



Group testing meets traitor tracing

Peter Meerwald, Teddy Furon

► To cite this version:

Peter Meerwald, Teddy Furon. Group testing meets traitor tracing. ICASSP, May 2011, Prague, Czech Republic. 10.1109/ICASSP.2011.5947280 . inria-00580899

HAL Id: inria-00580899

<https://inria.hal.science/inria-00580899>

Submitted on 21 Oct 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

GROUP TESTING MEETS TRAITOR TRACING

Peter Meerwald and Teddy Furon

INRIA Rennes Bretagne Atlantique, TEMICS
Campus de Beaulieu, Rennes, France

ABSTRACT

This paper links two a priori different topics, group testing and traitor tracing. Group testing, as an instantiation of a compressed sensing problem over binary data, is indeed easier than traitor tracing because the mixing model is far simpler. State-of-the-art algorithms for traitor tracing, including the celebrated probabilistic Tardos code, are applied to the group testing problem. They yield better than or competitive performance when compared to state-of-the-art algorithms.

Index Terms— Compressed sensing, Fingerprinting.

1. INTRODUCTION

Group testing and traitor tracing aim at the same goal: retrieving the identity of very few people in a huge population. In group testing, these few people are contaminated by a given disease while in traitor tracing, these people are dishonest users illegally distributing copyrighted material. This section briefly introduces both subjects and motivates that both are indeed similar instantiations of a noisy compressed sensing problem over binary data, the remaining part of the paper proposes a design of the contact matrix and reconstruction algorithms inspired by traitor tracing techniques.

1.1. Group testing

Group testing is usually introduced as the epidemiology problem (many other applications are listed for instance in [1, Sec. 1]) where we need to identify a small set of virally-infected people in a large population. Typically, blood samples are tested to screen out the infected persons. If the size N of the population is large, the cost of individual tests is prohibitive and a group testing strategy is deployed. It consists in performing $T \ll N$ tests on several pools of blood samples. The contact matrix $\mathbf{M}$ is the binary $T \times N$ matrix indicating which blood sample is involved in a given test: $M_{ij} = 1$ means that test i uses the blood of individual j , $M_{ij} = 0$ otherwise. The results of the test are stored in a binary vector $\mathbf{y}$ ($y_i = 1$ if the i -th test is positive, 0 otherwise). Infected persons are to be identified from $\mathbf{y}$ and $\mathbf{M}$. $\mathcal{K}$ denotes the set of the indices of

these K persons, and the N -bit long vector $\mathbf{x}$ the unknown indicating vector (*i.e.* $x_i = 1$ if individual i is infected). From a compressed sensing viewpoint, the goal is to recover the sparse input vector $\mathbf{x}$ explaining the output $\mathbf{y} = \mathbf{M} \otimes \mathbf{x}$ ($\otimes$ denotes binary matrix multiplication). In words, the output vector $\mathbf{y}$ solely depends on the K columns of the restriction of $\mathbf{M}$ to $\mathcal{K}$.

Many papers have proposed bounds on the minimal number of tests or practical designs (creation of $\mathbf{M}$ and reconstruction of $\mathbf{x}$) based on combinatorics. The pool design is usually based on a disjunct or separable contact matrix. However, a recent trend is to look at this problem from a probabilistic point of view [1, 2]. The result of a test should be seen as a random variable Y reflecting the fact that the test may not be perfect. Cheraghchi *et al.* introduced the concept of *activation*: one infected person in the pool will independently trigger the test with probability Π [1]. Atia and Saligrama preferred to speak of a probability of *dilution* $u = 1 - \Pi$; they also envisaged another setup with *additive noise* (q is the false positive test probability) [2]. Sejdinovic and Johnson recently considered both dilution and additive noise [3].

The introduction of randomness gives birth to an information theoretical study which states the asymptotical scaling of the number of tests T as a function of (N, K, u, q) [2, 3]. A random construction of $\mathbf{M}$ has also been proposed in [1, 3]. However, this construction must know in advance K .

