Fast Neighbor Discovery in MEMS FSO Networks

Michael Atakora and Harsha Chenji
School of Electrical Engineering and Computer Science, Ohio University, USA
ma136412, chenji@ohio.edu

***Abstract*—We investigate whether it is possible to achieve sub-millisecond latency for the discovery of multiple neighbors in laser-based Free Space Optical (FSO) networks. Given a large programmable array of micromirrors, we propose the use of adaptive boolean combinatorial group testing algorithms that are practical and efficient. The time taken scales as $O(N \log(L/N))$ for N neighbors even if N is unknown, but *no additional computation* (e.g., matrix inversion) is required. Compared to Raster and Lissajous pattern-based scanning, we report 99.92% and 87% reduction in latency, respectively, for an array of 10^6 micromirrors (approximately XGA resolution). We conclude that it is indeed possible to achieve sub-ms latency given realistic network parameters. Our proposed algorithms are evaluated in simulation, and are compared against state of art neighbor discovery schemes.**

I. Introduction

Laser-based Free Space Optical (FSO) networks have the potential to be interference-free while simultaneously providing several Tbps of throughput. This is possible due to highly directional laser beams, whose beamwidth is typically measured in micro or milli-radians, and whose bandwidth spans hundreds of GHz. In FSO-based mobile networks, this advantage becomes problematic due to the need for discovering and tracking these highly directional beams. FSO network hardware must be able to quickly steer beams in $3D$ space, be able to select a subset of multiple incident beams, be programmable enough to accommodate the algorithms required by higher layer protocols, and be able to demodulate/decode the incident power while suffering low loss.

MicroElectroMechanical Systems (MEMS) have found popularity in both radio frequency and optical wireless and mobile network [1]. A particular type of MEMS device, Digital Micromirror Device (DMD), consists of an array of micromirrors, each on the order of $10\mu m$ in width with an individually controllable tilt degree of freedom. These DMDs have been used in FSO networks to perform beam alignment and scanning. Each element in the array can be controlled individually, by sending a binary vector representing the per-mirror on/off state to a DMD controller. Through repeatedly sending different patterns and synchronizing with a R/G/B light source, a digital video projector can be built.

These DMDs can be to solve the neighbor discovery problem, where the locations and orientations of an unknown number of neighbors N is to be calculated by a receiver. The laser beams transmitted by neighbors will be incident on a subset of size $k < L$ mirrors of the L-sized DMD array at any instant. Our task is to find an efficient algorithm that finds this subset, and to investigate its scalability as $L \to 10^8$. It is not impossible to imagine an omnidirectional fly-eye receiver in the shape of a sphere that is tesselated with micromirrors; the key question is how neighbor discovery algorithms will perform at such scale, and if sub-millisecond discovery times are possible.

Since it is unlikely that the entire array is "lit" by the neighbors due to the very low beamwidths, $k \ll L$; in our simulations, $k \cong 400$ for $N = 5$ and $L = 10^6$. So, compressive sensing or other approaches that take advantage of sparsity, such as matrix completion, can be used. But in reality, k is typically unknown and may not even be sparse in certain extremes. Therefore, and for reasons explained in the next section, we investigate combinatorial group testing as a solution. We are interested in not just the theoretical asymptotic complexity, but also the performance of a practical implementation for a given k & L. Even if two algorithms have the same asymptotic performance, their time complexity could vary in the constant term - which could make a massive difference in the implementation. Additionally, k could vary non-linearly with network parameters of interest such as number of neighbors N, transmit power etc. The ultimate aim is to understand if sub-ms discovery time is possible, and if it is, then to understand the tradeoffs that need to be made.

We characterize the above problems and propose solutions. Our contributions are as follows: 1) casting neighbor discovery with unknown N as a boolean combinatorial group testing problem and proposing solutions that can handle unknown k and also non-sparsity; 2) improving implementation performance by finding different solutions that have the same asymptotic complexity but differ in the constant term; and 3) characterizing the tradeoffs needed for achieving sub-ms neighbor discovery latency in MEMS-based FSO networks.

The rest of this paper is laid out as follows: we motivate the need for our research and review related work in Section II. In Section III, we present a framework for neighbor discovery in FSO networks. Possible application scenarios and state of art techniques are reviewed. This is followed by a discussion of the proposed algorithms, based on boolean combinatorial group testing algorithms. A performance evaluation of these schemes is presented in Section IV (including discussion of sub-ms discovery), after which we discuss future work.

II. Background & Related Work

Let X (with $L = |X|$) be a set of items with a subset Y that are defective. In group testing, we perform a "test" by selecting a subset of X, and this query returns 1 if it contains at least one defective item. The aim is to reliably identify all the defective items while minimizing the number of queries or measurements M. Equivalently, $\mathbf{x}$ is a binary vector with

$k = |Y|$ ones. A binary query vector $\mathbf{\Phi}$ chooses positions in $\mathbf{x}$; the result of the query is $\bigvee_{i=0}^{L} x_i \phi_i$. Typically, L is very large while k is a few orders of magnitude smaller. Clearly, one could test each element of $\mathbf{x}$ individually, resulting in a time complexity of $O(L)$ queries, but research has shown that this can be improved. Solutions to the group testing problem can be classified into adaptive, where a query can depend on the result of previous queries, & non-adaptive, where the queries can be parallelized; into noisy, where the result of the query is affected by noise resulting in bit flips, & non-noisy; into probabilistic, where the items are identified correctly with a probability < 1 even if the queries are non-noisy, & deterministic.

