Sparse Group Testing Codes for Low-Energy Massive Random Access

Huseyin A. Inan, Peter Kairouz and Ayfer Ozgur
Stanford University
Email: {hinan1, kairouzp, aozgur}@stanford.edu

***Abstract*—We present random access protocols for machine-type communication where a massive number of low-energy wireless devices want to occasionally transmit short information packets. We focus on the device discovery problem, with extensions to joint discovery and data transmission as well as data transmission without communicating the device identities. We formulate this problem as a combinatorial group testing problem, where the goal is to exactly identify the set of at most d defective items from a pool of n items. We translate the energy constraint at the wireless physical layer to a constraint on the number of tests each item can participate in, and study the resulting "sparse" combinatorial group testing problem.**

The celebrated result for the combinatorial group testing problem is that the number of tests t can be made logarithmic in n when $d = O(\log n)$. However, state-of-the-art group testing codes require the items to be tested $w = \Omega(\frac{d \log n}{\log d + \log \log n})$ times. In our sparse setting, we restrict the number of tests each item can participate in by $w_{\max}$. We show that t decreases suddenly from n to $(d+1)\sqrt{n}$ when $w_{\max}$ is increased from d to $d+1$. If $w_{\max} = ld+1$ for any positive integer l such that $ld+1 \leq \sqrt[l+1]{n}$, we can achieve $t = (ld+1)n^{1/(l+1)}$. We also prove a nearly matching lower bound. These results reveal a favorable trade-off between energy and spectral efficiency for random access.

I. Introduction

A massive number of low energy wireless devices are envisioned to form the fabric of smart technologies and cyberphysical systems, enabling the Internet of Things (IoT). While this vision is expected to bring many business and social opportunities, it presents unprecedented challenges to wireless system and radio designers. One of these challenges is how to enable a massive number of devices to efficiently access the spectrum. Massive random access, in this emerging setting, differs significantly from that in more traditional systems, such as cellular and WiFi, due to a number of reasons.

First, the number of devices (or sensors) associated with a single access point can be much larger than that in traditional wireless systems. Further, each device can be sporadically active; i.e., has bursty traffic. Therefore, the number active devices at any given time can be much smaller than the total number of nodes in the system. Second, when active, each device can have only a few bits to communicate; for example, a single measurement value. Third, such wireless devices are severely energy-limited since they are expected to operate for many years on small batteries or harvest their energy from the environment (without relying on any traditional batteries). When the system size is large and the number of information bits to be communicated by each device (payload) is small, tasks such as discovering active devices, coordinating transmissions, resolving collisions, and estimating the channel can dominate the energy expenditure and throttle the network throughput.

In this paper, we tackle the following fundamental question: how do we allow a massive number of sporadically active low-energy wireless devices with small payloads to access the spectrum with minimal coordination and channel estimation overheads? Our design philosophy hinges on three principles: (a) combining the device discovery and data transmission phases, (b) embracing sensor collisions, (c) trading spectral efficiency for energy efficiency. The first principle allows us to combine device discovery, data transmission, and transmission without discovery in a single framework; the second eliminates the cost of collision resolution (feedback and retransmissions); the third principle allows us to discover a particularly favorable regime for achieving both energy and spectral efficiency.

A. Massive Random Access

We focus on the following model for random access: n devices (or sensors) are associated with a single access point. At most d of them can be active at any given time, where $n \gg d \gg 1$. We adopt the channel model in [1]: each device uses on-off signaling; i.e., it transmits a binary sequence of 0's and 1's. The access point simply detects whether or not there is energy on the channel in each time-slot. This leads to a (potentially noisy) Boolean OR-channel from the devices to the access point. This simple modulation and detection technique does not require any channel state information at the receiver, thus eliminating the need for channel training and estimation. To simplify the discussion, we mainly focus on the device discovery problem in this setting. However, we will argue in the next subsection that the solution we derive for this setting can be easily extended to jointly send device IDs and messages, and even to send messages without communicating the device IDs. Therefore, our problem formulation is as follows. Given n devices, design a length t binary signature for each device (i.e., $M_i \in \{0,1\}^t$ for $i = 1, \ldots, n$) such that from any $S \subset \{1, \ldots, n\}$ with $|S| \leq d$, we can exactly identify the set S (the active devices) from

$$Y = \bigvee_{i \in S} M_i.$$

Note that even though we require a zero-error code design, the energy detection at the receiver can be noisy in practice,

leading to a Boolean OR-channel with occasional bit flips of the output. The frequency of such errors will depend on the signal-to-noise ration (SNR) of the channels from each device to the access point. We ignore such errors in this paper. Note that this model also assumes that the active devices have synchronized transmissions. In practice, this can be achieved by sending a beacon signal from the access point and having devices synchronize their transmissions to the transmitted beacon signal.

B. Random Access and Group Testing

The problem statement above corresponds to a (non-adaptive) combinatorial group testing problem. The group testing problem consists of identifying a small set of d (or less) defective items from a large population of size n by performing tests on groups of items, rather than on individual items. For an unknown sequence $x \in \{0,1\}^n$ with at most d ones representing the defective items, we are allowed to test any subset $S \subseteq \{1,\ldots,n\}$ of the items. The result of a test S could either be positive, which happens when at least one item in S is defective (i.e., $\exists\ i \in S$ such that $x_i = 1$), or negative when all the items in S are not defective (i.e., $\forall i \in S$ we have $x_i = 0$). The goal is to design as few tests as possible so that we can exactly recover the unknown sequence x.

The original group testing framework was developed in 1943 by Robert Dorfman [2]. Back then, group testing was devised to identify which WWII draftees were infected with syphilis – without having to test them individually. In Dorfman's application, items represented draftees and tests represented actual blood tests. Over the years, group testing has found numerous applications, and many variations of the problem have been studied. Non-adaptive combinatorial group testing (CGT) refers to the fact that tests need to be designed ahead of time and the set of defective items needs to be recovered exactly. A non-adaptive CGT strategy can be represented by a $t \times n$ binary matrix M, where $M_{ij} = 1$ means that item j participates in test i. The test results vector Y is simply M multiplied by x, where linear summations are replaced by Boolean ORs. A necessary and sufficient condition for the design of a non-adaptive CGT strategy M is that of *separability*. A matrix M is *d-separable* if for any $x_1 \neq x_2$, d-sparse vectors, we have that $Mx_1 \neq Mx_2$. Note that this precisely corresponds to the device discovery problem defined above.

