On the evidence for a statistical-computational gap raised by the Group Testing Problem

Dissertation
zur Erlangung des Doktorgrades
der Naturwissenschaften

vorgelegt beim Fachbereich 12, Informatik und Mathematik
der Johann Wolfgang Goethe-Universität
in Frankfurt am Main

von
Oliver Gebhard
aus Frankfurt am Main

Frankfurt 2022
(D30)

vom Fachbereich 12, Informatik und Mathematik, der

Johann Wolfgang Goethe-Universität als Dissertation angenommen.

Dekan:
Prof. Dr. Martin Möller

Gutachter:
Prof. Dr. Amin Coja-Oghlan
(TU Dortmund)

Prof. Dr. Jonathan Scarlett
(National University of Singapore)

Datum der Disputation:

ACKNOWLEDGE

This thesis was supported by the German Research Fund grant CO 646/3. First of all I am grateful to Amin Coja-Oghlan for his constant guidance within my university career. He introduced me to the research on the topics relevant to this thesis. Over the years he always supported me within my personal and academic development.
It was always a pleasure to spent time with my colleagues Philipp Loick, Joon Lee, Tobias Kapetanopoulos, Noela Müller and Maurice Rolvien. Many of the results relevant to this thesis would not have been possible without their steady support. A special thank goes to Max Hahn-Klimroth and Jean Ravelomanana. First of all, I enjoyed working with Max on various research projects, appreciated his supportive attitude and, of course, he was a great travel companion. Second of all, I am grateful that Jean let me stay at his flat during my time in Dortmund as well as that he was my steady pillar in the home office period during the COVID-19 crisis. I enjoyed our joint walks and discussions.

I thank Jonathan Scarlett for accepting to act as examiner of this thesis. Furthermore, I enjoyed working with him during our joint project and appreciated the comments that helped improving the content of various papers relevant to this thesis.

Of course, I am also grateful for the steady support of my fellow PhD students from other research groups. I benefited from our fruitful discussions. A special thank goes to Alexander Molitor, Joel Kübler and Florin Boenkost.
Although we have not spent too much time together within our PhD I want to express my gratitude to Stephan Gardoll who was there for me from day 1 of university and stayed until the very end. Without him the time at university would probably not have been as successful and as much fun as it indeed was.

Research is a team sport. Therefore, I also thank my various co-authors from places all over the world. I enjoyed working with Petra Berenbrink, Oliver Johnson, Dominik Kaaser, Olaf Parczyk, Malin Rau, Nelvin Tan, Alex Wein and Ilias Zadik. Within the review process of our various papers the comments of the anonymous reviewers helped to improve the readability of the papers. Their comments supported me within my scientific development. A special thank in this context goes to Uriel Feige for pointing us into the right direction within the development of one of our paper.

Finally, I would like to express my gratitude to my friends outside university. The wide spread of their scientific disciplines was beneficial in various ways. They supported and encouraged me whenever I needed it. In this context I especially want to thank Jonas Haller, Theresa Liebetanz, Adrian Schlagert and Paul Strauch for many helpful comments on the content of this thesis. Finally I thank my entire family. A special thank goes to my parents. Their support within my entire university career was the basic pillar of my success.

CONTENTS

1. Introduction

In today's world, algorithms and data structures are everywhere. On a daily basis we all use computers for navigation, communication or entertainment. Thereby, we trust in the security of our data and willingly accept that algorithms try to use as much information as possible to, for instance, predict our next purchase. As almost every part of our daily life is nowadays influenced by algorithms, two fundamental questions immediately arise:

(1) Why do algorithms succeed or fail?
(2) Where are the limits of algorithms?

Although everyone is familiar with using algorithms on a daily basis, formulating, understanding and analysing these questions rigorously has been (and will remain) a challenging task for decades. Due to the fact that the modern society is a complex and fast-evolving system, answering these questions for the real world is very difficult. Therefore, one way of making steps towards the answer is the formulation of models that are portraying reality, but also remain easy to analyse.
But even problems formulated in a simplified setting, may turn out to be difficult to solve. The following question, raised in the 19th century [102], is a prominent example for such a problem:

Example 1.1 (Travelling Salesman Problem)**.** *"Given a list of cities and the distances between each pair of cities, what is the shortest possible route that visits each city exactly once and returns to the origin city?"*

