

Randomized Multi-Pulse Time-of-Flight Mass Spectrometry

Michael G. Moore, Andrew K. Massimino, and Mark A. Davenport
School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, Georgia
Email: {mmoore90,massimino,mdav}@gatech.edu

Abstract—Mass spectrometry is one of the primary methods for chemical analysis and serves as a fundamental tool in numerous scientific disciplines. In this paper we consider the design of time-of-flight mass spectrometers, which produce a stream of measurements which can be modeled by a convolution between the mass spectrum of interest and a specified pulsing pattern. Our goal is to reduce the total time necessary to analyze a sample to a given precision (or equivalently, given a fixed amount of time, to obtain a more precise estimate of the sample). We can do this by leveraging the structure that exists in typical mass spectra. In particular, since any given substance is usually composed of a relatively small number of distinct molecules, mass spectra tend to be relatively sparse. In this paper we perform an analysis of an idealized model of a time-of-flight mass spectrometer which uses a randomized pulsing pattern. Such an architecture has the potential to enable a new tradeoff between acquisition time and precision/dynamic range. We show that under certain natural conditions on the randomized scheme — namely, that the system does not pulse too often — this construction will lead to a system which satisfies certain desirable properties that are sufficient to ensure that sparse recovery is possible. In particular, we show that with high probability, the system will satisfy the conditions of a *bipartite expander graph* provided the pulsing rate is not too large. We then conclude with a range of simulations that support our theoretical analysis and demonstrate the practical viability of this approach.

I. INTRODUCTION

The goal of mass spectrometry is to determine the chemical composition of a substance by first ionizing a sample (the *analyte*) and then measuring the *mass-to-charge ratio* (MCR) of the constituent ions. This technique is extremely useful in numerous applications where precise chemical analysis is desirable, including chemistry, biology, and genetics [1]. A linear record of these ions, arranged by MCR, is called a *mass spectrum*.

In *time-of-flight mass spectrometry* (TOFMS) specifically, the measurement apparatus consists of an ionizer, a set of charged plates, a drift region, and a detector. The ionizer converts the analyte into ions, which are then accelerated by briefly charging a set of electric plates. Each ion gains a fixed kinetic energy from the acceleration, resulting in a velocity inversely proportional to the square root of the MCR. The accelerated ions travel through the drift region and arrive at the detector at different times depending on their velocity, which in turn depends on the MCR, allowing one to easily observe the mass spectrum. This setup is depicted in Fig. 1.




Fig. 1. A simplified view of a TOFMS device. Ions are produced and are then accelerated by a pulsed electric field. They travel through a drift region before arriving at a detector. The detector records the *mass spectrum*, the abundance of each ion in the sample. See [2] for more details.

A traditional TOFMS device will “fire” an electric pulse (accelerating the ions) and “listen” for the arrivals at the detector before repeating the process. Repetition is necessary because the output of a typical TOFMS system is often relatively noisy. Sources of this noise include a variable number of ion arrivals for a pulse event and modest variations in drift times for ions (two ions of the same species might arrive with slightly different delays). Thus one must combine the responses of many pulses to ensure a sufficiently precise output. In a traditional TOFMS system, the time between pulses is set to be sufficiently large to ensure that all the ions have time to travel through the device before the next pulse, so as to prevent “aliasing” of ions from different pulse events.

While simplifying the back-end processing, this design choice has significant costs in that it either greatly increases the total time necessary to analyze a sample or greatly limits the dynamic range (the allowable delay between the fastest and slowest ions). This can be problematic in certain circumstances, such as when examining a substance which may be dynamically evolving on a relatively short time scale.

Fortunately, mass spectra tend to be highly structured — in particular, they are often *sparse* (meaning that many MCRs are not represented in any individual analyte), so most of the system’s time is spent simply waiting while receiving no signal at the detector. In this paper, we consider a *multi-pulse* TOFMS system which exploits this sparsity and instead overlaps many pulses and sensing windows, reducing the total time required to determine chemical composition. Equivalently, multi-pulse TOFMS allows us to potentially increase the signal-to-noise ratio or improve the dynamic range given




Fig. 2. The observation matrix A is binary and Toeplitz and implements linear convolution with vector a . Each successive column is the previous column shifted down one index.

a total time constraint. This is analogous to the use of a staggered pulse repetition frequency to increase unambiguous range in RADAR systems [3]. We note that a variety of approaches to TOFMS have been proposed in the past that also exploit this structure (see, e.g., [4]–[6]).