If the k ones are equally likely to occur anywhere in the vector $\mathbf{x}$, then the number of possible outcomes are $\binom{L}{k}$, and the entropy is $\lceil \log_2 \binom{L}{k} \rceil \cong k \log_2(L/k)$ which establishes a lower bound on M. In the noiseless case, when k scales as $O(L^{1/3})$, M scales as $(k \log_2(L/k))(1 + o(1))$ in *both* the adaptive and non-adaptive settings [2]. This remains unchanged in the adaptive setting when k scales as $\Theta(L^{\theta})$, $\theta \in [1/3, 1]$. Adaptive algorithms that perform close to the lower bound are known to be practically usable and efficient in terms of implementation [2]. The best known adaptive algorithm has a complexity of $O(k \log(L/k)) + O(k)$ [3]. Competitive group testing [4], [5] deals with unknown k. With a c-competitive algorithm, M is within a constant factor $c \geq 1$ of the corresponding M when k is known. Du and Park [5] proposed the notion of strong competitiveness, and showed that $c \to 1$ as $L \to \inf$ for $k \geq 1$. Their algorithm achieves $M \leq k \log_2(L/k) + 4k$.

In compressed sensing (CS), a k-sparse signal $\mathbf{x}$ of dimension L can be recovered by constructing a measurement matrix $\mathbf{\Phi}$ such that the output $\mathbf{y} = \mathbf{\Phi x}$ can be decoded to recover $\mathbf{x}$. $\mathbf{y}$ is of size $M \times 1$ with $k \ll M \ll L$; thus $M = O(k \log(L/k))$ measurements are needed in the non-adaptive and noiseless case. In a sense, CS is a continuous version of group testing which is combinatorial. Practical applications of compressed sensing such as the single pixel camera [6] either require some knowledge of k for the minimization of M, or overestimate M for reliable signal recovery.

Decoding algorithms for $\mathbf{y}$ include l_0 pursuit which is optimal but NP-hard, l_1 pursuit which is a widely used heuristic, and techniques such as LASSO. Practically, such algorithms contribute additional latency to the process. For example, orthogonal matching pursuit and basis pursuit run in $O(kML)$ [7] and $O(L^3)$ [8] time respectively. Also, k is generally not known in advance (as is the case in this paper). Because of these two reasons, in this paper we only consider approaches which do not rely on prior knowledge of k. Even if an upper bound k on the number of incident mirrors is known in advance, the cost of decoding and recovery is prohibitive due to large $L \cong 10^6$.

Over the past few years there have been some efforts in the area of MEMS-based optical networking. In [9], [10], a demonstration of a 10 Gbps gimbal-less MEMS-based reconfigurable FSO link for data center applications is presented. These efforts utilize a secondary channel (visible red laser beam) for target acquisition. In our approach however, we do not use a secondary channel for acquisition. FSONet [11] uses a combination of galvo mirrors and a motorized rotation stage for beam steering. In addition, the acquisition scheme in [9], [10] requires some knowledge of initial node location. In contrast, the acquisition scheme presented in this work is totally oblivious to initial target location.

In MEMSEye [12], [13], a spiral scan pattern is employed for target acquisition. Spiral scan is an exhaustive and linear search approach and is therefore not suitable for extremely low latency applications such as neighbor discovery. An adaptive field of view MEMS-based wireless optical transmitter achieved by varying the lens-laser distance, and capable of delivering a throughput of 3 Gbps over 7 m is presented in [14]. The device built in [14] uses a feedback channel from the receiver to the transmitter for beam alignment and does not address the multi-node neighbor discovery problem.

A link establishment protocol in LED-based networks for a pair of nodes is presented in [15], [16]. This is achieved in-band via the exchange of frames of short length on steerable transceivers rotating at different speeds. In [16], the assumption is made that an omnidirectional RF link is available for the dissemination of transceiver orientation information. The three-way handshake discovery mechanism is extended to a 3D space [17]. These neighbor discovery efforts are limited to just a pair of nodes. In this paper however, the schemes presented are oblivious to the number of discoverable nodes.

III. Group Testing Framework for FSO Networks

In this section, we design a neighbor discovery, and clustering framework using group testing (GT) algorithms applied to array-oriented MEMS hardware such as DMDs. As mentioned previously, the decoding/recovery complexity of compressive sensing (CS) algorithms is prohibitive, and eliding it is the key to reducing latency. This is possible with GT algorithms which do not have any recovery complexity due to their boolean nature. Additionally, an upper bound on k for CS algorithms would mostly be inaccurate if the network topology is unknown or *ad hoc*.

We first present the problem formulation, followed by the chosen adaptive GT algorithm. As mentioned previously, this algorithm can be implemented in multiple ways, leading to schemes that have the same asymptotic complexity but different constants; we present one and two stage implementations of state of the art GT algorithms that can be applied to neighbor discovery.

A. Problem Formulation and System Model

Consider a DMD array with L elements organized as a two dimensional array of R rows and C columns (Figure 1b). This array is present at the receiver of each node in a FSO network. A given node has N neighbors whose locations and orientations are unknown, in addition to N being unknown. The aim of the neighbor discovery process is to reliably find N and the indices of the mirrors corresponding to each N.