Obviously, this problem is portraying reality as all of us already tried to solve this task in our daily life by combining, for instance, the groceries with the post office, the bakery and the barber shop. At the same time it turns out that this question is harder to analyse than one might expect at first glance. Indeed, this problem raised a lot of attention over the years and researchers were not able to entirely understand it even though they tried to tackle it from different angles [10]. Of course, it is natural to ask how the success of a travelling salesman is connected to the questions we raised about the fundamental understanding of algorithms.
For a mathematician the answer is quite obvious. The 'Travelling Salesman Problem' is deeply connected to the, so called, *P-NP-Problem*, which formulates one of the most important questions in modern computer science and its scientific relevance is immediate[1]. We will discuss this part of the answer in detail in Section 1.2. For others it might be unclear why this particular problem raises attention in the first place. Therefore, we will use the Section 1.1 to introduce the relevance of this problem on an intuitive level. We will see how this basically sets a framework to talk about the limits of algorithms. A reader familiar with the *P-NP-Problem* may skip this section.

1.1. **Why computational hardness matters?**

The scientific importance of the *P-NP-Problem* is unquestioned (see for instance [57] for details), but also from a non-scientific point of view we can point out its relevance. To emphasise what the *P-NP-Problem* is and why the answer might be relevant even from a non-mathematical point of view; let us introduce Alice and Bob.
First, let us assume that Bob wants to send a private message to Alice. For Bob and Alice it should be easy to send and receive the message. Nonetheless, for a third party it should be hard to access their communication. But how do we prevent the third party from spying? We just have to find an encoding-decoding scheme such that with the right information it remains easy for Alice and Bob to encode and decode their

[1] In the year 2000 the Clay Mathematics Institute announced a list of the 7 most important open mathematical problems and exposed $1'000'000 for solving one of them. So far only one of the "Millennial Problems" has been solved.

messages, but a third party must invest a huge amount of computational power to reveal the encrypted information. This idea actually builds the backbone of the entire data security upon the internet and is also the point where finding computational hard problems comes into play. As aforementioned, one of the most important questions in modern computer science is the, so called, *P-NP-Problem*. This problem asks whether or not in every problem with an easy to verify solution (NP) it is also easy to find a solution (P). Thus, computer scientists around the world wonder whether $P \neq NP$ or $P = NP$. For the message between Alice and Bob it is great if the latter would not be true. Why is that the case?
Let us for the moment assume that $P = NP$. Then for every computational problem instance, for which we can verify a given solution easily, producing such a solution is easy as well. Today's encoding-decoding schemes are based on the fact that one can use the schemes with the right information easily (as just verification of given information is needed), but getting access without it is hard (as we have to produce the information ourselves). A prominent example is the RSA crypto-scheme [86], which is based on the believe that factorisation of large numbers is supposed to be hard. Thus, $P = NP$ would (in theory) imply that one could access the communication of Alice and Bob even without their authorisation. Therefore, from a cryptography point of view and for the security of our data $P = NP$ would not be the most favourable outcome.
As so far no proof exists that either confirms or refutes the $P \neq NP$-conjecture, scientists try to find evidence that hints in one or the other direction. Taking our introductory *Travelling Salesman* example, one can indeed show that it is (together with other combinatorial problems) NP-hard [63]. So far, no efficient algorithm is known to succeed on these problem sets. This, of course, is no proof in one or the other direction as just because we do not know an efficient algorithm yet, does not imply that there is none. Therefore, the question remains far from being answered. But how can we provide evidence in one or the other direction?
One common way to go is the analysis of toy models. Most algorithms take noisy data as an input and then use a given procedure to solve their task. Therefore, taking such a toy model to pinpoint the circumstances under which this instance is easy, hard or impossible to solve by certain types of algorithms, is a step towards understanding the fundamentals of algorithmic performance. On the one hand, we may wonder how much data per noise is necessary such that a problem instance provides sufficient information to solve a given task and at what point the contained information is lost within the noise. On the other hand, we are interested whether we can find an algorithm that is able to use the provided information efficiently (we will explain the underlying notion of 'efficiently' in Section 1.2.2). We call the point beyond which there is sufficient information contained in the problem set to solve it, the *information-theoretic-threshold*. Furthermore, we say that a task is *information-theoretically-possible* beyond this point. The point beyond which there exists an efficient algorithm that can use the provided information to succeed, is called *algorithmic-threshold*. We call a problem *easy to solve* beyond this point. We refer the reader to Figure 1 for an illustration. The basic question is whether these two thresholds match or if there is a 'fundamental' gap between them. This gap would lead to an 'information-theoretically possible, but computationally hard' phase and thereby imply $P \neq NP$.
Unfortunately, we are so far not able to proof such a fundamental gap and only unsettled evidence for the existence of such gaps exists. To make the interplay of signal and noise more accessible we use the following example. Everybody knows the predictions made by Amazon, Netflix or Facebook:

- "You bought product X, you might be interested in product Y as well."
- "You watched movie P, you might also like the first season of Q."
- "You are friends with M, you might know N, too."