Here we focus on the technique of randomizing pulse times in order to overlap the “aliases” in a way that still allows us to distinguish different ions [6]. Although it is desirable to pulse frequently to maximize the number of ions measured (and hence increase the signal-to-noise ratio), if too many different ions strike the detector simultaneously it will not be possible to accurately assign the arrivals to the correct MCR bins. The frequency of pulses is a key parameter that effects a trade-off between analysis time and spectrum fidelity, but which has not yet been addressed in the existing literature. We believe that our theoretical analysis will give intuition as to how often we can pulse while still retaining the ability to identify the distinct ions composing a sample. Specifically, we do this by establishing conditions under which a random binary convolution matrix is the adjacency matrix of a bipartite expander graph.

II. A MATHEMATICAL MODEL OF MULTI-PULSE TOFMS

Let $x \in \mathbb{R}_+^n$ denote the nominal (time-discretized) response of the analyte, where each index determines the frequency of an ion species with a given TOF. If the analyte is excited by the pulse firing sequence $a \in \{0, 1\}^m$ (every 1 in a determines when a pulse is fired), then the detector response is the (noisy) convolution $y = a * x + \nu$ with $y, \nu \in \mathbb{R}^{m'}$ and $m' = m + n - 1$. Because convolution is a linear operation, we can construct a matrix A such that $Ax = a * x$. The sensing matrix $A \in \{0, 1\}^{m' \times n}$ is Toeplitz with a (zero-padded to length m') as the first column, as shown in Fig. 2.

The design of the sensing matrix is thus limited to the choice of a . While [6] selects the time between pulses to be a uniform random variable with a prescribed average rate, the following analysis is based on selecting a to be a Bernoulli random vector (in order to make firing times independent). We make the choice

$$a_i = \begin{cases} 1 & \text{with probability } p \\ 0 & \text{with probability } 1 - p \end{cases}$$

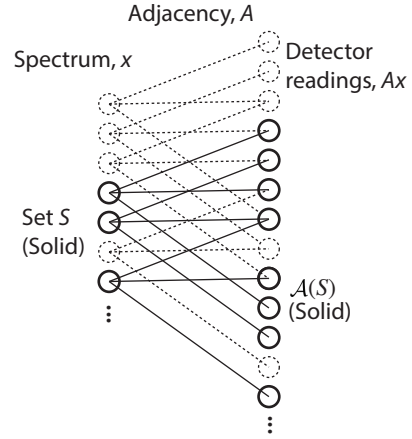


Fig. 3. A bipartite graph has a set of left and right vertices that are connected through an adjacency matrix A . The operator $\mathcal{A}(S)$ denotes the right vertices adjacent to S , a set of left vertices. The identical “pattern” of connections from each vertex corresponds to a Toeplitz choice of A , as in our case.

for $i \in \{1 \dots m\}$, which reduces the design of the sensing matrix to the choice of the single parameter p . We will assume that the parameters m and n (and m') are fixed. We also assume that the signal x (the mass spectrum) has sparsity k , by which we mean that no more than k of the n possible TOF bins are occupied by ion species for a given analyte.

Our goal is to understand how often we can pulse our system (i.e., how large we can set p) while still ensuring that we can reliably estimate the vector x (or the location of its nonzeros, depending on our error metric). One powerful technique to ensure that it is possible to recover sparse signals from our measurements is to show that A is the adjacency matrix of a bipartite expander graph [7]–[9].

A d -regular bipartite graph has n left vertices that each connect to d of the m right vertices (left degree d), as in Fig. 3. The adjacency matrix A of the graph determines how the left vertices (columns) are connected with the right vertices (rows) of the graph. In a TOF system, this matrix is Toeplitz, which means that each left vertex has the same “pattern” of connections (shifted according to the index).

A d -regular bipartite expander graph satisfies the property that for any choice of k left vertices the number of “adjacent” (connected) right vertices is almost as large as dk , suggesting that there are very few right vertices that connect to more than one of the chosen k left vertices. More formally, let S be a set of left vertices and $\mathcal{A}(S)$ return the set of right vertices connected to S by the adjacency matrix A . If $\mathcal{E}(k, \epsilon)$ is the set of all bipartite expander graph adjacency matrices and d is the number of nonzeros in each column of A then

$$A \in \mathcal{E}(k, \epsilon) \Leftrightarrow |\mathcal{A}(S)| \geq (1 - \epsilon)d|S| \quad \forall |S| \leq k. \quad (1)$$

Once we know that our sensing matrix is the adjacency matrix of a bipartite expander graph, this leads to a range of guarantees regarding our ability to recover sparse signals [7], [8]. Such matrices also satisfy the RIP-1 property [9].