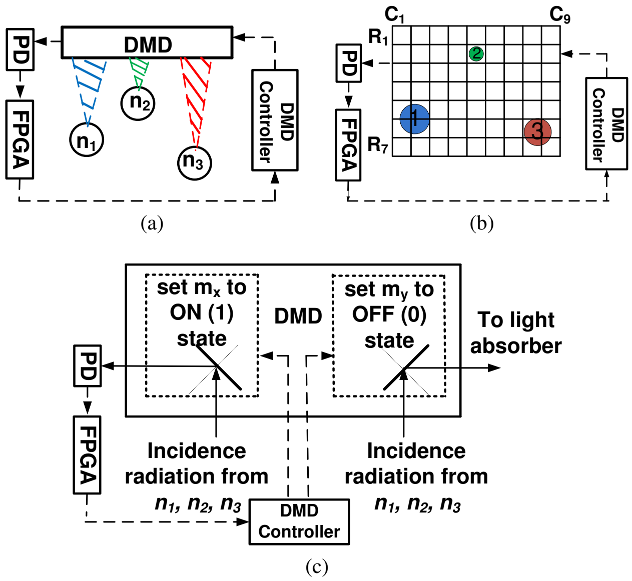

Figure 1: Architecture of the MEMS-based optical neighbor discovery system (a) Nodes n_1, n_2, and n_3 are within the field of view of the discoverer (b) Incident spots on discoverer's DMD array ($k = 9$, $L = 63$) (c) In the process of discovering n_1, n_2, and n_3, mirrors m_x where $x = 14, 37, 38, 46, 47, 53, 54, 62, 63$ have to be in the ON state at certain instances.

The transmit beams emanating from neighbors are incident on the DMD array. The diameter of the beam of wavelength λ emanating from a neighbor at a distance D is $2\omega = 4\lambda f_l/\pi d_i$ where f_l is the focal length of the lens having a diameter d_r, and $d_i = \min(d_r, 2D\tan\frac{\theta_v}{2})$ where θ_v is the vertical beamwidth. Note that it is assumed that for both nodes, the DMD is located at a distance f_l from it's lens. Thus, there are N such spots incident on the DMD array in question. In particular, $k \propto N$ elements of the array (of size L) are "lit" by the neighbors. Note that N and hence k is unknown.

The neighbor discovery problem is now to find the k array elements that are lit (Figure 1). In Figure 1a, the node performing the discovery has three neighbors (n_1, n_2 and n_3) within the field of view of its DMD array. The incident beams appear as Gaussian spots on the DMD array (Figure 1b). To accurately determine the collection of k incident mirrors, any solution has to set the state of individual mirrors using adaptive/non-adaptive patterns. This is achieved via the DMD Controller. Mirrors set to the `1` state reflect incident signals unto the photodetector (PD). In the `0` state, they reflect signals unto a light absorber as is highlighted in Figure 1c. In the adaptive case, the DMD controller sets mirror patterns based on PD current measurements from a Field Programmable Gate Array (FPGA) device. Neighbor discovery latency d_l is the duration it takes to identify all k incident mirrors; $d_l = \frac{M}{f_p}$ where f_p is the maximum bit pattern rate f_p of the DMD.

B. Application Scenarios and State of Art

We now present a scenario in which MEMS based optical neighbor discovery can be applied to, after which we discuss the state of art in discovery approaches.

Application Scenarios:- Before two nodes can communicate with each other, they initially need to identify the spatial directions via which both can be reached. This is the neighbor discovery stage and it involves the identification of all nodes within one hop of a reference node. Extremely low latency multi node neighbor discovery in FSO networks is a prerequisite for high throughput multicast in 5G and other emerging wireless technologies which support point to multipoint communications. A potential application of MEMS based neighbor discovery is the provision of an inter-rack networking solution in data center networks [10]. The current state of art employs scan patterns for neighboring Top-of-Rack switch discovery which either do not guarantee the detection of all k incident mirrors, or are exhaustive but with linear search patterns.

To clarify, "exhaustive" means that all L mirrors are tested, and "linear" means that each mirror is tested one-by-one (i.e., $M = L$ measurements or tests are performed). In contrast, in this paper, we propose an exhaustive approach which requires $M \ll L$ measurements. We assume that all nodes are static with transceivers whose transmit/receive pointing angles are unknown, and we do not address mobility. We now discuss state of art scanning techniques, and the M incurred by them.

Raster and Lissajous-pattern Scans:- The target acquisition problem in beam steerable MEMS-based FSO networks has been explored in [10]. In the considered scenarios, at most one target's exact location is to be determined subject to some knowledge of a constrained angular range in which the target can be precisely found. Directly applying the schemes in [10] to oblivious multiple node neighbor discovery is a non trivial problem, as we shall show. The scan patterns used in [10] are the raster, spiral and Lissajous patterns. The raster and spiral scans are exhaustive patterns and require $M = L$ measurements. But, armed with knowledge of the potential node's transceiver orientation, the search space can be significantly reduced leading to lower latency, using Lissajous patterns. They are defined by parametric equations of the form

$$x[i] = \frac{\sqrt{A}}{2}\sin\left(\frac{2\alpha\pi}{\sqrt{A}}i + \delta\right), \forall i \text{ s.t. } 1 \le i \le \lceil\sqrt{A}\rceil$$
$$y[j] = \frac{\sqrt{A}}{2}\sin\left(\frac{2\beta\pi}{\sqrt{A}}j\right), \forall j \text{ s.t. } 1 \le j \le \lceil\sqrt{A}\rceil \quad (1)$$

where A is the number of mirror elements in a square-shaped search space on the MEMS array, i, j are the indexes of the mirrors on the MEMS array.