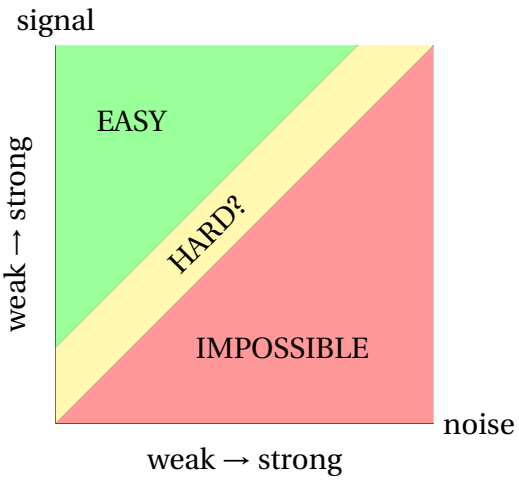

FIGURE 1. The diagram shows the interplay between signal strength (data quantity) and noise. It is often the case that for certain pairs there exists an algorithm that solves the task (green region) and for other pairs the noise is too strong such that no algorithm can solve the task (red region). In many cases, there is a gap where sufficient information is contained but no known algorithm succeeds (yellow region).

These predictions are made by algorithms that are designed to reveal our taste and behaviour from our data as well as the data of the users around us. Their predictions are based on the interplay of two measures, namely the quantity of the data (*signal*) and the *noise* within the data. If the signal is strong the algorithms will perform well even with a high level of noise and will start to fail if either the signal gets too weak or the noise too strong. In our example, one may think of the signal as all information we left on the internet that the algorithms may use. The noise comes in as, of course, not our entire life is spread on the internet and some parts of our personality and social interactions remain hidden from the algorithms. Therefore, the more information we reveal on the internet, the better the algorithms of Netflix, Amazon and Facebook can predict our taste. As those companies use their predictions for the purpose of earning money, they would like their algorithms to perform well even for low levels of signal and high levels of noise. All these prediction-algorithms are based on complicated models tailored for their explicit need and a mathematical rigorous analysis of their performance is (to the best of our knowledge) impossible. Nonetheless, we would like to understand why algorithms of this type succeed or fail.
Therefore, again scientists introduced simplified toy models to get a handle on the complexity and make steps towards understanding how these algorithms work and why they may start failing. Two well-known examples of these toy models for portraying the real-world in this scope are given by the *Principle Component Analysis* (PCA) and *Community Detection in the Stochastic Block Model* (SBM)[2]. We can, for the sake of the example, think of PCA as a simplification of the taste prediction rules for Amazon and Netflix. Additionally, the friend suggestions of Facebook are distantly related to the analysis of the SBM. For such toy models (even though very far away from the real world) researchers were able to pinpoint how signal and noise relate such that these problems are easy, (presumed to be) hard or impossible to solve. As not the focus of

[2]Note that the analogies are meant as visualisation of the need for toy models as step towards understanding algorithmic barriers in the real-world. The reader should be aware that there is a mayor difference between the methods used in real-world applications and the toy models used for research purposes. The results derived for these toy models do NOT directly apply to real-world networks.

this thesis, we refer the interested reader to [1, 60] for an overview about the results obtained for the PCA and SBM examples. Of course, one would like to analyse models and algorithms as close to the real world as possible. Due to the massive increase in complexity, this is so far only possible by introducing and analysing toy models distantly related to the real world. Over the last decades researchers from different fields shed light on the reasons why problems may be computational impossible, hard or easy. We will use Section 1.2 to introduce some of the concepts and will put our motivation into a rigorous perspective.