III. SENSING MATRIX AS AN EXPANDER GRAPH

We now show that under certain natural conditions on the choice of p , the sensing matrix A formed via the Toeplitz extension of a vector a with Bernoulli(p) entries will satisfy the conditions of a bipartite expander graph with high probability. In particular, Theorem 1 suggests that we must obey the constraint $p \lesssim \frac{\epsilon}{k}$. Thus, because the constant ϵ for an expander must satisfy $\epsilon < 1$, it is reasonable to choose $p \approx \frac{1}{k}$ (within some small constant factor).

Theorem 1. *Suppose that $m = \Omega(k^3 \log \frac{n}{k})$ and that $a \in \{0, 1\}^m$ has elements drawn as independent Bernoulli(p) random variables. Then the matrix A corresponding to linear convolution with a satisfies $A \in \mathcal{E}(k, \epsilon)$ with high probability provided that $kp < 2\epsilon\beta$ for some constant $\beta < 1$.*

Proof. Conversely to (1), a matrix is not an expander if there exists a set of left vertices S with $|S| = \ell$ for some $\ell \leq k$ such that that $|\mathcal{A}(S)| < (1 - \epsilon)d\ell$. Let s_i denote a single left vertex from set S . Note that $|\mathcal{A}(s_i) \cap \mathcal{A}(s_j)| = A_i^T A_j$, where A_i is the s_i^{th} column of A . Thus, we can use the inequality

$$|\mathcal{A}(S)| \geq \sum_{i=1}^{\ell} |\mathcal{A}(s_i)| - \sum_{\substack{i=1 \\ j=i+1}}^{\ell} A_i^T A_j$$

and the fact that $|\mathcal{A}(s_i)| = d$ for all i to arrive at the bound

$$\begin{aligned} \mathbb{P}(A \notin \mathcal{E}(k, \epsilon)) &\leq \sum_{\ell=1}^k \binom{n}{\ell} \mathbb{P} \left(\sum_{\substack{i=1 \\ j=i+1}}^{\ell} A_i^T A_j > \epsilon d \ell \right) \\ &\leq \left(\frac{en}{k} \right)^k \max_{\ell \leq k} \mathbb{P} \left(\sum_{\substack{i=1 \\ j=i+1}}^{\ell} A_i^T A_j > \epsilon d \ell \right) \\ &\leq \left(\frac{en}{k} \right)^k \max_{\ell \leq k} \max_{1 \leq i < j \leq \ell} \frac{\ell^2}{2} \mathbb{P} \left(A_i^T A_j > \frac{2\epsilon d}{\ell} \right), \quad (2) \end{aligned}$$

where (2) follows by the application of the union bound.

Note that $A_i^T A_j = \sum_{h=1}^{m-|s_i-s_j|} a_h a_{h+|s_i-s_j|}$ and that, though dependent, the terms of this sum can be split into two separate sums of independent Bernoulli random variables. These form two (dependent) binomial random variables. For example, $a_1 a_2 + a_2 a_3 + a_3 a_4 \dots = (a_1 a_2 + a_3 a_4 \dots) + (a_2 a_3 + a_4 a_5 \dots) = \Lambda_1 + \Lambda_2$.

Using the Chernoff bound for a binomial random variable $\Lambda \sim \text{Bin}(\frac{m}{2}, p^2)$ with $t > \frac{1}{2}mp^2$ results in the bound

$$\mathbb{P}(\Lambda \geq t) \leq \exp \left(-\frac{1}{mp^2} \left(t - \frac{mp^2}{2} \right)^2 \right)$$

and applying this to the split sums using the union bound gives

$$\mathbb{P} \left(A_i^T A_j \geq \frac{2\epsilon\beta mp}{\ell} \right) \leq 2 \exp \left(-\frac{m}{\ell^2} \left(\epsilon\beta - \frac{\ell p}{2} \right)^2 \right) \quad (3)$$

with the requirement that $\epsilon > \frac{\ell p}{2\beta}$. To use (3) in (2), we need to ensure that $\beta mp \leq d$. It is easy to recognize that $d \sim$