For example, in the scan of the entire MEMS array presented in Figure 2c, $A = L$, $\delta = \frac{\pi}{2}$, $\alpha = 3$ and $\beta = 2$. Note that for the single node location acquisition problem, even if some knowledge of node position is available, there is no guarantee that all k incident mirrors are identified. This stems from the observation that the Lissajous pattern follows a trajectory which does not sequentially activate all mirrors (one at a time) in its path. To absolutely guarantee that all k incident mirrors are identified, each mirror in the search space has to be switched to the `1` state at least once. As we

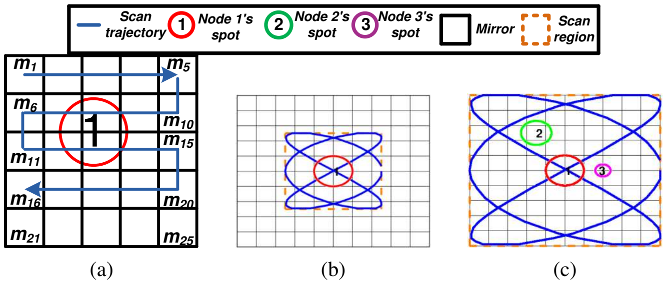

Figure 2: Neighbor discovery via Raster and Lissajous scan patterns (a) A raster scan pattern for the discovery of a single neighbor. Mirrors m_x where $x = 7, 8, 9, 12, 13, 14$ detect the presence of radiation when in the 1 state. (b) Lissajous scan with some knowledge of node's 1 location, a restricted search space is scanned (c) Lissajous scan with no knowledge of any node's location, the entire MEMS array is scanned.

observe in Figure 2c, blindly applying the $\alpha = 3$, $\beta = 2$ Lissajous pattern on the entire MEMS array results nodes 2 and 3 being undiscovered eventhough they are within the discoverer's FOV. Another interesting observation to bear in mind is that due to the sequential nature of the tests associated with Lissajous scan patterns, the number of measurements M does not depend on network parameters such as the number of neighbors N, the aperture focal length f_l and the average distance D from neighbor which determine the size of the Gaussian spot each neighbor generates on the MEMS device (i.e. k does not affect M). We analyze the performance of the $\alpha = 3$, $\beta = 2$ Lissajous pattern (LIS) vis-á-vis our proposed group testing based approaches in Section IV.

When there is no prior neighbor information, a simple raster scan (RAS) can be employed to identify all k incident mirrors. We use Figure 2a to illustrate this approach. The sequential scan starts by placing mirror m_1 into the 1 state, with all other mirrors in the 0 state. The trajectory of the scan if no incident signal is detected is the activation of mirrors m_1 to m_5 to m_6 to m_{10}...in that order, terminating at m_{25}. In the event incident signals are detected in row r_p, and assuming the spot on the MEMS array traverses q other rows, the scan ends at the end of row r_{p+q+1}. In Figure 2a, after all mirrors in the first row are scanned sequentially, no incident beam is detected. When mirrors m_x where $x = 7, 8, 9, 12, 13, 14$ are sequentially activated, each of them reflects incident signals unto the PD. For this example, p=2, q=1 and the scan terminates at the end of row r_4 (mirror m_{21}). For this scheme, M in the best case is $2C$. The best case corresponds to a spot which encompasses a single mirror in the first row. The worst case $M = R \times C$ occurs when (i) no incident spot is detected or (ii) either the last detected spot occurs in the $(R-1)^{th}$ or R^{th} row.

C. Group Testing-based Solution

It is easy to see that the RAS scheme, oblivious of N, requires $M = R \times C$ measurements but yields a 100% accurate N. LIS reduces M by only testing mirrors that are part of the chosen Lissajous pattern, which means means that N may be underestimated. Thus, the state of art approaches are either exhaustive but linear (e.g. raster) or do not guarantee the detection of all k incident mirrors (e.g. Lissajous scan). We now show how we can significantly improve upon the number of measurements M that are required, despite being exhaustive in nature and guaranteeing a 100% accurate N.

The optical neighbor discovery problem can be modeled as a group testing problem [18] (see Section II). The "lit" elements (Section III-A) are analogous to "defective" items, and the elements of the DMD array are the set of L items. A measurement (or a "test" in GT terminology) returns 1 if at least one element in the group is lit. Physically, this is implemented using a PD which measures the total energy incident (reflected) by the mirror elements. If the current generated is more than the noise floor, then it can be deduced that at least one element is lit. Note that the beams are assumed to be Gaussian, meaning that the intensity reduces exponentially as one moves away from the center of the beam, greatly reducing false positives.