The celebrated result for the group testing problem is that t can be made logarithmic in n, $t = O(d^2 \log_d^2 n)$ in [3], $t = O(d^2 \log n)$ in [4], [5]. This implies that group testing can provide drastic gains when $d \ll n$, say $d = O(\log n)$, compared to the naive approach of testing every item individually which results in $t = n$. Therefore, state-of-the-art group testing codes can be used for device discovery to minimize the length of the signature sequences, which is equivalent to minimizing the completion time of the device discovery phase. More precisely, given a non-adaptive CGT matrix M, we can use it for the following three tasks:

- Device discovery: assign the i^{th} column M_i to device i as its signature sequence.
- Joint discovery and data transmission: assign a disjoint set of columns of M to each device to form a codebook for data transmission. By transmitting one of its assigned codewords, each device can communicate both its identity and data, provided that there are less than d simultaneously transmitting devices. For example, by assigning two columns of M to each of $n/2$ devices, the sensors can communicate their identity and one bit of information. In general, the size of the codebook can be different for each device.
- Data transmission with non-identifiable transmitters: use the columns of M at all devices for data transmission; i.e., each device uses the same codebook formed by the columns of M. By construction, as long as there are less than or equal to d simultaneously transmitting devices, codewords can be recovered at the receiver but not the identity of the transmitting devices. A similar setting has been recently considered in [6].

The connection between group testing and random access has been recognized in the literature, in the early works of [7], [8] and more recently in [1], which is more closely related to our paper. However, different from our combinatorial non-adaptive setting, [7], [8] consider node discovery via the adaptive group testing framework (they rely on feedback from the receiver for resolving collisions). [1] considers neighborhood discovery via stochastic group testing where a small error is allowed in the discovery of the neighbors.

C. Energy vs Spectral Efficiency

In most IoT applications, the length of the codewords, which determines the spectral efficiency of the system, is not the only performance metric of interest. Typically energy efficiency is even more critical since devices rely on tiny batteries. Therefore, while it is desirable to minimize the time needed to identify active devices, it is also desirable to achieve this in an energy efficient manner; i.e., to minimize the total energy spent by each device. The energy spent is proportional to the number of "1"s in the codeword transmitted by each device. This motivates us to study group testing with constraints on the total number of tests that can be performed on each item. Using our previous notation, this corresponds to imposing a constraint on the total number of "1"s in each column of M. Note that if energy efficiency were the only metric of interest in this application, we could have resorted to the trivial solution that tests every item individually. This leads to a single "1" in each column of M. However, this results in codewords of length n. Since both of energy and spectral efficiency are of interest, we aim to understand the fundamental trade-off between them. For example, can allowing for a few more "1"s in each codeword significantly reduce the length of the codewords?

In the rest of the paper, we focus on designing combinatorial group testing codes with constraints on the total number of tests that need to be performed on each item. While in

the random access setting, such constraints are motivated by energy limitations, the total number of tests that can be performed on each item can be limited due to different reasons in other applications. For instance, the amount of blood or genetic material available from an individual can limit the number of tests that this individual can participate in.

D. Technical Contributions

In this paper, we explore the "sparse" setting for the non-adaptive combinatorial group testing problem where we restrict the number of tests each item can participate in by $w_{\max}$. We first revisit existing constructions based on d-disjunct matrices that achieve $t = O(d^2 \log n)$ [3] and $t = O(d^2 \log_d^2 n)$ tests [4]. A $t \times n$ binary matrix M is called d-disjunct if none of the columns of M is covered by the Boolean sum of any other d columns of M. It is easy to see that the set of at most d defective items can be decoded uniquely if the test matrix M satisfies this property. It is also easy to see that any testing matrix M that allows to uniquely recover up to d defective items must be $(d-1)$-disjunct [9]. We show that these constructions require items to be tested $w = \Omega\left(\frac{d \log n}{\log d + \log\log n}\right)$ times. In the other extreme, we show that if $w_{\max} \leq d$ then d-disjunct matrices must have $t = n$; i.e., testing every item individually is optimal. A natural question then is the following: how does t decrease as we increase $w_{\max}$ beyond the bare minimum (up to its values in state-of-the-art constructions)? In particular, can we slightly increase $w_{\max}$ beyond d and significantly reduce t, the number of tests needed? Surprisingly, the answer is positive.

We show that when $w_{\max} = d + 1$, the number of tests decreases drastically from $t = n$ to $t = (d+1)\sqrt{n}$. More generally, if $w_{\max} = ld + 1$ for any positive integer l such that $ld + 1 \leq \sqrt[l+1]{n}$, we can achieve

$$t = (ld + 1)n^{\frac{1}{l+1}}.$$

This implies that the fractional power of n can be reduced drastically by increasing $w_{\max}$ linearly in d. This result is most significant when $d = O(\log n)$. We achieve this performance by introducing a simple modification to Kautz and Singleton's construction, which shows that the field size in this construction can be used to trade between t and w_{max}. We then prove a nearly matching lower bound which shows that

$$t = \Omega(d^{\frac{2}{l+1}} n^{\frac{1}{l+1}}).$$

In particular, this shows that Kautz and Singleton's construction is order optimal (even up to an almost matching constant) when $w_{\max} = d + 1$. Given that Kautz and Singleton's construction is strictly suboptimal in the classical group testing setting, this finding is surprising.

