T)$ and $E(T)$ are defined respect to matrix M_{GS^c} . Since $p = \delta / \max(d, 2k)$, through Taylor expansion, one can check that for all $t \leq 2k$, it holds that

$$(1 - \delta/8)[1 - (1 - p)^t]r \geq (1 - \delta)(1 + \delta/8)ptr. \quad (6)$$

From (4) to (6) and the Union bound, we have

$$\Pr[|\mathcal{N}(T)| \leq (1 - \delta)E(T), \text{ given } T] \leq e^{-c_1 \delta^2 p|T|r}, \quad (7)$$

where c_1 is a constant independent of δ_1, p, k , and n .

From the Union bound, we have

$$\begin{aligned} \Pr[E_r^c] &\leq \Pr[D_r^c] + \sum_{t=1}^{2k} \binom{n}{t} \Pr[|\mathcal{N}(T)| \leq (1 - \delta)E(T), \\ &\quad \text{given } T \text{ with } |T| = t] \\ &\leq 2ne^{-\delta^2 pr/3} + \sum_{t=1}^{2k} e^{t(\log(n/t)+1) - c_1 \delta^2 ptr}, \end{aligned} \quad (8)$$

where the second inequality follows from (3) and (7).

Plugging (8) and (2) into (1), we have $\Pr[I_m^c] \rightarrow 0$ when $n \rightarrow \infty$, provided that

$$m \geq 2(d \log(n/d) + \log n)/(p(1 - \epsilon)(1 - p)^d \delta^2).$$

Since $p = \delta / \max(d, 2k)$, then $(1 - p)^d \geq 1/4$. Then when

$$m \geq 8 \max(d, 2k)(d \log(n/d) + \log n)/\delta^3,$$

with probability $1 - o(1)$, M is (d, k) -type identifiable. $\square$

Theorem 2 indicates that a randomly generated 0-1 matrix with $O((d + 1) \max(d, k) \log n)$ measurements is (d, k) -type identifiable with high probability. We compare this result with exiting ones in CS and GT in Table 1. We next provide an explicit measurement construction method based on expanders.

Theorem 3. M is (d, k) -type identifiable if it corresponds to a $(d + 2k, \epsilon/d, \gamma)$ -expander with $\frac{\epsilon(d+1)}{d - \epsilon(d+1)} < 1/6$.

Proof. $\forall S$ with $|S| = s \leq d$, and $\forall T \subseteq S^c$ with $|T| = t \leq 2k$, from the expansion property and $|\mathcal{N}(S)| \leq s\gamma$, we have

$$|\mathcal{N}(S \cup T)| - |\mathcal{N}(S)| \geq (1 - \epsilon(d + 1)/d)t\gamma. \quad (9)$$

Then the number of non-zero entries in each column i in $M_{\mathcal{N}^c(S)S^c}$ is between $(1 - \epsilon(d + 1)/d)\gamma$ and γ .

Since $M_{\mathcal{N}^c(S)T}$ has at most $t\gamma$ non-zero entries, from (9) one can check that $M_{\mathcal{N}^c(S)S^c}$ is a $(2k, \frac{\epsilon(d+1)}{d}, (1 - \frac{\epsilon(d+1)}{d})\gamma, \gamma)$ -expander. The claim follows. $\square$

From Theorem 3 with Proposition 1, an explicit construction of (d, k) -type identifiable matrix uses $O(d^{O(1)}(d + 2k)2^{O((d \log(\log(n/\epsilon)))^3)})$ measurements, and this number is larger than that in Theorem 2 with random construction.

Algorithm 2 Recovery algorithm for up to h errors

Input: y, M

- 1 For each $S \subseteq [n]$ with $|S| \leq d$, if $|\mathcal{N}(S) \setminus \mathcal{F}(y)| + |\mathcal{F}(y) \setminus \mathcal{N}(S)| \leq h$, S is the set of corruptions, denoted by D .
 - 2 $R = [n] \setminus (\mathcal{N}(D) \cup \mathcal{F}(y))$, $M_r = M_{RD^c}$.
 - 3 $\mathbf{x}_r = \arg \min_{\mathbf{z}} \|\mathbf{z}\|_1$ s.t. $M_r \mathbf{z} = \mathbf{y}_R$.
 - 4 **Return:** Corruptions D , uncorrupted values $\mathbf{x}_r$.
-

4. ERRONEOUS MEASUREMENTS

We next consider the case the measurements contain errors. Let α_i^T denote the i th row of M . We consider two types of errors: (1) $y_i = \infty$ when $\alpha_i^T \mathbf{x} \in \mathcal{R}$, and (2) $y_i \in \mathcal{R}$ when $\alpha_i^T \mathbf{x} = \infty$. We assume that the total number of these two types of errors is at most h , and these errors can happen at arbitrary unknown locations. The goal is to design measurement matrix M such that all (d, k) -type signals can be correctly recovered *no matter where the h errors are*.

