

# Randomized Group Testing for Acoustic Source Localization

William Mantzel<sup>e</sup>, Justin Romberg<sup>e</sup>, and Karim Sabra<sup>m</sup>

<sup>e</sup>Department of Electrical and Computer Engineering

<sup>m</sup>Department of Mechanical Engineering

Georgia Institute of Technology

## ABSTRACT

Undersea localization requires a computationally expensive partial differential equation simulation to test each candidate hypothesis location via matched filter. We propose a method of batch testing that effectively yields a test sequence output of random combinations of location-specific matched filter correlations, such that the computational run time varies with the number of tests instead of the number of locations. We show that by finding the most likely location that could have accounted for these batch test outputs, we are able to perform almost as well as if we had computed each location’s matched filter. In particular, we show that we can reliably resolve the target’s location up to the resolution of incoherence using only logarithmically many measurements when the number of candidate locations is less than the dimension of the matched filter. In this way, our random mask pattern not only performs substantially the same as cleverly designed deterministic masks in classical batch testing scenarios, but also naturally extends to other scenarios when the design of such deterministic masks may be less obvious.

**Keywords:** matched-field processing, undersea detection, group testing, compressed sensing

## 1. INTRODUCTION

Matched field processing continues to serve as one of the most widely used methods for localizing undersea targets. However, as the models governing undersea acoustic interactions become more sophisticated, often requiring fine-grain solutions to partial differential equations, the tradeoff between run time and performance begins to worsen, perhaps unnecessarily. When the source, receiver, and fluid are stationary with respect to each other, the principle of *reciprocity* can be employed to solve these differential equations on a *per-receiver* basis instead of a *per-source* basis, resulting in dozens rather than thousands of PDEs to be solved. In other cases, this principle may only approximately apply and it may be necessary to use another technique to mitigate the computational bottleneck.

Here we consider efficiently testing groups of locations simultaneously to reduce the number of PDE simulations that must be performed. Though the rigorous analysis has come from recent years, the idea of group testing dates back to Syphilis testing during World War II. Robert Dorfman and David Rosenblatt noted the wastefulness of performing so many blood tests to detect Syphilis when it was almost certain that any given blood test would turn up negative.<sup>1</sup> Instead, they pointed out that the blood from several patients should be pooled together and tested simultaneously. The testing procedure was sensitive enough that even great dilutions of the antigen could be positively detected, so that an infection could be detected from even a single patient out of a large mixture. The key idea was that if a test turned up positive, at most 1 test was wasted, but a negative test ended up clearing the names of all whose blood was mixed. Note that a positive test is not necessarily wasted if all but one of the people whose blood was mixed were exonerated in other mixtures that tested negative.

The intuition shared in this approach is echoed in now-conventional wisdom from the Information Theory community that a bit is well spent if it distinguishes between two equally likely events. In this case, the “bit” is a test that either turns up positive or negative. In particular, it is necessary to pool a number of patients inversely proportional to the incidence rate to ensure that a group test is roughly equally likely to be positive or negative (e.g. a group on the order of 100 patients when 1% of the patients have syphilis).

---

E-mail: {willem, jrom, karim.sabra}@gatech.edu

Although it is simple to construct a single test that is equally likely positive as negative, it is somewhat more difficult to construct a set of tests that will carry the most information conditioned upon information already conveyed by previous tests. In other words, the results of one group test may be reliably predicted from a set of other group tests. There are clever ways to construct a set of group tests that try to maximize the information. However, one of the simplest techniques that remains one of the most robust and versatile is to simply construct each group as a random subset of the whole. It is this random construction that we will explore and analyze in this paper.

There has been much interest lately in using random measurements to maximize the amount of information captured in the measurement process. Indeed, recent work in compressed sensing has shown how random measurements may be used to effectively capture and reconstruct sparse or compressible information, shifting the workload from hardware to software, motivated by the assumption that hardware measurements are costly compared to software computations.<sup>2</sup> This paper employs a variation on the theme, relieving the computational workload in software instead of the sensing workload in hardware.

## 2. RANDOMIZED GROUP TESTING

Our problem is stated as follows. The signal could have been emitted from any of  $M$  hypothetical locations specified by range and depth  $\{(r_m, d_m)\}$  with corresponding Green's function  $G_m$  that is a column vector defined by the concatenation of the discrete-frequency responses from source location  $m$  to the set of receivers.  $\mathbf{G} \in \mathbb{R}^{N \times M}$  is defined as the matrix such that the  $m$ th column is  $G_m$ , normalized to have unit norm. The actual location is  $\gamma$  and the actual received signal is  $y = G_\gamma$ .

The standard approach of Matched-Field Processing (MFP) is to estimate the source location  $m$  as the one that maximizes the inner product  $\langle y, G_m \rangle$ .<sup>3</sup> Because these vectors are the vertical concatenation of their constituent frequency responses, this inner product effectively sums the frequency correlations over all receivers. Though these correlations may be high inside a small neighborhood around  $\gamma$ , they generally decay in magnitude very quickly away from this *region of ambiguity* and are generally expected to be relatively small.