E. Paper Organization

The remainder of this paper is organized as follows. In Section II, we present the needed prerequisite material and describe two common combinatorial group testing constructions. We also show that both constructions require items to be tested $w = \Omega(\frac{d \log n}{\log d + \log\log n})$ times. The main results of our paper are formally presented in Section III. The proofs are deferred to Section IV. We provide, in Section V, a brief survey of important results on combinatorial group testing and a detailed comparison with a recent paper [10] that also considers the sparse setting, albeit in a stochastic rather than combinatorial framework. Finally, we conclude our paper in Section VI with a few interesting and nontrivial extensions.

II. Preliminaries

For any $t \times n$ matrix M, we use M_j to refer to its j'th column and M_{ij} to refer to its (i,j)'th entry. For an integer $m \geq 1$, we denote the set $\{1, \ldots, m\}$ by $[m]$. The Hamming weight of a row or a column of M will be simply referred to as the "weight" of the row or column.

A. Non-adaptive Combinatorial Group Testing

Our paper focuses on non-adaptive combinatorial group testing (CGT). A non-adaptive CGT strategy can be represented by a $t \times n$ binary matrix M, where $M_{ij} = 1$ means that item j participates in test i. We will occasionally refer to M as a group testing code (or codebook) and its i^{th} column M_i as the i^{th} codeword. The test results vector Y is simply M multiplied with by x, where linear summations are replaced by Boolean summations. A necessary condition for the design of a non-adaptive CGT strategy M is that of *separability*. A matrix M is *d-separable* if for any $x_1 \neq x_2$, d-sparse vectors, we have that $Mx_1 \neq Mx_2$. Unfortunately, the d-separability condition does not lead to tractable, explicit, and efficiently decodable constructions of M for an arbitrary value of n. To circumvent this issue, a stronger condition on M is needed. This condition is known as *d-disjunctiveness* [3], [11]. We say that a binary codeword M_1 covers a binary codeword M_2 if $M_1 \vee M_2 = M_1$. A binary matrix is *d-disjunct* if any Boolean sum of up to d columns of M does not cover any other column not included in the sum. The d-disjunctiveness property ensures that we can recover up to d summed codewords from their Boolean sum. This can be naively done using a *cover decoder*. A cover decoder simply scans through the columns of M, and checks whether or not the test results vector Y covers a particular column. If column i is covered by Y, item i is declared defective. When M is d-disjunct, the cover decoder succeeds at identifying all the defective items, while achieving a zero false positive rate. Interestingly, it turns out that $(d+1)$-separable matrices are also d-disjunct [9]. Therefore, even though disjunctiveness is stronger than separability, the two conditions are essentially equivalent.

We define $t(d,n)$ to be the smallest t needed for a binary $t \times n$ matrix M to be d-disjunct. Notice that naturally, $t(d,n) \leq n$ because we can always use the identity matrix $M = I_n$ to identify any $1 \leq d \leq n$ defectives among n items. A classical result in the non-adaptive CGT literature shows that $t(d,n) \geq \Omega(d^2 \log_d n)$ [5], [12], [13]. Several explicit and randomized constructions of d-disjunct matrices have been developed over the past 50 years with the most efficient known constructions achieved $t = O(d^2 \log n)$ [4], [5], [14].

B. Relevant Lower Bounds

We now summarize two known lower bounds on the minimum number of tests. These bounds imply that individual testing is necessary whenever $d = \Omega(\sqrt{n})$ or $w_{\max} \leq d$, where $w_{\max}$ is the maximum number of tests an item participates in.

Proposition 1: For all n and d, the following bound on $t(d, n)$ holds

$$t(d, n) \geq \min\left\{ \binom{d+2}{2}, n \right\}.$$

Proposition 1 suggests that we need $d = O(\sqrt{n})$ to be able to design a d-disjunct matrix with $t < n$. The proof of this proposition is due to D'yachkov and Rykov, and can be found in [11].

Proposition 2: If $w_{\max} \leq d$, then $t(d, n) = n$.

The above proposition shows that one cannot do better than individual testing when the maximum number of tests an item can participate in is less than or equal to d. Hence, we focus our attention on a setting where $w_{\max} \geq d + 1$ and show that $t(d, n)$ suddenly transitions from n to $\sqrt{n}$ when $w_{\max} = d+1$. The proof of Proposition 2 can be found in [9].

C. Disjunct Matrices via Error Correcting Codes

A q-nary error-correcting code is a code whose codewords consist of q basic symbols [15]. Binary codes are a special case of q-nary codes with $q = 2$. Consider a q-nary code with $n = q^k$ codewords of length $t = k + r$. Denoting the minimum distance between the codewords as $d_{\min}$, one can show that $d_{\min} \leq r + 1$ from the following observation. Fix any k positions in the codewords. If any two codewords have the same symbols in these positions, then it must be the case that $d_{\min} \leq r$. Otherwise, we must observe all possible q^k sequences in the k fixed positions. In this case, some of the codewords will differ by only one position on the fixed k positions. Hence, $d_{\min} \leq r + 1$. We state this formally in the following theorem [16].

Theorem 1: A q-nary code with $n = q^k$ codewords of length $t = k + r$ must satisfy $d_{\min} \leq r + 1$.

Codes with $d_{\min} = r + 1$ and $n = q^k$ are called maximum distance separable (MDS) codes [16]. Reed-Solomon codes [17] are a known class of MDS codes with the constraint that $q \geq t$. When concatenated with a nonlinear code, Reed-Solomon codes lead to d-disjunct group testing codes. In what follows, we will use the subscript q in the parameters of the Reed-Solomon codes to separate them from the group testing codes that will be constructed shortly. To recap, Reed-Solomon codes achieve a minimum distance of $d_{\min} = r_q + 1$ with a code length of $t_q = k_q + r_q$ and a number of codewords equal to $n_q = q^{k_q}$, provided that $q \geq t_q$ and q is any prime power.

We can convert a Reed-Solomon code into a group testing code using the following method introduced by Kautz and Singleton in [3]. We replace each codeword symbol $i \in \{1, 2, \ldots, q\}$ by e_i, a length q binary sequence with a single nonzero entry in the i^{th} position. Thus, a Reed-Solomon code is transformed into a binary code of length $t = qt_q$ by concatenating it with an "identity code." The minimum distance of the resultant binary code is double that of the Reed-Solomon code; i.e., $d_{\min} = 2(r_q + 1)$. This is because any two distinct q-nary symbols will differ in two positions in their corresponding length q binary sequences. Note that the number of codewords remains the same $n = n_q = q^{k_q}$, and all the binary codewords have the same weight $w = t_q$. This construction will be referred to as the Kautz and Singleton construction.