1.2. **Computational hardness.** The answer to the question whether or not a certain problems are hard to solve, is highly correlated with the scientific background of the person one asks. Therefore, we have to start with introducing the notion of hardness that we work on within this thesis. In fact, we are interested in computational hardness of problems. But how do we measure hardness and how do we classify a problem as hard or easy? Intuitively speaking for a given problem we measure, how long algorithms need to compute a proper solution to the problem set on the worst case input. The most popular classification is given by the notion of 'easy to solve' (P) vs. 'easy to verify' (NP). While algorithms are used as some kind of cure-all black box in colloquial language these days, we have to clarify how we actually formally think about an algorithm. From this definition we will be able to formulate and distinguish the two classes rigorously. We therefore shortly introduce the concept of the *Turing Machine*. One can imagine a Turing Machine as a tape that is divided into squares. This tape comes together with a set of symbols and a set of states. Our Turing Machine takes one of the states and each of the squares can either be empty or contain one of the available symbols. Now, our Turing Machine can move along the tape, read one symbol at a time, overwrite it according to the current state of the Turing Machine and change its current state. We refer the reader to [96] for more details. We say a problem belongs to the class P if there exists such a Turing Machine that can solve the problem by applying its procedure step by step, but for an input of size n it should only require a number of steps polynomial in n. At the same time, the class NP consists of all problems for which there exists a Turing Machine that can verify a given solution of the problem in a poly(n) steps. Obviously, $P \subset NP$ holds. But the question if $P \supset NP$ also holds remains open. In terms of Figure 1 we would like to know whether the two lines coincide or if the difference (yellow region) is of fundamental nature. In the following we will introduce some concepts and results that might carry evidence about the natural reasons of computational hardness. The analysis of computational hardness is closely related to understanding the behaviour of complex and chaotic systems. Another scientific discipline that tries to answer such questions is physics. We will use the following section to emphasise how physics relates to computational hardness.

1.2.1. *How understanding physics might help?* Starting with the fundamental work of Aristotle (384- 322 BC) and the seminal observations of Galileo (1561- 1642 AD), Isaac Newton (1643- 1727 AD) and Albert Einstein (1879- 1955 AD), for centuries physicists tried to understand complex real-world phenomena. Of course, their journey is far from being over, but one could get the idea that they might be able to support us on our journey of understanding computational hardness as we obviously entered the physicist's playground of analysing complex systems. Physicists usually try to explain their observations through theoretical models. The probably best known example of such a observation-theory pair is given by Newton's apple and the theory of gravity. Obviously, we do not find a direct explanation for all phenomena we observe around us. Therefore, similar to the theory of computational hardness, physicists started to introduce toy models to portray the real world in a simplified way. In particular, we will use

this section to emphasise the connection between our objective and the field of statistical physics and we do so, by focusing on the, so called, *spin-glass theory*. One can imagine the spin-glass theory as the physics equivalent of the toy model analysis done within the research of computational hardness.
Intuitively speaking, the theory of spin-glasses tries to explain, how macroscopic phenomena of particle systems are driven by the local interactions of the particles as well as the external influences affecting the system (like for instance temperature or air pressure). We start with introducing the common terminology and concepts used in this context and will use them to draw the line from physics to computational hardness. The general introduction is based on [21, 76, 83, 97]. A spin-glass is a set of particles that can take a magnetic orientation. For simplicity one may think of it as ± 1 and will call it spin from now on. Furthermore, the particles can interact with each other randomly. Therefore, such a particle system exhibit a certain form of disorder and it is often represented through a Hamiltonian H. Here, we think about the interactions through a random coupling J and the spin configuration of our N particles through a vector $\sigma = \{\sigma_1, ..., \sigma_N\}$. Then, the Hamiltonian is just a model-specific function

$$H : (J, \sigma) \mapsto H(J, \sigma).$$

To make the term 'model-specific' a little bit more accessible, we give the following example:

Example 1.2. *[The Sherrington-Kirkpatrick (SK) model] Consider a spin configuration $\sigma = \{\sigma_1, ..., \sigma_N\}$ on N particles with $\sigma_i \in \{\pm 1\}$ and J_{ij} as independent Gaussian variables. The Hamiltonian has the following form:*