GT algorithms which achieve close to the information theoretic lower bound are already well known. We adapt the algorithm provided in [19] (called GT-A in this paper) which requires a total of $M \leq 1.65k(\log_2 \frac{L}{k} + 1.031) + 5$ tests. It (Alg. 1) takes as input a collection D of L mirror elements and produces a subcollection of D' of all k incident elements. In line 1, the state S of all mirrors in D is set to 1. This ensures that at this stage, all mirrors focus light onto the PD. From lines 2 to 27, the algorithm considers situations in which there are at least three undetected incident mirror elements. The algorithm then proceeds to measure the power from activating disjoint collections of mirror elements of size $3, 12, 48, \ldots, 2^h + 2^{h+1}$ (h is the group size parameter) till the presence of optical irradiance is detected by the PD.

The presence or absence of a beam is determined by a simple threshold comparison test. We compare the measured output to the PD's noise floor. If a beam is not detected (line 6), the states of all elements in that collection S_X is updated to 0, and we no longer need to test those elements. In the event that $h = 10$, we test for the presence of a beam with the input set being all mirrors with a state of 1 (lines 7-9).

In the case that the presence of irradiance is detected in the previous steps, D' is updated with the collection of incident elements out of a maximum of 3 elements in a 1 state that produces this signal detection via lines 12 to 20. When the input set consists of more than 3 mirror elements with a state of 1 (lines 22 to 27), binary search is used to extract a single incident element. In lines 28 to 30, if there are at most 2 elements with a state of 1, then, each of them is checked independently for the presence of irradiation.

In addition to GT-A, we implement another GT algorithm, [20] called GT-B in this paper, which is conceptually similar to [19] but requires $M \leq (1.5+\epsilon)k(\log_2 \frac{L}{k} + 1.09) + \frac{0.4}{\epsilon}$ tests with $0 < \epsilon < 0.01$. Unlike in [19], whereby at least one incident mirror element out of a possible three is identified (lines 12-20 of Alg. 1), in [20], three out of four mirror elements are individually tested to determine the presence or otherwise of radiation. In addition, lines 22-25 are modified in [20] to identify a set of mirror elements of cardinality $4^{(h-1)}$ which contains at least one single incident element.

Algorithm 1: A GT based discovery algorithm

Input: A collection D of L mirror elements
Output: A subcollection D' of all k incident mirrors elements
1 $D' \leftarrow \emptyset$
  Set state of all L elements to `1` i.e. $S_{1,\ldots,L} \leftarrow 1$
2 **while** $L \geq 3$ **do**
3   $h \leftarrow 0$
4   **do**
5     $X \leftarrow$ pick from $D \min(2^h + 2^{h+1}, L)$ elements with `1` state
      if $(y_X \leq y_t)$ **then**
6       $S_X \leftarrow 0$; $h \leftarrow h + 2$; $L \leftarrow L - 1$;
7   **if** $(h = 10)$ **then**
8       $X \leftarrow$ pick from D all elements with `1` state
        if $(y_X \leq y_t)$ **then**
9         $S_X \leftarrow 0$; $L \leftarrow L - |X|$;
10 **while** $((L \neq 0) \vee (y_X \leq y_t))$;
11 **if** $(y_X > y_t)$ **then**
12   **if** $(h = 0)$ **then**
13      $u, v, w \leftarrow 3$ elements from X with `1` state
        if $(y_u \geq y_t)$ **then**
14        $D' \leftarrow D' \cup u$;
15   **if** $(y_v \geq y_t)$ **then**
16        $D' \leftarrow D' \cup v$;
17   **if** $(y_u \leq y_t) \wedge (y_v \leq y_t)$ **then**
18        $D' = D' \cup w$; $S_{u,v,w} \leftarrow 0$; $L \leftarrow L - 3$
19   **else**
20        $S_{u,v} \leftarrow 0$; $L \leftarrow L - 2$
21   **if** $(h > 0)$ **then**
22      $X' \leftarrow$ pick from $D \min(2^h, |X|)$ elements with `1` state
        if $(y_{X'} > y_t)$ **then**
23        $X \leftarrow X'$
24   **else**
25        $X \leftarrow X - X'$; $S_{X'} \leftarrow 0$; $L \leftarrow L - |X'|$; Element $d \leftarrow$ Apply binary search to X
          $D' \leftarrow D' \cup d$; $S_d \leftarrow 0$; $L \leftarrow L - 1$;
26 **while** $L \neq 0$ **do**
27   $x \leftarrow$ an element of D with a state of `1`; $L \leftarrow L - 1$;
     if $(y_x \geq y_t)$ **then**
28     $S_x \leftarrow 0$; $D' \leftarrow D' \cup x$;

D. Hierarchical Implementation and Clustering

The above algorithm can be implemented in various stages, leveraging the two dimensional layout (R rows, C columns) of the DMD array. We present multiple ways to implement GT. These schemes are: One-Stage Group Testing (OSGT-A, OSGT-B), and Two-Stage Group Testing (TSGT-A, TSGT-B). Note that the schemes presented in this paper are fully adaptive i.e in addition to subsequent stages using the outcomes of previous stages to decide which mirrors should be tested, all tests within a stage are adaptive. In all these schemes, a scan pattern is generated by setting individual mirror states to either

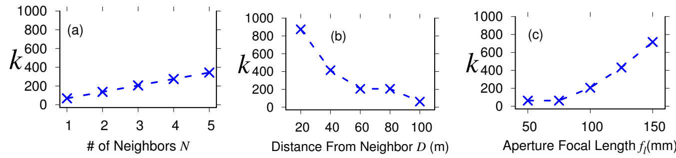

Figure 3: Avg. num. of incident mirrors k vs. network parameters. (a) Effect of N on k for N=1-5, (b) Effect of D on k for D=20-100, (c) Effect of f_l on k for f_l=50-150 mm.

a `1` or `0`. In transitioning from one state to the other, it takes a period of duration equal to the mirror switching time t_s.