1.2. Traitor tracing

An application of traitor tracing, *a.k.a.* passive fingerprinting, is the Video-on-Demand scenario where a movie is distributed to N users. During the purchase, each user indeed receives a personal copy of the movie. The versioning process usually divides the movie into T video blocks, and it hides a bit into each block thanks to a watermarking technique. Therefore, a unique codeword $\mathbf{M}_j = (M_{1j}, \dots, M_{Tj})^t$ is sequentially hidden in the copy of the j -th user. A collusion of size K is a group of dishonest users $\mathcal{K} = \{j_1, \dots, j_K\}$ mixing their versions to forge a pirated copy. Since the partition of the movie is a priori not secret, the pirates can exchange video blocks. The T -bit long vector $\mathbf{y}$ decoded from the pirated content is thus a mix of the K codewords such that $y_i \in \{M_{ij} | j \in \mathcal{K}\}, \forall i \in [T]$ ($[T] \triangleq \{1, \dots, T\}$). Therefore,

Thanks to ANR agency for funding the project Medievals.

like in the group testing field, the output vector $\mathbf{y}$ depends on the K columns of the restriction of $\mathbf{M}$ to $\mathcal{K}$, and the goal of the accusation process is to identify these K dishonest users.

Similar to the work of Atia & Saligrama, Moulin has bounded the performance of a traitor tracing scheme from an information theoretic viewpoint [4]. In particular, he proved that the single decoder is outperformed by joint decoders. The first type of accusation takes a decision on a user-by-user basis whereas the latter type works on groups of users.

As for a practical construction of a traitor tracing code, Tardos was the first to exhibit an optimum solution [5], nowadays celebrated as the Tardos code, in the sense that its length T has the optimum scaling in $O(K^2 \log(N/\epsilon))$, where ϵ is the probability of accusing an innocent user. The code construction is fully probabilistic.

1.3. Similarities and differences

The similarity between the two fields lies in the mathematical model capturing how $\mathbf{y}$ depends on the K code-words $\{\mathbf{M}_j | j \in \mathcal{K}\}$. This is usually done in traitor tracing by the collusion strategy $\boldsymbol{\theta} \in [0, 1]^{K+1}$ with $\theta_\sigma \triangleq \mathbb{P}_{Y_i}[1 | \sum_{j \in \mathcal{K}} M_{ij} = \sigma]$, $0 \leq \sigma \leq K$. Translated into group testing terminology, a test result is a binary r.v. whose distribution depends on the number of infected involved in the test. For instance, applied to Sejdinovic's setup [3], we have $\theta_\sigma = 1 - (1 - q)u^\sigma$. If $(q, u) = (0, 0)$, then this $\boldsymbol{\theta}$ is known, in the field of traitor tracing, as the *all-one* collusion strategy.

In group testing, parameters (u, q) are known because they depend on the biological nature of the test. The problem then solely depends on the numbers of infected: the number of tests has been shown to be scaling as $O(K \log N)$ [2]. However, in traitor tracing, the colluders are free to choose their collusion strategy provided it complies with the marking assumption: $\theta_0 = 0$ and $\theta_K = 1$. In words, when they all have the i -th video block containing the same hidden symbol, this symbol is decoded in the pirated movie since they can't modify the embedded symbol. The performance of the traitor tracing code must be guaranteed whatever the collusion strategy. This is achieved by focusing on the worst case attack [6] or resorting to an accusation process statistically orthogonal to $\boldsymbol{\theta}$ as originally proposed by Tardos [5]. The optimum code length T has been shown to be scaling as $O(K^2 \log N/\epsilon)$, where ϵ is the probability of accusing an innocent, due to this nuisance parameter $\boldsymbol{\theta}$ [4]. This renders traitor tracing more difficult than group testing. On the other hand, traitor tracing aims at finding some colluders, whereas the goal of group testing is the exact recovery of the K infected.

2. THE CONTACT MATRIX

2.1. Generation

The population is composed of N individuals, among them K infected persons. There will be T tests. The contact matrix

is constructed like a Tardos fingerprinting code. Vector $\mathbf{p} = (p_1, \dots, p_T)^t$ are instances of i.i.d. r.v.s drawn according to $f(p)$ s.t. $0 < p_i < 1$, $\forall i \in [T]$. The goal of this section is to determine an appropriate distribution. Then, elements of the contact matrix are independently drawn s.t. $\mathbb{P}_{M_{ij}}[1] = p_i$.