$$H = -\frac{1}{\sqrt{N}} \sum_{i,j=1}^{N} J_{ij} \sigma_i \sigma_j$$

As the Sherrington-Kirkpatrick model (Example 1.2) is the probably most-studied spin-glass model, we will use it to access the ideas of spin-glass theory on an exemplary level. In contrast to purely ferromagnetic systems (adjacent particles tend to take the same spin, $J_{ij} = \mathbb{1}\{\sigma_i \neq \sigma_j\}$) and purely anti-ferromagnetic systems (adjacent particles tend to take opposite spins, $J_{ij} = \mathbb{1}\{\sigma_i = \sigma_j\}$), one can imagine a spin-glass as a mixture between ferromagnetic and anti-ferromagnetic (as in terms of Example 1.2 the J_{ij} are Gaussian variables and may take positive as well as negative values). Due to this mixture of interactions each particle receives contradicting preferences from its neighbours and the equilibrium state of the system is not as obvious as in the two pure models. Assume we start with an initial spin configuration σ. Now, together with the interactions induced by J we see that the system (σ, J) exhibits a certain amount of contradictions for different σ_i. One may think of these contradictions as internal energy of the system. Now, the system tries to reduce this internal energy by changing the spins over time. In the end, the system would like to reach one of its ground states σ^*, the states of least energy[3]. In the best case, this would end in a spin configuration that has no more contradictions. Due to the randomness of the interactions (and the thereby occurring contradicting preferences) it is usually not possible to reach such a state. The best we can hope for is a state with the least energy σ^*. Also reaching such a ground state might turn out to be a non-trivial task for the particle system. This is due to the existence of, so called, meta-stable configurations σ_M. These states exhibit the least energy within all configurations of a certain distance, but the internal energy of σ_M exceeds the energy of σ^*. Thus, locally they look like states of least energy, but they are not the optimal state upon the entire configuration space. Even though, these configurations σ_M are not the best possible outcome, leaving them is not easy for the

[3]For the moment we assume that the temperature T is tending to zero ($T \to 0$).

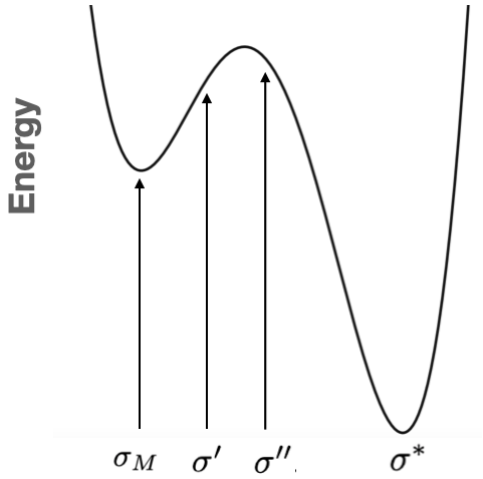

FIGURE 2. Consider four configurations $\sigma_M, \sigma', \sigma''$ and σ^*. Here σ_M denotes a meta-stable configuration, σ^* denotes the configuration of least energy (ground state) and σ' as well as σ'' denote intermediate configurations. Even though σ_M is not the best possible a change $\sigma_M \to \sigma'$ is no improvement as σ' remains in the attraction radius of σ_M. From an energetic point of view the system tends to return to σ_M as it is looking to achieve a lower energy. A change $\sigma_M \to \sigma''$ would help. In this case the system reaches a point that leads to an attraction to the state of lowest energy σ^*. This additional push to get to σ'' instead of σ' may turn out to be difficult or may take a very long time. Note that this illustration assumes the case temperature $T \to 0$.

system. At first glance a variation from σ_M would increase (instead of reduce) the energy and it would take multiple inferior steps until the system overcomes this energetic barrier. Therefore, the system might get stuck in these meta-stable configurations for a very long time as leaving them does not seem reasonable for the system from a energetic point of view. We refer the reader to Figure 2 for an illustration.
In this context, the main objective of physicists is the analysis of the Hamiltonian $H(\sigma, J)$ and its behaviour in the large system limit as $N \to \infty$. In particular, they are interested in $\min_\sigma H(\sigma, J)$. So far, the only influential parameters for the system were given by the spin configuration σ and the interactions induced by J. Thereby, we ignored (or simplified) external influences that might change the system's behaviour. Intuitive candidates for such parameters are the temperature and air pressure. A standard example one can have in mind at this point is water. As we all know, water appears in different physical forms (solid, liquid, gaseous) for different temperatures. Furthermore, air pressure influences, for instance, the boiling point of water[4]. A standard choice of these influences in the field of spin-glass theory is given by the temperature T and the external field h.
At this point, the physicists' dream is the analysis of their favourite spin-glass model, but apparently this turns out to be a non-trivial task, in general. Unfortunately, many properties of the spin-glass cannot be evaluated exactly. A notable exception in this regard is the Nishimori-line (a sub-space of the parameter space), where many properties can be evaluated exactly for certain spin-glass models [81, 82]. The questions one may ask are quite similar to the ones we already mentioned for computational hardness. On the one hand, we may ask under which circumstances the spin system's behaviour is actually coming from the internal particle interactions (spin-glass phase)

[4]On the Mount Everest we only need $71°C$ instead of the $100°C$ necessary on sea level. This is due to the difference in air pressure.