In *One-Stage Group Testing (OSGT-A)*, Algorithm 1 is applied to the entire collection of DMD mirrors $|D| = R \times C$ i.e. $L = R \times C$, which discovers all k incident mirrors in a single step. In *Two-Stage Group Testing (TSGT-A)*, we first discover the "lit" rows by treating all the elements of a row as a single element & applying GT-A, followed by the "lit" columns and then take the intersection. We apply algorithm 1 to all rows to obtain the collection of rows of mirrors D'_r which are incident to a beam. In this step, $|D| = L = R$ and $D' = D'_r$. We then reapply algorithm 1 to all columns to obtain D'_c. In this step, $|D| = L = C$ and $D' = D'_c$. The application of algorithm 1 to rows and columns constitute the first stage. In the second stage, using all mirrors in the intersection of D'_r and D'_c as input, algorithm 1 is subsequently applied to obtain $\mathcal{D}'$. For OSGT-B and TSGT-B, GT-B replaces GT-A.

IV. Performance Evaluation

In this section we analyze and compare simulation results, including achieving sub-ms discovery times. The schemes compared are: One-Stage Group Testing (OSGT-A, OSGT-B with $\epsilon = 0.01$) and Two-Stage Group Testing (TSGT-A, TSGT-B with $\epsilon = 0.01$). We benchmark these algorithms against the $\alpha = 3$, $\beta = 2$ Lissajous (LIS) and raster (RAS) pattern-based acquisition scans of [10]. We measured the performance using number of measurements M, number of incident mirrors k and neighbor discovery time d_l as metrics, with a Java based simulator. Each data point is sampled across 1000 random runs. The parameters we use for the analysis are number of neighbors N, aperture focal length f_l, number of elements L in DMD array and average distance from the neighbor D.

The realistic default values (and ranges) used are: $N = 3$ (1-5), $f_l = 100$ mm (50-150 mm), $L = 10^6$ (10^4-10^8) each with an area of 10^{-10} m^2 and $D = 60$ m (20-100 m). Additionally, we used an optical wavelength $\lambda = 1550$ nm, a lens aperture diameter $d_r = 30$ mm, a photodetector (PD) sensitivity of -30 dBm which determines the threshold for detection, a transmit power $P_t = -2$ dB, and a maximum binary pattern rate f_p of 48K binary patterns/sec. Note that M for LIS and RAS depend only on L and are independent of k.

A. Number of Incident Mirrors

Incident mirrors are those which can each individually generate a PD current greater than the noise threshold, when its incident light is directed to the PD; k is the number of

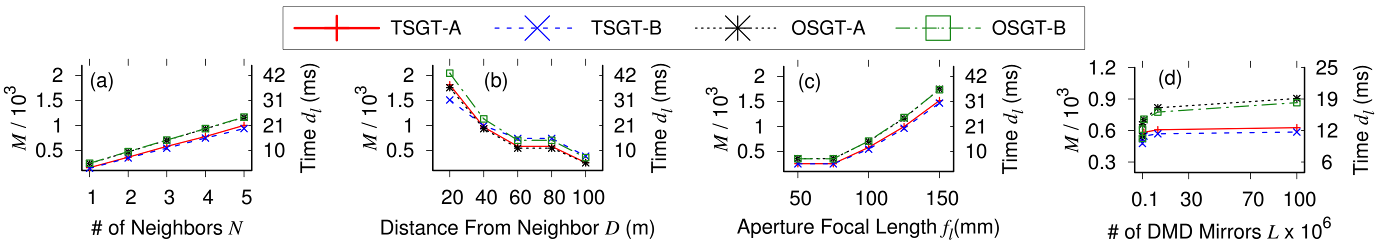

Figure 4: Average number of measurements M and discovery time d_l versus various network parameters. (a) Effect of N for N=1-5, (b) Effect of D for D=20-100, (c) Effect of f_l for f_l=50-150 mm., (d) Effect of L for L = 10^4-10^8.

incident mirrors. In Figure 3a, it is intuitive that as the number of neighbors N within the field of view of the discoverer increases, k increases linearly. This is because a cluster of mirrors is associated with a discovered node. Assuming each neighbor is equidistant from the discoverer, then $k \propto N$. As the distance D of a neighbor increases, k decreases as presented in Figure 3b. This observation is a consequence of the spot size equation $2\omega = \frac{4\lambda f_l}{\pi d_i}$ (see subsection III-A). In addition to these results, we found through simulation that transmit P_t does not significantly affect k at reasonable neighbor distances. This is because the $\frac{1}{e^2}$ spot diameter does not depend on the received optical power but on f_l and d_r.

B. Number of Measurements M and Scalability

Generally as k increase, M increases for each of the evaluated schemes. The complexity of the OSGT schemes is $M = O(k\log(\frac{L}{k}))$. TSGT schemes have a measurement complexity of $M = O(|D'_r|\log(\frac{R}{|D'_r|}) + |D'_c|\log(\frac{C}{|D'_c|}) + k\log(\frac{|D'_r|\times|D'_c|}{k}))$. For any particular GT-based neighbor discovery algorithm, two stage schemes perform better than their corresponding one stage counterparts (e.g. TSGT-A requires a lower M compared to OSGT-A). The reason for this is that, the first stage of two stage schemes reduces the search space for the second stage.