Because computing these individual correlations with  $G_m$  directly is computationally costly, we instead compute the correlations with  $K$  combinations of these Green's functions, yielding the vector  $z \in \mathbb{R}^K$ . These combinations are described in the columns of the *mask matrix*  $\Phi \in \mathbb{R}^{M \times K}$  that is constructed as a random matrix whose entries are independently drawn from  $\{-1, 1\}$  with equal likelihood. The  $k$ th such measurement is described as:

$$z(k) = \left\langle \sum_m \Phi_k(m) G_m, y \right\rangle \quad (1)$$

$$= \langle \mathbf{G} \Phi_k, y \rangle \quad (2)$$

$$= \langle \Phi_k, \mathbf{G}^T G_\gamma \rangle. \quad (3)$$

These batch correlations are then described in matrix notation as:

$$z = (\mathbf{G} \Phi)^T y = \Phi^T \mathbf{G}^T G e_\gamma, \quad (4)$$

where the delta vector  $e_\gamma$  is one at the true location of the transmitter  $m = \gamma$  and zero elsewhere.

Intuitively, this combination  $\mathbf{G} \Phi_k$  in some sense tests groups of Green's functions simultaneously. Indeed, if  $\Phi$  were a 0/1 i.i.d. Bernoulli matrix, each test vector  $\mathbf{G} \Phi_k$  is viewed as the sum of some random subset of columns in  $\mathbf{G}$ . This may be conceived as a "chorus" of sources transmitting from these subset of locations for each test. The key idea is that it is as computationally intensive to simulate a "chorus" of sources as a single source. By computing the "full-chorus" test measurement,  $\bar{z} = 1^T \mathbf{G}^T G_\gamma$  (where  $1^T$  denotes a row vector of ones), the 0/1 measurements  $z_k$  may be easily converted to  $\pm 1$  measurements as:  $2z_k - \bar{z}$ . Though the 0/1 measurements are more useful for purposes of illustration, the  $\pm 1$  measurements are more useful for purposes of analysis, as well as generally giving better detection results, for reasons analogous to the choice of antipodal pairs in BPSK communication systems over on/off pairs.




Figure 1. Shown in (a) is the ambiguity function  $\mathbf{G}^T \mathbf{G}_\gamma$  over 512 locations. Shown in (b) is a random approximation to this ambiguity function  $\Phi \Phi^T \mathbf{G}^T \mathbf{G}_\gamma$  using 16 tests. Although the latter is quite a poor approximation visually, it shares one key characteristic with the former: they both have peak values at  $(r_\gamma, d_\gamma) = (6\text{km}, 70\text{m})$ .

### 3. RECOVERY

Though computing  $z = (\mathbf{G}\Phi)^T y$  is much easier than computing  $\mathbf{G}^T y$  directly, it is of lower dimension and generally less informative with regard to the location of the source. Nevertheless, there are techniques we can use to make reliable estimates for the source location.

A straightforward method simply searches the  $\Phi$  matrix for a row with a matching sign sequence. There is an analogy to group testing in looking for the agent common to all tests that returned “positive”. The rationale here is that because  $\langle G_m, G_\gamma \rangle$  is small in magnitude except where  $m = \gamma$  where the correlation peaks at 1, each  $z(k)$  component will be dominated by  $\Phi_k(\gamma)$  and will tend to have the same sign. One problem with such an approach is that the solution is not unique when another row shares the same sign sequence with the  $\gamma$ th row. This uniqueness of the  $\gamma$ th row occurs with probability  $(1 - 2^{-K})^M$ , so that this possibility of non-uniqueness may only be reliably ruled out when  $M < 2^K$ .

On the other hand, it is also possible that there is no row with matching sign sequence. This means that for at least one of the tests, the off-diagonal terms of  $\mathbf{G}^T \mathbf{G}$  were able to dominate the diagonal term causing a flipped sign. One sensible approach to mitigate this is to look for signs to flip back in  $z$  that would result in a matching row in  $\Phi$ , starting with the elements of  $z$  with smallest magnitude, i.e. the ones that were the most likely flipped. This process may be formalized by simply searching for the row in  $\Phi$  that is closest in Euclidean space to the test measurement  $z$ :

$$\arg \min_m \|\Phi^T e_m - z\|. \quad (5)$$

This minimization is equivalent to a maximization over  $\Phi \Phi^T \mathbf{G}^T \mathbf{G}_\gamma$ :

$$\arg \min_m \|\Phi^T e_m - z\|^2 \quad (6)$$

$$= \arg \min_m K - 2e_m^T \Phi \Phi^T \mathbf{G}^T \mathbf{G}_\gamma + \|z\|^2 \quad (7)$$