2.2. Probabilities

As already mentioned, the assumptions of [3] lead to model $\theta_\sigma = 1 - (1 - q)u^\sigma$. We now focus on a test with $p_i = p$ and drop the index i . Thanks to the probabilistic construction of the contact matrix, the probability that σ infected are involved in this particular test is $\mathbb{P}_\Sigma[\sigma | p] = \binom{K}{\sigma} p^\sigma (1 - p)^{K - \sigma}$. $\mathbb{P}_A[a]$ denotes the probability that r.v. A takes the value a . Therefore, the probability that this test is positive is

$$\mathbb{P}_Y[1 | p, K] = \sum_{\sigma=0}^K \theta_\sigma \binom{K}{\sigma} p^\sigma (1 - p)^{K - \sigma} \quad (1)$$

which nicely simplifies to $1 - (1 - q)(1 - p + up)^K$. In the same way, assume we know the identity of $L < K$ infected, and among them, $\rho < L$ are involved in this test. The probability that the test is positive given this extra information is

$$\mathbb{P}_Y[1 | p, \rho, L, K] = \sum_{\sigma=\rho}^{K-L+\rho} \theta_\sigma \binom{K-L}{\sigma-\rho} p^{\sigma-\rho} (1-p)^{K-L-\sigma+\rho}$$

which simplifies to $1 - (1 - q)(1 - p + up)^{K-L} u^\rho$. In particular, for $L = 1$, we have:

$$\begin{aligned} \mathbb{P}_Y[1 | p, 0, 1, K] &= 1 - (1 - q)(1 - p + up)^{K-1} \quad (2) \\ \mathbb{P}_Y[1 | p, 1, 1, K] &= 1 - (1 - q)(1 - p + up)^{K-1} u \quad (3) \end{aligned}$$

2.3. Analysis of the single decoder

Thanks to (1), (2) and (3), we compute the amount of information $I(Y; X | p)$ that this test result brings about the identity of one infected person. This figure bounds the performance of the single decoder. Since the first step of our reconstruction algorithm (see next section) relies on this decoder, it is of utmost importance to maximize this figure of merit. Fig. 1 shows the typical behavior of this quantity as a function of p for the setups of [1](top) and [3](bottom). Define $p^*(K) \triangleq \arg \max_{p \in (0,1)} I(Y; X | p)$. We were not able to find a close form expression of its maximum but a good approximation is $I(Y; X | p^*(K)) \propto K^{-1}$. This is a strong difference w.r.t. traitor tracing where $\max_{p \in (0,1)} I(Y; X | p) \propto K^{-2}$ for the worst attack [7]. Since we do *not* a priori know the value of K at the generation of $\mathbf{M}$, we cannot set $p_i = p^*(K)$, $\forall i \in [T]$. However, given a range $K \in [\underline{K}, \overline{K}]$, we choose $f(p)$ as the uniform distribution over $[p^*(\underline{K}), p^*(\overline{K})]$.

3. THE RECONSTRUCTION ALGORITHM

Once vector $\mathbf{y}$ is observed, the reconstruction algorithm aims at identifying the K infected. In traitor tracing, the decoder of




Fig. 1. $I(Y; X|p)$ in nats as a function of p for different values of K . (top) Setup of [1]: $q = 0$ and $u = 0.2$; (bottom) Setup of [3]: $q = 0.01$ and $u = 0.05$.

a Tardos code computes a score for each codewords. It then accuses people whose score is above a threshold. Originally, Tardos proposed a pivotal quantity with respect to θ , in the sense that the distribution of the scores does not depend on this nuisance parameter [5]. In group testing, there is little ambiguity on θ because it only depends on K (we assume that (u, q) is known). Therefore, we prefer a ‘learn and match’ strategy which consists of estimating K first and then use an optimal detector. This has been tested for traitor tracing with some difficulties because identifiability issues appear: $\exists \theta_1 \neq \theta_2$ s.t. $\mathbb{P}_Y[1|p, \theta_1] = \mathbb{P}_Y[1|p, \theta_2]$, $\forall p \in (0, 1)$ [7, Prop. 4]. Here, it is not the case provided $u \neq 1$ and $q \neq 1$.

3.1. Estimation of K

The maximum likelihood estimator is given by:

$$\hat{K} = \arg \max_K \sum_{i=1}^T \log \mathbb{P}_Y[y_i|p_i, K].$$

There is no close form expression, but since p_i are distributed over a short interval, an approximation is $\hat{K}_0 = \log(\Delta/(1 + \Delta)/(1 - q))/\log(1 - \mathbb{E}(p)(1 - u))$, with $\Delta = \sum_{y_i=0} \log(1 - p_i(1 - u))/\sum_{y_i=1} \log(1 - p_i(1 - u))$. We then refine this by looking for a maximum $\hat{K}$ of the likelihood around $\hat{K}_0$.