If M is a $(d, 2h + 1)$ -disjunct matrix, then one can identify d corruptions in the presence of at most h errors [21]. Then one sufficient condition for identifying (d, k) -type signals from measurements that contain h errors is as follows,

Proposition 3. *If M is $(d, 2h + 1)$ -disjunct, and for every S of up to d corruptions and for every H with up to h errors, there exists $G \subseteq [m] \setminus (\mathcal{N}(S) \cup H)$ s.t. M_{GS^c} corresponds to a $(2k, \delta, \gamma_l, \gamma_u)$ -expander with $\delta\gamma_u/\gamma_l < 1/6$, then all (d, k) -type signals can be recovered in the presence of h errors.*

The proof follows clearly from previous discussions and is skipped. The recovery algorithm for such matrices is stated in Algorithm 2. We prove that these matrices can be obtained through random construction with the same probability p as that in Theorem 2. The bound of the number of measurements needed is as follows.

Theorem 4. *One can identify all (d, k) -type signals from $m = O(\max(d, k)(d \log n + h \log(\max(d, k)) + h \log \log n))$ measurements that contain h errors in arbitrary locations.*

Proof. The proof follow the same line as that for Theorem 2, and we skip the details. Let $\hat{I}_m$ denote the event that for every set S with $|S| \leq d$ and every set H with $|H| \leq h$, (1) $\mathcal{N}^c(S) \geq (1 - \epsilon)(1 - p)^d m$, (2) $M_{\mathcal{N}^c(S)S^c}$ has at least $(1 - \delta)p(1 - \epsilon)(1 - p)^d m$ nonzero entries in each column, (3) $M_{G'S^c}$ corresponds to a $(2k, \delta, (1 - \delta)pr', [(1 + \delta_2)pr'])$ -expander for a fixed G' in $\mathcal{N}^c(S) \setminus H$ with $|G'| = r' = (1 - \epsilon)(1 - p)^d m - h$ entries. If $\hat{I}_m$ happens, and if it holds that

$$2h + 1 \leq (1 - \delta)p(1 - \epsilon)(1 - p)^d m, \quad (10)$$

then one can identify all (d, k) -type signals from m measurements that contain at most h errors at arbitrary locations.

One can check that $\Pr[\hat{I}_m^c] \rightarrow 0$ when $n \rightarrow \infty$ provided that m is as stated in the Theorem. And (10) follows for this choice of m . Then the claim follows. $\square$

5. NUMERICAL RESULTS




Fig. 1. Identification of corruptions




Fig. 2. Recovery of sparse signals with corruptions

We fix $n = 1000$ and the number of ‘1’s in each column of M to be 5, and randomly generate a 0-1 matrix M with $m = 600$ and $m = 800$ respectively. We first consider the performance of identifying corruptions in Fig. 1. For each d , we randomly choose the locations of the corruptions, and the results are averaged over 500 runs. When $m = 600$, 10 corruptions can be correctly identified.

In Fig. 2, we fix d and increase the number of non-zero entries k . The locations of corruptions and non-zero entries are randomly chosen, and non-zero entries are sampled as i.i.d. Gaussian random variables. Algorithm 1 is applied to recover (d, k) -type signals. Let $\mathbf{x}^*$ contain the uncorrupted entries, and let $\mathbf{x}_r$ denote our reconstruction. $\|\mathbf{x}^* - \mathbf{x}_r\|_2 / \|\mathbf{x}^*\|_2$ is the normalized recovery error of the uncorrupted part. The results are averaged over 100 runs. When $m = 600$, we can recover all $(5, 220)$ -type signals or all $(10, 200)$ -type signals.

6. CONCLUSION

We considered recovering sparse link delay values from path delay measurements in the presence of link failures and for the first time formulated it into a CS problem with corruptions. We provided bounds of the number of nonadaptive measurements needed to identify both corruptions and real entries. Explicit constructions and efficient recovery algorithms are also provided. One ongoing work is to explore construction methods with fewer measurements.