$$= \arg \max_m e_m^T \Phi \Phi^T \mathbf{G}^T \mathbf{G}_\gamma. \quad (8)$$

In other words, we are looking for the row in  $\Phi$  that is the most correlated with our test measurement vector  $z$ . Such an approach has parallels in compressed sensing solutions (e.g. iterative thresholding<sup>4</sup>), and seems to be a reasonable approach. Indeed, because  $\mathbb{E}[\Phi \Phi^T] = K\mathbf{I}$ , we have in expectation that  $\mathbb{E}[\Phi \Phi^T \mathbf{G}^T \mathbf{G}_\gamma] = \mathbf{G}^T \mathbf{G}_\gamma$ , so this approach seems to give a reasonable approximation to the desired ambiguity function, as illustrated in Fig. 1.




Figure 2. The probability of unique recovery for various scenarios are shown here for either 10 or 20 receivers and either 4 or 64 frequency points.

#### 4. EXPERIMENTAL RESULTS

To simulate the performance of our proposed approach in a plausible environment, we ran a series of tests in a constant-depth ocean environment modeled using normal mode equations,<sup>5</sup> shown in Fig. 2. Here we placed either 10 or 20 uniformly spaced receivers between 0 and 200 meters. The frequency response was taken over either 4 or 64 discrete frequency points per receiver. We used between 1 and 64 test measurements to try to estimate the source location. The  $M$  locations were chosen as some subset of a grid that, at the finest scale, varies in depth from 10 to 190 meters in 12 meter increments, and varies in range from 5 to 8.1 kilometers in 100 meter increments, yielding  $M = 512$  locations at the finest scale. To test other values of  $M$  decreasing in powers of 2 we simply subsampled this grid alternatively in range and depth. For each fixed value of  $M$  and corresponding grid, we generated 100 different random masks  $\Phi$ . For each mask  $\Phi$ , we averaged the probability of success over all possible locations  $\gamma$ , i.e. the fraction of  $\Phi\Phi^T\mathbf{G}^T\mathbf{G}$  (over all columns) that are diagonally dominant. For our purposes, success was defined when the *unique* maxima was achieved at  $\gamma$ .

The main point on these graphs to note is that the curves only shift to the right by 1 test every time  $M$  is doubled up to a certain point where  $M \simeq N$ . This is the point of diminishing returns, where each further factor of 2 increase in  $M$  requires a matching factor of 2 increase in  $K$ . More specifically, this point of transition occurs when the rank of the Grammian matrix  $\mathbf{G}^T\mathbf{G}$  reaches its peak. Incremental improvements occur thereafter for further tests (increasing  $K$ ), as the Grammian becomes progressively more well-conditioned.

## 5. CONCLUSION

In conclusion, we have leveraged a strategy for our purposes that is as familiar to classical group testing as it is to modern compressed sensing, using a small number of random compilations to tease out the needle in the haystack. Although we have presented this approach in the context of acoustic localization, this work may be generalized to any detection scenario with a low incidence rate (sparsity) such that it takes as much effort to compute a weighted sum of inner products as a single inner product.

These results presented suggest a wide applicability in practical scenarios. Indeed, the occasional high competing correlations (in some cases as much as 0.4) did not necessarily prevent localization.

Ultimately though, we are constrained by the fundamental limits in the ambiguity function, and the overall performance of our approach essentially depends on how diagonally dominant the Gramian (correlation) matrix is. At best, our approach is able to reliably localize the source inside a number of candidate locations using only logarithmically many measurements. At worst, even if we are not able to reliably and uniquely determine the location when there is some ambiguity of the location due to correlated Green's functions, we are at least able to substantially reduce our search space over which to run a classical Matched-Field Processing algorithm, using only on the order of a dozen tests.

## ACKNOWLEDGMENTS

Much thanks to Alexis Dumortier for his early work on this project and to Shaun Anderson for producing source code to simulate the underwater acoustics and for patiently enduring a variety of questions on the matter.

## REFERENCES

- [1] Du, D. and Hwang, F., [*Combinatorial group testing and its applications*], World Scientific Pub (2000).
- [2] Candès, E., Romberg, J., and Tao, T., “Stable signal recovery from incomplete and inaccurate measurements,” *Communications on Pure and Applied Mathematics* **59**(8), 1207 (2006).
- [3] Baggeroer, A., Kuperman, W., and Mikhalevsky, P., “An overview of matched field methods in ocean acoustics,” *IEEE Journal of Oceanic Engineering* **18**(4), 401–424 (1993).
- [4] Blumensath, T. and Davies, M., “Iterative thresholding for sparse approximations,” *Journal of Fourier Analysis and Applications* **14**(5), 629–654 (2008).
- [5] Schroeder, M., Rossing, T., Dunn, F., Hartmann, W., Campbell, D., and Fletcher, N., [*Springer Handbook of Acoustics*], Springer Publishing Company, Incorporated (2007).