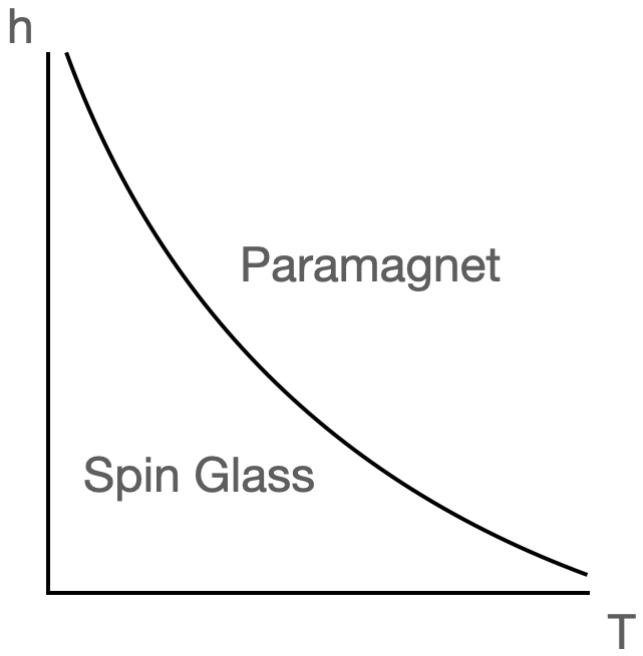

FIGURE 3. The behaviour of the Sherrington-Kirkpatrick model is depending on the temperature T and the external field h. Here, we see under which circumstances this model behaves like a spin-glass and under which the behaviour is driven by the external influences. The separation line is called Almeida-Thouless line.

and at what point the behaviour is just driven by the external influences T and h (paramagnetic phase). On the other hand, within the spin-glass phase finding and analysing the behaviour of the system's ground states and its physical properties is the next task. Thus, the physicist also tries to find the threshold beyond which there is some information about the spin configuration contained. Furthermore, she wants to find a way to analyse the contained information.

To make this connection even more accessible, we will emphasise it by considering the Sherrington-Kirkpatrick model (Example 1.2) as a toy model. After being introduced in 1975 by D. Sherrington and S. Kirkpatrick, this model evolved to one of the most popular and well-studied spin-glass models. We will highlight some of the results to provide an intuitive connection between spin-glasses and computational hardness research.

The first obvious connection is given by the separation of the parameter space given by the Almeida-Thouless (AT) line [36] (see Figure 3 for illustration). For the SK-model [23, 59, 100] rigorously established that this transition boundary is, indeed, mathematically meaningful. Above the line the external influences are too strong and interactions only play an insignificant role for the physical properties of the system. One may think of the, so called, para-magnetic phase as the impossible regime (red area in Figure 1). Thus, the AT-line can be seen as an information-theoretic threshold within the SK-model.

Now, moving below the AT-line we enter the, so called, spin-glass phase, where indeed interesting phenomena based on local interaction influence the system's physical properties. At this point, one may choose a property of the system and ask a) *Does the system contain sufficient evidence about the chosen property?* and b) *Can we evaluate this property efficiently?* These questions again sound quite familiar. Within the spin-glass phase of the SK-model some interesting properties were predicted as well as rigorously established.