7. REFERENCES

- [1] E. Candès and T. Tao, “Decoding by linear programming,” *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [2] E. Candès and T. Tao, “Near-optimal signal recovery from random projections: Universal encoding strategies?,” *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [3] D. Donoho, “Compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [4] D. Donoho, “High-dimensional centrally symmetric polytopes with neighborliness proportional to dimension,” *Discrete Comput. Geom.*, vol. 35, no. 4, pp. 617–652, 2006.
- [5] T. Bu, N. Duffield, F. L. Presti, and D. Towsley, “Network tomography on general topologies,” in *Proc ACM SIGMETRICS*, 2002, pp. 21–30.
- [6] A. Coates, A.O. Hero III, R. Nowak, and B. Yu, “Internet tomography,” *IEEE Signal Process. Mag.*, vol. 19, no. 3, pp. 47–65, 2002.
- [7] Y. Chen, D. Bindel, H. H. Song, and R.H. Katz, “Algebra-based scalable overlay network monitoring: Algorithms, evaluation, and applications,” *IEEE/ACM Trans. Netw.*, vol. 15, no. 5, pp. 1084–1097, 2007.
- [8] N. Duffield, “Network tomography of binary network performance characteristics,” *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5373–5388, 2006.
- [9] A. Gopalan and S. Ramasubramanian, “On identifying additive link metrics using linearly independent cycles and paths,” *IEEE/ACM Trans. Netw.*, vol. 20, no. 3, pp. 906–916, 2012.
- [10] H. X. Nguyen and P. Thiran, “Using end-to-end data to infer lossy links in sensor networks,” in *Proc. IEEE INFOCOM*, 2006, pp. 1–12.
- [11] M. Coates, Y. Pointurier, and M. Rabbat, “Compressed network monitoring for ip and all-optical networks,” in *Proc. ACM SIGCOMM IMC*, 2007, pp. 241–252.
- [12] J. Haupt, W.U. Bajwa, M. Rabbat, and R. Nowak, “Compressed sensing for networked data,” *IEEE Signal Process. Magazine*, vol. 25, no. 2, pp. 92–101, 2008.
- [13] W. Xu, E. Mallada, and A. Tang, “Compressive sensing over graphs,” in *Proc. IEEE INFOCOM*, 2011.
- [14] M. Wang, W. Xu, E. Mallada, and A. Tang, “Sparse recovery with graph constraints: Fundamental limits and measurement construction,” in *Proc. IEEE INFOCOM*, 2012, pp. 1871–1879.
- [15] P. Babarczi, J. Tapolcai, and P. Ho, “Adjacent link failure localization with monitoring trails in all-optical mesh networks,” *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 907–920, 2011.
- [16] M. Cheraghchi, A. Karbasi, S. Mohajer, and V. Saligrama, “Graph-constrained group testing,” *IEEE Trans. Inf. Theory*, vol. 58, no. 1, pp. 248–262, 2012.
- [17] N.J.A. Harvey, M. Patrascu, Yonggang Wen, S. Yekhanin, and V.W.S. Chan, “Non-adaptive fault diagnosis for all-optical networks via combinatorial group testing on graphs,” in *Proc. IEEE INFOCOM*, 2007, pp. 697–705.
- [18] J. Tapolcai, B. Wu, P. Ho, and L. Rónyai, “A novel approach for failure localization in all-optical mesh networks,” *IEEE/ACM Trans. Netw.*, vol. 19, pp. 275–285, 2011.
- [19] R. Dorfman, “The detection of defective members of large populations,” *Ann. Math. Statist.*, vol. 14, pp. 436–440, 1943.
- [20] M. Cheraghchi, A. Hormati, A. Karbasi, and M. Vetterli, “Group testing with probabilistic tests: Theory, design and application,” *IEEE Trans. Inf. Theory*, vol. 57, no. 10, pp. 7057–7067, 2011.
- [21] D. Du and F. K. Hwang, *Combinatorial Group Testing and Its Applications (Applied Mathematics)*, World Scientific Publishing Company, second edition, 2000.
- [22] A.C. Gilbert, B. Hemenway, A. Rudra, M.J. Strauss, and M. Wootters, “Recovering simple signals,” in *Proc. ITA*, 2012, pp. 382–391.
- [23] E. Porat and A. Rothschild, “Explicit non-adaptive combinatorial group testing schemes,” in *Proc. ICALP*, pp. 748–759, 2008.
- [24] L. Applebaum, W.U. Bajwa, R. Calderbank, and S. Howard, “Choir codes: Coding for full duplex interference management,” in *Allerton*, 2011, pp. 1–8.
- [25] R. Berinde, A. Gilbert, P. Indyk, H. Karloff, and M. Strauss, “Combining geometry and combinatorics: a unified approach to sparse signal recovery,” *arxiv:0804.4666*, 2008.
- [26] M. Capalbo, O. Reingold, S. Vadhan, and A. Wigderson, “Randomness conductors and constant-degree lossless expanders,” in *Proc. ACM STOC*, 2002, pp. 659–668.
- [27] M. Mitzenmacher and E. Upfal, *Probability and computing: Randomized algorithms and probabilistic analysis*, Cambridge University Press, 2005.