In one stage schemes however, the search space is the entire collection of mirrors in the DMD array. On average OSGT-B requires about 29% more tests compared to TSGT-B (this is similar to results for OSGT-A and TSGT-A) . From the results in Figure 4, two stage schemes have a lower discovery time compared with their one stage counterparts. This is because after the first stage, the search space reduces leading to a smaller discovery time. Discovery latency d_l, plotted on the right side Y-axes of Figure 4 is $\propto M$.

We observe that GT-B outperforms GT-A (e.g. TSGT-B (OSGT-B) performs better than TSGT-A (OSGT-A)). For sparse scenarios ($k \ll L$; $k \approxeq 400$ for $N = 5$ and $L = 10^6$ in our simulations) in which the distribution of the coordinates mirrors which form a cluster are dependent on each other (e.g. follow a Gaussian distribution in this work), many unnecessary tests on sets which do not contain incident mirror elements can be avoided by adaptively testing larger disjoint sets. For GT-B, disjoint sets of size 4, 16, ..., 4^{m-1} are tested till a set of mirror elements which contains at least one incident element incident is discovered in the m^{th} test. For GT-A, if a disjoint set of size $3(4^{n-1})$ is discovered during the n^{th} test, then disjoint sets of size 3, 12, ..., $3(4^{n-2})$ tested to contain no incident element in the preceding $n-1$ tests. Note that since $t \in \mathbb{Z}^+$, $4^t > 3(4^{t-1})$, $m <= n$ and GT-B saves n-m tests over GT-A.

Scheme	$L = 10^4$	$L = 10^5$	$L = 10^6$	$L = 10^7$	$L = 10^8$
LIS	10.6ms	34.7ms	111.9ms	355.1ms	1124.8ms
RAS	208.3ms	2.1s	20.8s	208.3s	2083.3s
OSGT-A	11.4ms	13.6ms	14.6ms	17.1ms	18.8ms

Table I: Average neighbor discovery time d_l for various schemes for L = 10^4-10^8.

Investigating scalability of d_l with L is one of the goals of this paper. In Figure 4d, increasing L from 10^4 to 10^8 results in a 25% increase in M for TSGT-A, a 23% rise in M for TSGT-B, a 66% increase in M for OSGT-A, and a 65% increase in M for OSGT-B. The logarithmic trends in Figures 4d is because $M \propto \log L$. The logarithmic scaling results are particularly important since we envisage the use of omnidirectional fly-eye receivers in the shape of a sphere that is tesselated with tens/hundreds of millions of micromirrors. We note that in all scenarios, with these group testing algorithms, the neighbor discovery time is less than 45ms.

C. Comparison with the State of Art

In Table I, we compare the performance of one of our proposed GT-based schemes (OSGT-A) to the Lissajous pattern (LIS, subsection III-B) and raster scan (RAS, subsection III-B) based discovery methods. We select OSGT-A for comparison because it has a slightly worse scalability performance relative to our other TSGT-A, TSGT-B, OSGT-B algorithms as is shown in Figure 4d. Recall that RAS is an exhaustive-but-linear approach when applied to oblivious neighbor discovery, and the number of measurements $M = L$ since L mirrors have to be individually scanned to guarantee the discovery of all N neighbors. LIS is quasi exhaustive, so as the search space L increases, the number of mirrors in the trajectory of Lissajous pattern path also increases. With LIS, all mirrors in this path are individually scanned for the presence of an incident signal. Note that LIS discovers $\leq N$ neighbors in the oblivious case since some mirrors are never placed in the 1 state (see subsection III-B). The adaptive OSGT-A and by extension all our proposed algorithms significantly outperforms both LIS and RAS as $L \rightarrow 10^8$ w.r.t. neighbor discovery latency while guaranteeing the detection of all N neighbors in the oblivious case despite being an exhaustive scheme.

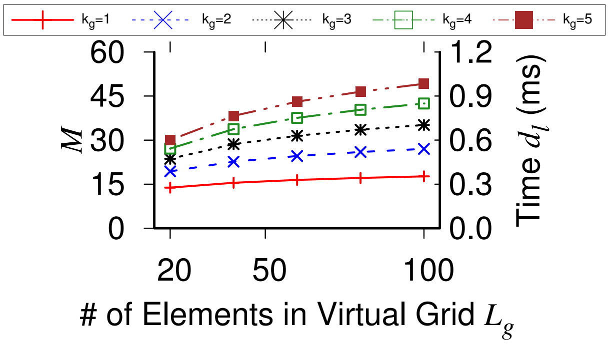

Figure 5: Illustrating sub-ms detection using the virtual grid concept for $k_g = 1-5$ over different L_g.