$\text{Bin}(m, p)$ and that for $\beta < 1$

$$\mathbb{P}(d \leq \beta mp) \leq \exp \left(-\frac{mp}{2} (1 - \beta)^2 \right). \quad (4)$$

A union bound over (3) and (4) ensures that for any $\epsilon > \frac{kp}{2\beta}$ and $\beta < 1$,

$$\begin{aligned} \mathbb{P}(A \notin \mathcal{E}(k, \epsilon)) &\leq \exp \left(-\frac{mp}{2} (1 - \beta)^2 \right) + \\ &\exp \left(k \log \frac{en}{k} + 2 \log k - \frac{m}{k^2} \left(\epsilon\beta - \frac{kp}{2} \right)^2 \right). \quad (5) \end{aligned}$$

This probability is small provided that $m = \Omega(k^3 \log \frac{n}{k})$. \square

IV. SIMULATIONS

To complement the theoretical analysis above, we conduct a range of simulations under a simple noise model where the arrival of ions at the detector is modeled as a Poisson process with rate determined by the sensing matrix A and the underlying mass spectrum x . In particular, we will model our observations as

$$y_i \sim \text{Poisson}(\lambda_i) \quad \lambda \triangleq Ax + \gamma \quad (6)$$

for a known background noise power $\gamma \in \mathbb{R}_+$ and k -sparse unknown signal $x \in \mathbb{R}_+^n$.¹ The observation model in (6) leads to the negative-log-likelihood

$$\mathcal{L}(x|y) = \sum_{i=1}^{m'} \lambda_i - y_i \log \lambda_i + \log y_i!$$

We can then compute the constrained and sparsity-regularized maximum-likelihood estimate via:

$$\hat{x} = \arg \min_{x \in \mathbb{R}_+^n} \mathcal{L}(x|y) + \alpha \|x\|_1 \quad (7)$$

using SPIRAL [11] or other convex optimization techniques. We do not use the more common regularized least-squares optimization for this problem because least-squares is poorly suited to our Poisson observation model [12].

To examine the practical impact of the choice of p on our ability to recover k -sparse signals under the observation model in (6) we conduct synthetic simulations over a range of combinations of k and p . We set $m' = 50\,000$, $n = 5\,000$, and $\gamma = 0.05$. We generate candidate k -sparse signals x in which the locations of the k nonzeros are selected uniformly at random and the values are drawn from the truncated exponential distribution with range $[0.1, 1]$ (so that, for example, the chance of being in the range $[0.1, 0.2]$ is equal to the chance of being in the range $[0.2, 0.4]$). After generating our observations y according (6), we compute an estimate \hat{x} via the optimization problem (7). As a post-processing step (to aid in detecting the

¹It is possible to consider a slightly more realistic model as in [6], which employs a compound Poisson–Erlang model that also accounts for additional measurement noise in the detector/sampling system. Moreover, in practice ion impacts tend to spread across multiple detector time-samples (rather than having an ideal impulsive response) and ions also possess a small amount of variability in drift times. Thus, in a real TOFMS system the use of a cluster-sparse model might be more appropriate [10]. Here we ignore these possible refinements in the interest of simplicity.




Fig. 4. Normalized ℓ_1 error, miss rate, and false discovery rates for regularized maximum-likelihood recovery of the system. The dashed line corresponds to the choice $kp = 2$ and closely aligns with the optimum p over a wide range of k in terms of error and miss rate.

locations of the nonzeros) we threshold \hat{x} at one-quarter the minimum abundance so that any entry $\hat{x}_i < 0.025$ is set to zero. The regularization parameter used is $\alpha = m'p$.

We examine the performance of this system over 500 trials for each (k, p) pair, with 6 values of k per decade and 8 values of p per decade. We report the average normalized ℓ_1 error $\frac{1}{k}\|x - \hat{x}\|_1$, miss rate $\frac{1}{k}\sum_i \mathcal{I}(x_i > 0, \hat{x}_i = 0)$, and false discovery rate $\frac{1}{\sum_i \mathcal{I}(\hat{x}_i > 0)}\sum_i \mathcal{I}(\hat{x}_i > 0, x_i = 0)$ of the recovery.² In Fig. 4 we show how these three statistics vary as a function k and p . All MATLAB code used for these simulations is available for download at <http://users.ece.gatech.edu/~mdavenport/software/>.