Let us, for instance, assume that we may be interested in the expected spin of the i-th particle. A way to deal with this question was addressed in [99] by D. Thouless, P. Anderson and R. Palmer. Intuitively speaking, they proposed a system of equations (TAP-equations) to calculate the expected spin of the the i-th particle given the interactions J and the other $N-1$ particles. The next step is to obtain a solution of this system and, indeed, E. Bolthausen proposed an iterative construction for solving the TAP-equations [17] for certain regimes of T and h within the spin-glass phase.
With this in mind, we recall that, indeed, we might be able to benefit from the insights of physics. An intuitive connection between physics and computational hardness is given by the fact that one can imagine a spin configuration $\{\sigma_1,...,\sigma_N\}$ as binary input to the algorithm and the ground states σ^* as solutions obtained by an algorithm. Furthermore, one could imagine the existence of meta-stable states σ_M as a natural barrier that not only keeps a particle system from reaching its ground state but also causing algorithms to get stuck while solving a task. While the impossible phase corresponds to the para-magnetic phase, for the spin-glass phase one can ask whether we can say something meaningful about certain properties of our spin-glass. This corresponds to finding the 'information-theoretically possible' regime in the computational hardness research. Therefore, in our spin-glass example an immediate question is how the ground states σ^* relate to each other and whether meta-stable states occur. Thus, beyond the information-theoretic threshold as soon as sufficient information are contained in the problem set, the solution space geometry of our problem set might matter. It might hint to some 'natural' phenomena (based on the energy landscape of our particle system) that cause algorithmic barriers and computational hardness. We again employ the SK-model as example to access the theory on an explanatory level and transfer the predictions afterwards.
In the 1980s, G. Parisi[5] proposed a way to handle the analysis of the spin-glass phase of the SK-model. While the analysis is straight forward in the para-magnetic phase, the obtained solution is not satisfactory in the spin-glass phase. Therefore, G. Parisi proposed the idea of *Replica Symmetry Breaking*. Within the spin-glass phase he assumed the existence of many well-separated stable ground states[6]. The obtained predictions are in line with the theory of physics and indicate that, indeed, the configuration space geometry is getting more complex in the spin-glass phase. Furthermore, these predictions were rigorously established for the SK-model [11, 84].

At this point one might wonder why we should care about these predictions for the SK-model as the model itself seems quite distant from being relevant for the real world and, of course, the fact that the interactions only happen pairwise and that each particle influences all the others indicates that there might be models closer to the real world than the SK-model. Therefore, physicists started to apply their methods to models that avoid these connectivity issues. A prominent example in this regard is the Bethe-Lattice [75].

Due to the close relation of the questions raised in spin-glass theory and the admired goals in the scope of algorithmic hardness, physicists started to formulate and analyse many of the toy models known for computational hardness with their techniques.
The probably most influential transfer was proposed by connecting the separation properties for the ground states of spin-glasses and the existence of meta-stable states with the solution space of a broad class of well-known toy models [64, 73].

[5] In 2021 he received the Nobel price in physics for his work on complex systems.
[6] By well-separated we mean that they satisfy the following inequality: $\left|\sigma_1^*-\sigma_2^*\right| \leq \max\left(\left|\sigma_1^*-\sigma_3^*\right|,\left|\sigma_3^*-\sigma_2^*\right|\right)$. This is often referred to as ultrametricity.

We will use the following section to show how scientists tried to turn these predictions into rigorous evidence for computational hardness.

1.2.2. *On rigorous evidence for Computational hardness.* In the previous section, we took a brief detour in the physics' world. We have seen that it might not be the worst idea to listen to the physicists' intuition while working our way along the desired P-NP-Problem as their experience on complex systems might help. As already mentioned, within the research of computational hardness we normally consider toy models to pinpoint potentially hard regimes. We will use this section to motivate where the physics ideas can be applied to computational models and that one can actually turn these concepts into rigorous proofs. At this point, computer scientists and physicists are not as far apart as one might expect. Both want to understand complex systems and due to its complexity they portray the real world through toy models. For computer scientists a constraint satisfaction problem (CSP) is the equivalent to what a spin-glass means to the physicists. In the end, they also want to deal with their favourite CSP and reveal all relevant information about it. Consider n variables $x_1, ..., x_n$ (for simplicity one can think of x_i as binary) and m constraints $a_1, .., a_m$. Now, we connect the variables with constraints and equip each constraint with a function that takes the values of its neighbouring variables and maps it to either satisfied or unsatisfied. In this case, the connection process as well as the function given to the constraints are problem specific. For the moment, Figure 4 provides a simplified example one can have in mind while thinking about such problems. One may ask whether or not there exists a configuration $(x_1, ..., x_n)$ that meets all constraints. Furthermore, one would like to know how to find such a configuration. Obviously, the answer of such questions are depending on the ratio $\alpha = \frac{m}{n}$ between the number variables and the number of constraints. It seems, of course, tempting to take the physicists tool box and throw it