D. Achieving Sub-ms Discovery Time

Our implemented algorithms yield discovery times in the order of tens of *ms* for GT based neighbor discovery schemes. We believe that this result can be significantly improved to sub-*ms* time using techniques such as 1) reducing L and k by creating virtual grids within the DMD, which we call successive approximation; and 2) techniques such as random hashing. In 2), assuming large non-contiguous arrays of mirrors are selected from L, the probability of detecting groups which are not incident to a beam is high if $k << L$. With 1), virtual non-overlapping grid elements that correspond to a collection of mirrors can be created and tested at a go for incident radiation. Once energy is detected, a refinement step is applied to the virtual grid element to accurately identify incident mirrors. To buttress our position that sub-ms detection latency is indeed possible, we simulated (Figure 5) a scenario whereby mirrors in the DMD array were grouped into L_g elements in a virtual grid, with the area of a single element in the grid ≊ the average 2ω. If there are k_g spots on the DMD array, then the problem translates into finding k_g out of L_g virtual elements that are incident to radiation. Results in Figure 5 which is based on the theoretic M upper bound for OSGT-A indicate that sub-ms is indeed achievable for some k_g and L_g.

In addition to these enhancements, we assumed in this work that the coordinates of the k incident mirrors are independent of each other. However, these k are clustered into P (=N) spots which correspond to the N neighbors. In reality, the coordinates of incident mirrors in each cluster are not independent of each other. Relaxing this assumption would further reduce significantly the number of measurements (and discovery latency), since we no longer need to discover all k mirrors to reconstruct the spatial directions of neighbors.

V. Conclusion and Future Work

For fast neighbor discovery in MEMS-based FSO networks, exhaustive search schemes do not perform efficiently. We formulated the neighbor discovery problem as a boolean adaptive combinatorial group testing problem and provided solutions based on group testing. We showed via simulation that the discovery of neighbors using MEMS-based optical devices can be achieved with a latency of a few ms while sub-ms detection is possible. We also envision that this latency can be significantly improved when additional techniques such as randomized hashing and/or successive approximation methods are employed.

The next steps in our work are integrating an efficient scheme to track a given neighbor, investigating successive approximation methods and randomized hashing to further speed up the discovery time, and the building of a hardware prototype that achieves sub-millisecond detection time. Finally, all these components will be combined into a generalized rendezvous protocol.

Acknowledgment: This work has been supported in part by NSF under grant number CNS-1657279.

References

[1] W. Davis. 18 - optical mems for displays in portable systems. In D. Uttamchandani, editor, *Handbook of Mems for Wireless and Mobile Applications*, Woodhead Publishing Series in Electronic and Optical Materials. Woodhead Publishing, 2013.

[2] J. Scarlett. Noisy adaptive group testing: Bounds and algorithms. *IEEE Transactions on Information Theory*, 2018.

[3] N. H. Bshouty, V. E. Bshouty-Hurani, G. Haddad, T. Hashem, F. Khoury, and O. Sharafy. Adaptive group testing algorithms to estimate the number of defectives. *Proceedings of Algorithmic Learning Theory*, 07–09 Apr 2018.

[4] Y. Cheng, D.-Z. Du, and F. Zheng. A new strongly competitive group testing algorithm with small sequentiality. *Annals of Operations Research*, Jun 2015.

[5] D.-Z. Du and H. Park. On competitive group testing. *SIAM J. Comput.*, October 1994.

[6] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk. Single-pixel imaging via compressive sampling. *IEEE Signal Processing Magazine*, March 2008.

[7] J. A. Tropp and A. C. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on Information Theory*, Dec 2007.

[8] R. G. Baraniuk. Compressive sensing [lecture notes]. *IEEE Signal Processing Magazine*, July 2007.

[9] M. Kavehrad. Mems-based reconfigurable optical wireless networking in data centers. *2017 IEEE Photonics Conference (IPC)*, Oct 2017.

[10] Y. L. Peng Deng, Mohsen Kavehrad. Mems-based beam-steerable fso communications for reconfigurable wireless data center. *Proc. SPIE*, 2017.

[11] M. Curran, M. S. Rahman, H. Gupta, K. Zheng, J. Longtin, S. R. Das, and T. Mohamed. Fsonet: A wireless backhaul for multi-gigabit picocells using steerable free space optics. MobiCom '17.

[12] V. Milanovi, A. Kasturi, N. Siu, M. Radojii, and Y. Su. memseye for optical 3d tracking and imaging applications. *2011 16th Intl. Solid-State Sensors, Actuators and Microsystems Conf.*, June 2011.

[13] V. M. and. Fast and high-precision 3d tracking and position measurement with mems micromirrors. *2008 IEEE/LEOS Intl. Conf. on Optical MEMs and Nanophotonics*, Aug 2008.

[14] P. Brandl, S. Schidl, A. Polzer, W. Gaberl, and H. Zimmermann. Optical wireless communication with adaptive focus and mems-based beam steering. *IEEE Photonics Technology Letters*, Aug 2013.

[15] M. Khan and M. Yuksel. Maintaining a free-space-optical communication link between two autonomous mobiles. WCNC '14.

[16] M. Khan, S. Bhunia, M. Yuksel, and S. Sengupta. Los discovery in 3d for highly directional transceivers. MILCOM '16.

[17] M. Khan and M. Yuksel. In-band los discovery between drones using highly directional transceivers. DroNet'18.

[18] D.-Z. Du and F. K. Hwang. *Combinatorial Group Testing and Its Applications*, 2000.

[19] D.-Z. Du, G.-L. Xue, S.-Z. Sun, and S.-W. Cheng. Modifications of competitive group testing. *SIAM J. Comput.*, February 1994.

[20] J. Schlaghoff and E. Triesch. Improved results for competitive group testing. *Combinatorics, Probability and Computing*, 2005.