Consider a binary code M with minimum codeword weight of $w_{\min}$. We define $\lambda_{\max}$ to be the maximum number of overlapping ones between any two codewords in M. In the coding theory literature, $\lambda_{\max}$ is commonly referred to as the maximal correlation of M. A central result in group testing demonstrates that M is d-disjunct as long as $\lambda_{\max} d + 1 \leq w_{\min}$. This can be seen from the following simple argument. Take any $d+1$ codewords and fix one codeword among them. The number of overlapping ones between the fixed codeword and the rest of the codewords is at most $d\lambda_{\max}$. Since the minimum weight satisfies $w_{\min} \geq d\lambda_{\max} + 1$, this codeword cannot be covered by the rest of the codewords. Thus, M must be at least d-disjunct. We state this formally in the following lemma.

Lemma 1: A binary code M with codewords of minimum weight $w_{\min}$ and maximal correlation $\lambda_{\max}$ is $\left\lfloor \frac{w_{\min}-1}{\lambda_{\max}} \right\rfloor$-disjunct.

Observe that in Kautz and Singleton's construction, we have that

$$\lambda_{\max} = w - d_{\min}/2 = t_q - r_q - 1 = k_q - 1.$$

Therefore, Kautz and Singleton's construction provides us with a group testing code that is $\left\lfloor \frac{t_q-1}{k_q-1} \right\rfloor$-disjunct.

Theorem 2: Kautz and Singleton's construction provides a $t \times n$ d-disjunct matrix where $t = O(d^2 \log_d^2 n)$ with constant column weight $w = \Omega\left(\frac{d \log n}{\log d + \log\log n}\right)$ and constant row weight $\rho = \Omega\left(\frac{n}{d \log_d n}\right)$.

Proof: To obtain a d-disjunct code using Kautz and Singleton's construction, we set $t_q = q$, and choose q and k_q such that $d = \left\lfloor \frac{q-1}{k_q-1} \right\rfloor$. Note that $n = q^{k_q}$ and $q = \Theta(dk_q)$. Hence, $q = \Theta(d \log_q n)$ or $q \log q = \Theta(d \log n)$. Since $q \geq d$, we get that $q = O(d \log_d n)$. Note that $t = qt_q = q^2$, therefore $t = O(d^2 \log_d^2 n)$. The corresponding binary code has constant column weights $w = t_q = q$. Note that $q \log q = \Theta(d \log n)$ is related to the famous Lambert W function [18] and using $W(x) \geq \log x - \log\log x$, we get $q = \Omega\left(\frac{d \log n}{\log(d \log n)}\right)$ or equivalently $w = \Omega\left(\frac{d \log n}{\log d + \log\log n}\right)$. From our earlier discussion on MDS codes, recall that achieving a minimum distance of $r_q + 1$ requires that for any arbitrary k_q rows of this code, the chosen $k_q \times q^{k_q}$ matrix must include all q^{k_q} possible assignments of q-nary symbols in the columns. It follows that any row of a Reed-Solomon code must include every q-nary symbol an equal number of times. More precisely, each q-nary symbol must present q^{k_q-1} times in all rows. Therefore, the corresponding binary code has a constant row

weight of $\rho = n/q$. Since $q = O(d \log_d n)$, it follows that $\rho = \Omega\left(\frac{n}{d \log_d n}\right)$. ∎

A different line of work introduced by Porat and Rothscheld in [4] constructs $t \times n$ d-disjunct matrices with $t = O(d^2 \log n)$. Their approach is based on q-nary codes that meet the Gilbert-Varshamov bound where the alphabet size is $q = \Theta(d)$. As in the Kautz and Singleton construction, their inner code is the identity code. The resulting binary code has the property that all the columns have the same weight of $w = \Theta(d \log n)$. Furthermore, the maximum row weight satisfies $\rho_{\max} = \Omega(n/d)$.

Theorem 3: The explicit construction by Porat and Rothscheld in [4] achieves a $t \times n$ d-disjunct matrix where $t = O(d^2 \log n)$ with constant column weight $w = \Theta(d \log n)$ and maximum row weight $\rho_{\max} = \Omega(n/d)$.

Compared to Kautz and Singleton's construction, one can observe that Porat and Rothscheld's construction is better in the regime where $d = O(\log n)$. However, Kautz and Singleton's construction is better when $d = \Theta(n^\alpha)$ for some constant $\alpha \in (0, 1/2)$. In this regime, Kautz and Singleton's construction meets the fundamental lower bound (see Section V); therefore it is order optimal. Nevertheless, the regime where $d = \Theta(n^\alpha)$ has not received considerable attention in the group testing literature because t cannot scale logarithmically in n when $d = \Theta(n^\alpha)$. This is why we focus our attention on the more common $d = O(\log n)$ regime in which Kautz and Singleton's construction is strictly suboptimal.

III. Our Results

In this section, we formally present our results and discuss their implications. The detailed proofs are deferred to section IV.

We focus on a model where each item can participate in a limited number of tests. This is equivalent to restricting the codewords (columns of M) to have a limited number of "1"s. Recall, from the discussion in Section II, that if the codewords have a Hamming weight that is bounded by d, one cannot do better than the identity matrix; i.e., $t = n$. Hence, we are interested in the regime where $w_{\max} > d$.

Given that it is impossible to achieve $t < n$ when $w_{\max} \leq d$, it is natural to ask: what happens when $w_{\max} = d + 1$? The following theorem presents our first (somewhat surprising) result for this case.