3.2. Single decoder with likelihood ratio test

Once K has been estimated, the score of each individual j is indeed a test between the hypotheses:

- $\mathcal{H}_0$: This individual is not infected. The test results are independent of her blood sample: $\mathbb{P}_{Y, M_j} = \mathbb{P}_Y \mathbb{P}_{M_j}$.

- $\mathcal{H}_1$: This individual is infected. The test results depend on her blood sample: $\mathbb{P}_{Y, M_j} = \mathbb{P}_{Y|M_j} \mathbb{P}_{M_j}$.

The score of individual j is then the log-likelihood ratio computed thanks to (1), (2) and (3):

$$s_j = \sum_{i=1}^T \log \frac{\mathbb{P}_Y[y_i|p_i, M_{ij}, 1, \hat{K}]}{\mathbb{P}_Y[y_i|p_i, \hat{K}]}. \quad (4)$$

This decoder is more complex than the distance decoder proposed in [1], but a careful implementation takes one second to yield 10^6 scores on a regular computer. Yet, this decoder allows exact recovery only if the K infected get ranked first.

3.3. Advanced decoders

We propose two improvements of the single decoder. The first idea is to go for a joint decoder which computes scores for τ -tuples. This has never been done, as far as we know, due to the prohibitive complexity of browsing all the $\binom{N}{\tau}$ combinations. We will tend to this gradually. In the first step, we use the previous single decoder and isolate in the set $\mathcal{S}^{(2)}$ the $|\mathcal{S}^{(2)}| = \lceil \sqrt{2!N} \rceil$ persons having the highest scores. There are $\binom{|\mathcal{S}^{(2)}|}{2} = O(N)$ pairs in $\mathcal{S}^{(2)}$ so that the computation of their scores is ‘affordable’. After having sorted these pair scores in decreasing order, we include in $\mathcal{S}^{(3)}$ individuals involved in the pairs having the highest scores, and stop when the size $|\mathcal{S}^{(3)}|$ equals $\lceil \sqrt[3]{3!N} \rceil$. Denote by $r^{(3)}$ the number of pairs we had to browse to fill set $\mathcal{S}^{(3)}$. The idea is to gradually discard the less likely people while maintaining a list $\mathcal{S}^{(k)}$ of suspects short enough to allow the computation of $\binom{|\mathcal{S}^{(k)}|}{k} = O(N)$ scores over bigger k -tuples in the next step. For a given k -tuples $\mathcal{K}_\ell$, define $\rho_{i\ell} \triangleq \sum_{j \in \mathcal{K}_\ell} M_{ij}$, $\forall i \in [T]$ to compute the score as the following log-likelihood ratio:

$$s_\ell^{(k)} = \sum_{i=1}^T \log \frac{\mathbb{P}_Y[y_i|p_i, \rho_{i\ell}, k, \hat{K}]}{\mathbb{P}_Y[y_i|p_i, \hat{K}]}. \quad (5)$$

Exact recovery implies that no infected is discarded at any step, which is a weaker condition than for the single decoder.

The second idea consists in deeming as infected the most likely individuals. It is then possible to include this side information in the computation of the new scores. At the end of the $(k - 1)$ -th iteration of the joint decoder, $|\mathcal{S}^{(k)}|$ persons are suspected because they belong to the first $r^{(k)}$ $(k - 1)$ -tuples (Note that we fix $|\mathcal{S}^{(k)}|$ so that $r^{(k)}$ is indeed a random variable). While filling the set $\mathcal{S}^{(k)}$ we also count the number of times individuals in this set appear in the $r^{(k)}$ first $(k - 1)$ -tuples. Usually these tuples are hybrid, *i.e.* composed of infected and sound persons. But whereas we always find the same K infected, the identities of the sound persons are very different from a hybrid tuple to another. Therefore, infected people have a high number of appearances whereas sound people have low figures. In our algorithm, we accuse

only one individual $j^{(k)}$ if the number of appearances are sufficiently unevenly distributed. Denote $\mathcal{I}^{(k)}$ the set of deemed infected at iteration k . We have $\mathcal{I}^{(1)} = \mathcal{I}^{(2)} = \emptyset$ (no side information is available for the single and the pair decoder) and $\mathcal{I}^{(k+1)} = \mathcal{I}^{(k)} \cup \{j^{(k)}\}$ if someone is accused at round $k \geq 2$. Define $\sigma_i^{(k)} = \sum_{j \in \mathcal{I}^{(k)}} M_{ij}, \forall i \in [T]$. Then the score for k -tuple $\mathcal{K}_\ell$ with this side information is:

$$s_\ell^{(k)} = \sum_{i=1}^T \log \frac{\mathbb{P}_Y[y_i | p_i, \rho_{i\ell} + \sigma_i^{(k)}, k + |\mathcal{I}^{(k)}|, \hat{K}]}{\mathbb{P}_Y[y_i | p_i, \sigma_i^{(k)}, |\mathcal{I}^{(k)}|, \hat{K}]} \quad (6)$$

Exact recovery implies that only truly infected are included in the side information.

4. EXPERIMENTS

In this section we compare our proposed Tardos single and joint, side-informed decoders with two group testing setups reported by Cheraghchi *et al.* ([1], $N = 100000$, $K = 10$, $u = 0.2$ and $q = 0$) and Sejdinovic & Johnson ([3, 2], $N = 5000$, $K = 50$, $u = 0.05$, $q = 0.01$).

Fig. 2 shows the probabilities for exact recovery of all K infected individuals depending on the number of tests T for the Cheraghchi setup (*top*) and Sejdinovic setup (*bottom*). Here, the ‘Tardos’ contact matrix is constructed using the optimum $p^* = 0.0684$ and $p^* = 0.0138$, resp. At least 1000 trials have been performed at the various numbers of tests.

For the Cheraghchi setup, our decoders outperform the distance decoder [1], reducing the number of tests for exact recovery from 3000 to about 1200 tests using the single, and about 700 tests using the joint, side-informed Tardos decoder. In the Sejdinovic setup, the joint, side-informed Tardos decoder is competitive with the belief propagation reconstruction approach [3]; about the same number of tests are required for exact recovery.

We remark that both previous experimental setups assume K to be known which is unrealistic in practice. Experiments confirm that the performance of our decoders indeed degrades only slightly relying on the proposed estimate $\hat{K}$. For the sake of fair comparison, however, we stick to the scenarios reported in the literature.

5. CONCLUSION

This paper makes for the first time the connection between traitor tracing and group testing stressing their similarities and differences. Traitor tracing accusation processes are used for the identification of infected person and are shown to be competitive to previous reconstruction algorithm. Our future work aims at creating a measure of confidence about this reconstruction, associating each individual with an estimation of the probability that she is infected.




Fig. 2. Comparison of our single and joint, side-informed decoders with two group testing setups.

References

- [1] M. Cheraghchi, A. Hormati, A. Karbasi, and M. Vetterli, “Group testing with probabilistic tests: Theory, design and application,” in *Proc. 47th Allerton Conf. on Commun., Control and Computing*, Monticello, IL, USA, Sept. 2009, ext. version on arXiv:1009.3186v1.
- [2] G. Atia and V. Saligrama, “Boolean compressed sensing and noisy group testing,” *submitted to IEEE Trans. Inf. Theory*, 2009, arXiv:0907.1061v3.
- [3] D. Sejdinovic and O. Johnson, “Note on noisy group testing: asymptotic bounds and belief propagation reconstruction,” in *Proc. 48th Allerton Conf. on Commun., Control and Computing*, Monticello, IL, USA, Oct. 2010, arXiv:1010.2441v1.
- [4] P. Moulin, “Universal fingerprinting: capacity and random-coding exponents,” in *Proc. IEEE Int. Symposium on Inf. Theory*, Toronto, ON, Canada, July 2008, pp. 220–224.
- [5] G. Tardos, “Optimal probabilistic fingerprint codes,” in *Proc. 35th ACM Symposium on Theory of Computing*, San Diego, CA, USA, 2003, pp. 116–125.
- [6] T. Furon, L. Pérez-Freire, A. Guyader, and F. Cérrou, “Estimating the minimal length of Tardos code,” in *Proc. 11th Information Hiding Workshop*, Darmstadt, Germany, June 2009, vol. 5806 of *LNCs*, pp. 176–190.
- [7] T. Furon and L. Pérez-Freire, “Worst case attacks against binary probabilistic traitor tracing codes,” in *Proc. IEEE Int. Workshop on Information Forensics and Security*, London, UK, Dec. 2009, WIFS’09, pp. 46–50.