V. DISCUSSION

As we can see from Fig. 4, the ℓ_1 error and miss rate are optimized for most k by choosing $p \approx \frac{2}{k}$ (though this will vary slightly with m and n). The contours of the false discovery rate in Fig. 4 follow a slightly different trend, but also appear to be minimized by choosing $p \propto \frac{1}{k}$. This agrees closely with the analysis performed in Section III.

This result also has an intuitive explanation: the quantity kp corresponds to the expected number of unique species simultaneously arriving at the detector. For $kp \ll 1$, most observation intervals have zero impact events and one would expect that increasing p to take better advantage of the limited observation time should lead to improved estimates of the mass spectrum. Conversely, for $kp \gg 1$ multiple species are likely to be arriving simultaneously in nearly all observation intervals. It is not hard to imagine that for large kp , the constant barrage of ions will make it difficult to disentangle the signal and determine the abundance, or even presence, of any individual species.

Note that a traditional TOFMS system, in which the drift region is allowed to clear completely between pulses, is roughly akin to the choice $p = \frac{1}{n}$. Thus, using a uniformly random pulsing pattern, one can accelerate acquisition by a factor of $\frac{n}{k}$ compared to traditional TOFMS. One can imagine

that similar results concerning the pulse-rate might apply when the time between pulses is restricted to some sufficiently broad range (since most spectrometers have a minimum interval between pulses) or when ion drift times suffer small variations.

ACKNOWLEDGMENTS

This work was supported by grants NRL N00173-14-2-C001, AFOSR FA9550-14-1-0342, NSF CCF-1409406 and CCF-1350616, and gifts from LexisNexis Risk Solutions and Lockheed Martin.

REFERENCES

- [1] R. Cotter, "The new time-of-flight mass spectrometry," *Analytical Chemistry News & Features*, vol. 71, no. 14, pp. 445A–451A, 1999.
- [2] M. Wolff and W. Stephens, "A pulsed mass spectrometer with time dispersion," *Review of Scientific Instruments*, vol. 24, no. 8, pp. 616–617, 1953.
- [3] M. Richards, *Fundamentals of Radar Signal Processing*. McGraw-Hill, 2005.
- [4] C. Cocq, "High duty cycle pseudo-noise modulated time-of-flight mass spectrometry," 2001, US Patent 6,198,096.
- [5] A. Brock, N. Rodriguez, and R. Zare, "Hadamard transform time-of-flight mass spectrometry," *Analytical Chemistry*, vol. 70, no. 18, pp. 3735–3741, 1998.
- [6] M. Ibrahim, A. Montanari, and G. Moore, "Accelerated time-of-flight mass spectrometry," *IEEE Trans. Signal Processing*, vol. 62, no. 15, pp. 3784–3798, 2014.
- [7] R. Berinde, P. Indyk, and M. Ružić, "Practical near-optimal sparse recovery in the L1 norm," in *Proc. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Sep. 2008.
- [8] A. Gilbert and P. Indyk, "Sparse recovery using sparse matrices," *Proc. IEEE*, vol. 98, no. 6, pp. 937–947, 2010.
- [9] R. Berinde, A. Gilbert, P. Indyk, H. Karloff, and M. Strauss, "Combining geometry and combinatorics: A unified approach to sparse signal recovery," in *Proc. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Sep. 2008.
- [10] V. Cevher, P. Indyk, C. Hegde, and R. Baraniuk, "Recovery of clustered sparse signals from compressive measurements," in *Proc. Sampling Theory and Applications (SampTA)*, Marseille, France, May 2009.
- [11] Z. Harmany, R. Marcia, and R. Willett, "This is SPIRAL-TAP: Sparse Poisson Intensity Reconstruction Algorithms – Theory and Practice," *IEEE Trans. Image Processing*, vol. 21, no. 3, pp. 1084–1096, 2012.
- [12] D. Motamedvaziri, M. Rohban, and V. Saligrama, "Sparse signal recovery under Poisson statistics," in *Proc. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Oct. 2013, pp. 1450–1457.

²The ℓ_1 error is presented because it is more sensitive to small errors in the recovery than the ℓ_2 error and is less easily dominated by the performance on only the largest few coefficients.