Theorem 4: For all integers $d, n \geq 2$ such that $d + 1 \leq \sqrt{n}$ and $\sqrt{n}$ is a prime power, there exists a $t \times n$ matrix that is d-disjunct with constant column weight $w = d + 1$ such that

$$t = (d + 1)\sqrt{n}.$$

On the other hand, for any integer $d, n \geq 2$, a $t \times n$ matrix that is d-disjunct with maximum column weight $w_{\max} \leq d+1$ must satisfy

$$t \geq \min\left\{\sqrt{nd(d+1)}, n\right\}.$$

The proof of the above theorem can be found in Section IV-A.

A few comments are in order. First, Theorem 4 shows that by increasing $w_{\max}$ from d to $d + 1$, we suddenly get $t = \Theta(d\sqrt{n})$ instead of $t = n$. Second, the achievability result in Theorem 4 is obtained by changing the field size from $q = O(d \log_d n)$ to $q = \sqrt{n}$ in Kautz and Singleton's construction. Kautz and Singleton's construction is strictly suboptimal in the non-sparse setting with $d = O(\log n)$. It is interesting that a simple modification to this well known construction makes it optimal (even up to an almost matching constant) in the sparse setting.

We next investigate the more general case where the codeword weights are bounded by $w_{\max} = ld+1$, for some integer $l > 1$. The achievability result in the next theorem is again obtained by tuning the field size in the Kautz and Singleton's construction. In this case we can show that this construction is nearly optimal.

Theorem 5: For every integer $d \geq 2$ and $n > \binom{d+2}{2}$, there exists a $t \times n$ matrix that is d-disjunct with constant column weights $w = ld + 1$ such that

$$t = (ld + 1)\ \sqrt[l+1]{n},$$

where l is any integer satisfying $ld + 1 \leq \sqrt[l+1]{n}$ and $\sqrt[l+1]{n}$ is a prime power.

On the other hand, for any integer $d \geq 2$ and $n > \binom{d+2}{2}$, a $t \times n$ matrix that is d-disjunct with maximum column weight $w_{\max} \leq ld+1$, where $l > 1$ is any integer such that $ld + 1 \leq \sqrt[l+1]{n}$, must satisfy

$$t \geq \left(\frac{(l-1)^{l+1}(d-1)^{l+1}}{2e^l(l-1)(d-1)^{l-1} + 1}\right)^{1/l+1}\ \sqrt[l+1]{n}.$$

The proof of the above theorem can be found in Section IV-B.

Note that as we increase the weights as a multiple of d (i.e., $w_{\max} = ld + 1$), the minimum number of required tests decreases exponentially in l. As we see from Theorem 5, for a fixed l the lower bound we get is $\Theta(d^{\frac{2}{l+1}}\ \sqrt[l+1]{n})$ whereas the upper bound is $\Theta(d\ \sqrt[l+1]{n})$. While we have a matching lower and upper bounds in terms of the scaling with respect to n, there is an increasing gap of $d^{\frac{l-1}{l+1}}$ between them, which approaches d for large l.

IV. Proofs for Sparse Codewords

A. Proof of Theorem 4

We begin with the lower bound. We consider M to be a $t \times n$ d-disjunct matrix. We will separate the columns of this matrix into 2 groups $\mathcal{N}_1, \mathcal{N}_2 \subseteq [n]$ such that $\mathcal{N}_1 \cup \mathcal{N}_2 = [n]$ and $\mathcal{N}_1 \cap \mathcal{N}_2 = \emptyset$. We define a row $i \in [t]$ to be private for a column j, if j is the only column in the matrix having a one on row i. If a column M_j has weight at most d, then it must have at least one private row, otherwise we can find at most d columns such that their union will span M_j which contradicts with d-disjunctiveness. Now consider all the columns that have weights equal to $d + 1$. It may well be possible that some of them also have private rows. Hence, we construct the first set $\mathcal{N}_1$ such that it includes the columns whose weights are less than or equal to d and the ones that have weights equal to $d+1$ such that they have at least one private row. The second set $\mathcal{N}_2$ consists of the rest of the columns; i.e., the ones that

have weights equal to $d+1$ and do not have any private row. Defining w_j to be weight of the column j for $1 \leq j \leq n$, more formally we write

$$\mathcal{N}_1 = \{j \in [n] \mid w_j \leq d \text{ or } w_j = d+1 \text{ and } M_j \text{ has at least one private row}\},$$
$$\mathcal{N}_2 = \{j \in [n] \mid w_j = d+1 \text{ and } M_j \text{ has no private row}\}.$$

Note that by construction, $\mathcal{N}_1 \cup \mathcal{N}_2 = [n]$ and $\mathcal{N}_1 \cap \mathcal{N}_2 = \emptyset$, hence $n = |\mathcal{N}_1| + |\mathcal{N}_2|$. In the following, we will bound the size of both sets $\mathcal{N}_1$ and $\mathcal{N}_2$.

We note that each column in the set $\mathcal{N}_1$ has at least one private row and by definition of the private row it cannot be shared by two distinct columns. Since there could be at most t private rows, we have $|\mathcal{N}_1| \leq t$.

We now consider the set $\mathcal{N}_2$. We generalize the definition of the private row to the private set as follows. A private set for a column is defined as a set of position of ones such that no other column can cover these positions by itself; i.e., no other column has all ones in these positions. We claim that all size-2 sets of positions of ones of a column in set $\mathcal{N}_2$ must be private; i.e., all pairs of positions of ones must be private for a column in set $\mathcal{N}_2$. We prove this by contradiction. Assume there exists a column in the set $\mathcal{N}_2$ such that it has at least one pair of positions of ones which is not private. This means that there exists another column which can cover these positions. Note that any column in the set $\mathcal{N}_2$ has weight $d+1$ and has no private row, therefore, there are $d-1$ positions of ones except this pair and we can find at most $d-1$ columns that can cover union of these positions. This yields that we can find at most d columns that can cover all $d+1$ positions of ones of this column which contradicts with the d-disjunctiveness. Note that there are $\binom{d+1}{2}$ number of pairs of position of ones and by definition of a private set it cannot be shared by two distinct columns. We further note that each column in the set $\mathcal{N}_1$ will have a private row and it must be the case that the columns in the set $\mathcal{N}_2$ must have a zero in these rows, therefore, there could be at most $\binom{t-|\mathcal{N}_1|}{2}$ number of private pairs. Hence, we have

$$|\mathcal{N}_2|\binom{d+1}{2} \leq \binom{t-|\mathcal{N}_1|}{2}.$$

Therefore,

$$|\mathcal{N}_2|d(d+1) \leq (t-|\mathcal{N}_1|)(t-|\mathcal{N}_1|-1) \leq (t-|\mathcal{N}_1|)^2.$$

Defining $n_1 \triangleq |\mathcal{N}_1|$, this gives us

$$t \geq n_1 + \sqrt{(n-n_1)d(d+1)} \tag{1}$$

Note that $0 \leq n_1 \leq t \leq n$. One can take the second derivative of the right-hand side of (1) and observe that it is negative for $0 \leq n_1 \leq t \leq n$ which means it is a concave function of n_1 and it will be minimum at either $n_1 = 0$ or $n_1 = t$. Therefore,

$$t \geq \min\left\{\sqrt{nd(d+1)}, t + \sqrt{(n-t)d(d+1)}\right\}.$$

Noting that $t \geq t + \sqrt{(n-t)d(d+1)}$ only when $t = n$, one can observe that

$$t \geq \min\left\{\sqrt{nd(d+1)}, n\right\}.$$

For the achievability, we use the Kautz and Singleton construction. Let us now consider the case where all weights are equal to $d+1$ by choosing $w = t_q = d+1$ and $k_q = 2$. With this choice, since $n = q^{k_q}$, we obtain $q = \sqrt{n}$ and therefore $t = (d+1)\sqrt{n}$. Note that in order to satisfy the requirement $q \geq t_q$ where q is any prime power, we must ensure that $d+1 \leq \sqrt{n}$ and $q = \sqrt{n}$ is any prime power. This completes the proof for the achievability.

B. Proof of Theorem 5

We begin with the lower bound. We consider M to be a $t \times n$ d-disjunct matrix. The idea will be similar to $l = 1$ case. We will separate the columns into $l+1$ groups and construct the sets $\mathcal{N}_i$ for $i = 1, \ldots, l+1$ such that $\mathcal{N}_1 \cup \ldots \cup \mathcal{N}_{l+1} = [n]$ and $\mathcal{N}_i \cap \mathcal{N}_j = \emptyset$ for any $i, j \in [l+1]$ such that $i \neq j$. We will construct the sets $\mathcal{N}_1, \ldots, \mathcal{N}_{l+1}$ as follows. We will keep the first set $\mathcal{N}_1$ as the columns whose weights are less than or equal to d and the ones that have weights equal to $d+1$ such that they have at least one private row. For $i = 2, \ldots, l$, the set $\mathcal{N}_i$ consists of the columns that satisfies one of the following two conditions: Either its weight is between $(i-2)d+2$ and $(i-1)d+1$ and it has no private set of size $i-1$ or between $(i-1)d+2$ and $id+1$ and it has at least one private set of size i. Finally, the last set $\mathcal{N}_{l+1}$ consists of the columns whose weights are between $(l-1)d+2$ and $ld+1$ and they have no private set of size l. More formally,

$$\mathcal{N}_1 = \{j \in [n] \mid w_j \leq d \text{ or } w_j = d+1 \text{ and } M_j \text{ has at least one private row}\},$$
$$\mathcal{N}_i = \{j \in [n] \mid (i-2)d+2 \leq w_j \leq (i-1)d+1 \text{ and } M_j \text{ has no private set of size } i-1 \text{ or } (i-1)d+2 \leq w_j \leq id+1 \text{ and } M_j \text{ has at least one private set of size } i\}, \quad \text{for } i = 2, \ldots, l,$$
$$\mathcal{N}_{l+1} = \{j \in [n] \mid (l-1)d+2 \leq w_j \leq ld+1 \text{ and } M_j \text{ has no private set of size } l\}.$$

Note that by construction, $\mathcal{N}_1 \cup \ldots \cup \mathcal{N}_{l+1} = [n]$ and $\mathcal{N}_i \cap \mathcal{N}_j = \emptyset$ for any $i, j \in [l+1]$ such that $i \neq j$, hence $n = |\mathcal{N}_1| + \ldots + |\mathcal{N}_{l+1}|$. In the following, we will bound the size of these sets.

Recalling the discussion in the previous section, we have $|\mathcal{N}_1| \leq t$. Consider the sets $\mathcal{N}_i$ for $i = 2, \ldots, l$. For any column $j \in \mathcal{N}_i$, if we have $(i-1)d+2 \leq w_j \leq id+1$, then by construction M_j has at least one private set of size i. For the case $(i-2)d+2 \leq w_j \leq (i-1)d+1$, we claim that all the sets of positions of ones of size i must be private for the column M_j. We can similarly show this by contradiction. Assume there exists a set of positions of ones of size i such that it is not private for the column M_j. Then we can find a column that can cover these positions. Since by construction of set N_i, the column M_j has no private set of size $i-1$, one

can find at most $((i-1)d+1-i)/(i-1) = d-1$ columns that will cover the rest of the positions of ones. Hence we have at most d columns covering the column M_j which contradicts the d-disjunctiveness. Therefore, we obtain that all the columns in the set N_i must have at least one private set of size i. Since the private sets cannot be shared among columns and we have at most $\binom{t}{i}$ private sets of size i, it yields $|\mathcal{N}_i| \leq \binom{t}{i}$. For the last set $\mathcal{N}_{l+1}$, similar arguments apply and for each column, it should be the case that all the set of positions of ones of size $l+1$ must be private. Since $w_j \geq (l-1)d+2$ for $j \in \mathcal{N}_{l+1}$, we have $|\mathcal{N}_{l+1}|\binom{(l-1)d+2}{l+1} \leq \binom{t}{l+1}$. Therefore,

$$\begin{aligned}
n &= |\mathcal{N}_1| + \ldots + |\mathcal{N}_{l+1}| \\
&\leq \sum_{i=1}^{l} \binom{t}{i} + \frac{\binom{t}{l+1}}{\binom{(l-1)d+2}{l+1}} \\
&\overset{(i)}{\leq} \left(\frac{et}{l}\right)^l + \frac{t \ldots (t-l)}{((l-1)d+2)\ldots((l-1)(d-1)+1)} \\
&\overset{(ii)}{\leq} \frac{e^l t^l}{l^l} + \frac{t^{l+1}}{((l-1)(d-1))^{l+1}} \\
&\overset{(iii)}{\leq} \frac{e^l t^l}{(l-1)^l} \frac{t}{(d-1)^2/2} + \frac{t^{l+1}}{((l-1)(d-1))^{l+1}} \\
&= t^{l+1}\left(\frac{2e^l}{(l-1)^l(d-1)^2} + \frac{1}{(l-1)^{l+1}(d-1)^{l+1}}\right) \\
&= t^{l+1}\left(\frac{2e^l(l-1)(d-1)^{l-1}+1}{(l-1)^{l+1}(d-1)^{l+1}}\right)
\end{aligned}$$

where (i) is due to the inequality $\sum_{i=0}^{l}\binom{t}{i} \leq \left(\frac{et}{l}\right)^l$ for $t \geq l \geq 1$, (ii) is bounding all the terms in the numerator by t and denominator by $(l-1)(d-1)$ and in (iii) we use Proposition 1 and $\binom{d+2}{2} \geq \frac{(d-1)^2}{2}$. This completes the lower bound.

For the achievability, we can use the Kautz and Singleton construction. Consider the case where all weights are equal to $ld+1$ by choosing $w = t_q = ld+1$ and $k_q = l+1$. With this choice, since $n = q^{k_q}$, we obtain $q = \sqrt[l+1]{n}$ and therefore $t = (ld+1)\sqrt[l+1]{n}$. Note that in order to satisfy the requirement $q \geq t_q$ where q is any prime power, we must ensure that $ld+1 \leq \sqrt[l+1]{n}$ and $q = \sqrt[l+1]{n}$ is any prime power. This completes the proof for the upper bound.

V. Related Work

Group testing algorithms can be broadly partitioned into adaptive or non-adaptive [9], [19]. In adaptive group testing, the tests are designed sequentially, meaning that the j^{th} test is a function of the outcomes of the $j-1$ previous tests. In non-adaptive group testing, the tests are designed and fixed a priori. Even though adaptive group testing offers more freedom in design, it is known that it does not improve upon non-adaptive group testing by more than a factor d in the number of required tests [9], [19]. In addition to adaptivity, group testing algorithms can be partitioned into combinatorial [9], [19] or probabilistic [20]–[23]. Combinatorial group testing approaches recover the set of defective items with probability one. In contrast, the probabilistic approach allows for a small probability of making a mistake that asymptotically (in the number of items) goes to zero. Relative to combinatorial group testing, the probabilistic approach requires a factor of d less tests when an ε probability of error is tolerable [20].

A large body of work in group testing focuses on designing d-disjunct matrices with minimal $t(d,n)$. When $d \geq 2$ and $\binom{d+2}{2} < n$, it is known that

$$t(d,n) \geq \frac{(d+1)^2}{12\log d}\log n = \Omega\left(\frac{d^2}{\log d}\log n\right).$$

One of the earliest constructions of explicit disjunct matrices due to Kautz and Singleton achieves $t = O(d^2\log_d^2 n)$ [3]. Their construction uses a Reed-Solomon code concatenated with a nonlinear identity code. A more recent construction by Porat and Rothscheld achieves $t = O(d^2\log n)$ [4]. These are the best lower and upper bounds known on $t(d,n)$, with a $\log d$ gap to optimality.

Much of the recent research efforts have focused on designing testing strategies that can be decoded in poly(t)-time, while preserving the order of t as much as possible. [24] gives a result with an efficient decoding time, however, using $O(d^4\log n)$ tests. The first result that achieves an efficient decoding time while matching the $O(d^2\log n)$ bound on the number of tests was recently presented in [14].

To the best of our knowledge, our problem formulation is novel and has not been widely explored in the group testing literature. The only exception is a recent paper by Gandikota et al. [10]. However, Gandikota et al. focus (for the most part) on statistical approaches that provide lower and upper bounds on the number of tests such that the testing achieves an ϵ-error on decoding the defective set while our approach is purely combinatorial. They use information theoretic techniques based on Fano's inequality to prove lower bounds for the ϵ-error case and then specialize their bound for the 0-error case. For the construction of d-disjunct matrices with constraints on the column and row weights, they use randomized designs while we provide explicit constructions. For explicit constructions, they refer to [25], however, the construction in [25] is highly suboptimal as we discuss next.

Table I compares the results provided in this paper with the ones presented in [10]. When the weights of the columns are constrained with $w = ld+1$, the term that depends on n in our lower bound is $\sqrt[l+1]{n}$ whereas [10] provides $\sqrt[ld+1]{n}$ which is significantly weaker. In terms of upper bound, we provide an explicit construction achieving $O\left(d\sqrt[l+1]{n}\right)$ which is better than the randomized construction in [10]. Note for example that when $l = 1$ our upper bound gives $d\sqrt{n}$ while their upper bound gives $d^{1+\frac{2}{d+1}n}$. The explicit construction [25] referred in [10] provides an upper bound of $O(n^{\frac{1}{(ld+1)^{1/d}}})$ which is substantially weaker than the results in this paper.

VI. Conclusion & Discussion

In this paper, we departed from the classical combinatorial group testing framework and studied the fundamental trade-off between t and (d,n) in a setting where there is a constraint on the number of tests per item (sparse codewords).

Setting	This paper		Gandikota et al. [10]	
	Lower Bound	Upper Bound	Lower Bound	Upper Bound
Sparse Codewords with $w = ld + 1$	$\Omega\left(d^{\frac{2}{l+1}}\sqrt[l+1]{n}\right)$	$O\left(d\sqrt[l+1]{n}\right)$ (Explicit)	$\Omega\left(d^2(\frac{n}{d})^{\frac{1}{ld+1}}\right)$	$O\left(d^2(n\left(\frac{n}{d}\right)^d)^{\frac{1}{ld+1}}\right)$ (Randomized)

TABLE I
COMPARISON OF NON-ADAPTIVE SPARSE GROUP TESTING RESULTS.

We proved that by allowing the number of tests per item to be $d + 1$ instead of d, one can achieve $t = \Theta(d\sqrt{n})$ (instead of $t = n$), establishing a sharp transition in t. We then demonstrated that Kautz and Singleton's construction, which is known to be strictly sub-optimal in the classical group testing setting, is order optimal in this setting. We also presented lower bounds on the number of tests with nearly matching constructions when the number of tests per item increases linearly with d.

There are a number of nontrivial extensions to our work. Firstly, as the number of tests per item increases linearly with d (i.e., $w = ld+1$), the gap between our lower bounds on t and the nearly matching upper bounds increases as a function of d. It would be interesting if one can come up with sharper lower bounds or improved constructions that could achieve better performance. Secondly, our results have exclusively focused on constructions. We paid no attention to decoding complexity, which (at best) can be on the order of t (or polynomial in t). One interesting venue for future research would be to study what decoding procedures could exhibit such a decoding performance.

REFERENCES

[1] Jun Luo and Dongning Guo, "Neighbor discovery in wireless ad hoc networks based on group testing," in *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*. IEEE, 2008, pp. 791–797.

[2] Robert Dorfman, "The detection of defective members of large populations," *The Annals of Mathematical Statistics*, vol. 14, no. 4, pp. 436–440, 1943.

[3] W. Kautz and R. Singleton, "Nonrandom binary superimposed codes," *IEEE Transactions on Information Theory*, vol. 10, no. 4, pp. 363–377, October 1964.

[4] Ely Porat and Amir Rothschild, "Explicit non-adaptive combinatorial group testing schemes," *Automata, Languages and Programming*, pp. 748–759, 2008.

[5] Noga Alon, Dana Moshkovitz, and Shmuel Safra, "Algorithmic construction of sets for k-restrictions," *ACM Transactions on Algorithms (TALG)*, vol. 2, no. 2, pp. 153–177, 2006.

[6] Y. Polyanskiy, "A perspective on massive random-access," in *2017 IEEE International Symposium on Information Theory (ISIT)*, June 2017, pp. 2523–2527.

[7] T. Berger, N. Mehravari, D. Towsley, and J. Wolf, "Random multiple-access communication and group testing," *Communications, IEEE Transactions on*, vol. 32, no. 7, pp. 769 – 779, jul 1984.

[8] Jack K Wolf, "Born again group testing: Multiaccess communications," *Information Theory, IEEE Transactions on*, vol. 31, no. 2, pp. 185–191, 1985.

[9] Dingzhu Du, Frank K Hwang, and Frank Hwang, *Combinatorial group testing and its applications*, vol. 12, World Scientific, 2000.

[10] V. Gandikota, E. Grigorescu, S. Jaggi, and S. Zhou, "Nearly optimal sparse group testing," in *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Sept 2016, pp. 401–408.

[11] A. G. D'yachkov and V. V. Rykov, "A survey of superimposed code theory," *Problems Control Inform. Theory/Problemy Upravlen. Teor. Inform.*, vol. 12, pp. 229–242, 1983.

[12] Arkadii Georgievich D'yachkov and Vladimir Vasil'evich Rykov, "Bounds on the length of disjunctive codes," *Problemy Peredachi Informatsii*, vol. 18, no. 3, pp. 7–13, 1982.

[13] Zoltán Füredi, "Onr-cover-free families," *Journal of Combinatorial Theory, Series A*, vol. 73, no. 1, pp. 172–173, 1996.

[14] Piotr Indyk, Hung Q Ngo, and Atri Rudra, "Efficiently decodable non-adaptive group testing," in *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 2010, pp. 1126–1142.

[15] W. W. Peterson, *Error-Correcting Codes*, 1961.

[16] R. C. Singleton, "Maximum distance p-nary codes," *IEEE Trans. on Information Theory*, vol. IT-10, pp. 116–118, April 1964.

[17] I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," *Journal of the Society for Industrial and Applied Mathematics*, vol. 8, no. 2, pp. 300–304, 1960.

[18] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth, "On the lambertw function," *Advances in Computational Mathematics*, vol. 5, no. 1, pp. 329–359, Dec 1996.

[19] Hung Q Ngo and Ding-Zhu Du, "A survey on combinatorial group testing algorithms with applications to dna library screening," *Discrete mathematical problems with medical applications*, vol. 55, pp. 171–182, 2000.

[20] G.K. Atia and V. Saligrama, "Boolean compressed sensing and noisy group testing," *Information Theory, IEEE Transactions on*, vol. 58, no. 3, pp. 1880–1901, 2012.

[21] C. L. Chan, P. H. Che, S. Jaggi, and V. Saligrama, "Non-adaptive probabilistic group testing with noisy measurements: Near-optimal bounds with efficient algorithms," in *2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Sept 2011, pp. 1832–1839.

[22] Dino Sejdinovic and Oliver Johnson, "Note on noisy group testing: Asymptotic bounds and belief propagation reconstruction," *CoRR*, vol. abs/1010.2441, 2010.

[23] Kangwook Lee, Ramtin Pedarsani, and Kannan Ramchandran, "Saffron: A fast, efficient, and robust framework for group testing based on sparse-graph codes," in *Information Theory (ISIT), 2016 IEEE International Symposium on*. IEEE, 2016, pp. 2873–2877.

[24] Venkatesan Guruswami and Piotr Indyk, *Linear-Time List Decoding in Error-Free Settings*, pp. 695–707, Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.

[25] Anthony J. Macula, "A simple construction of d-disjunct matrices with certain constant weights," *Discrete Mathematics*, vol. 162, no. 1, pp. 311 – 312, 1996.