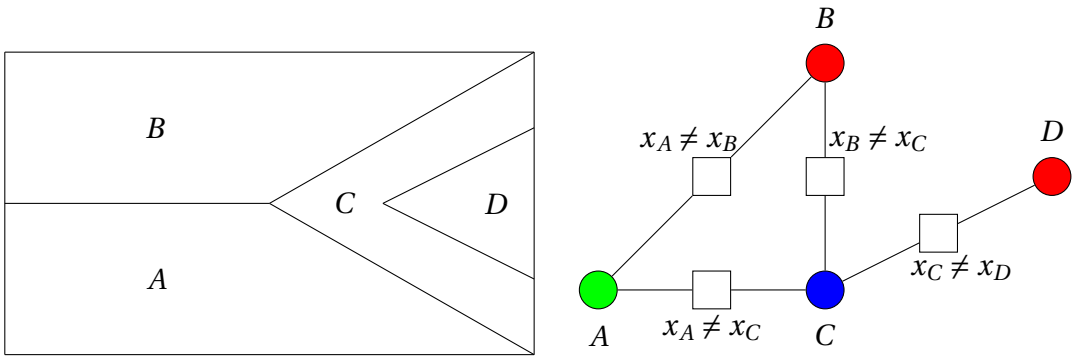

FIGURE 4. Assume we would like to colour our favourite map, we have 3 colours available and we would like to do it in a way such that countries with a joint border are coloured differently. Now we can encode the information provided by the map (countries as circles, common borders by lines). For each border we add the corresponding constraint ensuring that countries corresponding to the particular border are coloured differently (rectangle). We check that indeed the right-hand-side shows a proper colouring of our map. Note that (even beside permutation) the colouring is not unique as we could recolour area D and use green instead. For completeness, we draw attention to the fact that we would not be able to colour the given map with only two colours (due to region C).

at your favourite CSP, but all that glitters is not gold. On the one hand, their predictions are often correct and provide a good starting point for rigorous research. On the other hand, we have to be careful as their ideas are often non-rigorous and hard to

prove. The intuitive connection between CSPs and spin-glass systems is quite obvious. Of course, instead of variables $x_1, ..., x_n$ and constraints $a_1, ..., a_m$ in a CSP we can think about a system of n interacting particles. Furthermore, finding a proper solution corresponds to a configuration of $x_1, .., x_n$ such that there are no contradictions, thus automatically leading to a state of least energy σ^*. Over the years physicists used their methods to handle all kinds of different CSPs. While some are dealing with particular problem sets like our introductory example about the travelling salesman [70] or colouring certain underlying structures [45], the probably most influential contributions [64, 73] analyse a more general framework. First of all, they predict the ratio $\alpha = \frac{m}{n}$ below which there exist a solution for a broad class of CSPs. Secondly, they use the separation property of the ground states as well as the existence of meta-stable states to predict that beyond a certain α' it will be hard to find a solution even though there exists one. Since then researchers tried to pour these predictions into a solid and rigorously proven framework. Over the years a rich body of literature emerged and proved many of the physicists' predictions about the solution space of a broad class of CSPs rigorously [2, 3, 4, 28, 33, 34]. Interestingly, many problems exhibit the predicted phenomena of a mismatch between the existence of solutions and the ability of finding them efficiently with the state-of-the-art algorithms. Famous member of this group are finding a *large independent set* [27], or *planted clique* [9] in Erdös-Renyi graphs and the *random k-SAT problem* [25]. Even though all of the three examples are worth a closer look, we will employ the random k-SAT problem as explanatory example.

Example 1.3 (Random k-SAT)**.** *Given n variables $x_1, ..., x_n$ and m constraints $a_1, ..., a_m$. Each $a_j = (a_{1j}, ..., a_{kj})$ now chooses its k variables uniformly at random from $\{x_1, ..., x_n, \neg x_1, ..., \neg x_n\}$.*

(1) For each a_j we connect the $a_{1j}, ..., a_{kj}$ with a logical OR ($\vee$).

(2) We connect all $a_1, ..., a_m$ with a logical AND ($\wedge$).

This leaves us with a problem instance of the following form:

$$\Phi = (a_{11} \vee a_{21} \vee ... \vee a_{k1}) \wedge (a_{12} \vee a_{22} \vee ... \vee a_{k2}) \wedge ... \wedge (a_{1m} \vee a_{2m} \vee ... \vee a_{km})$$

The target is to find a configuration $(x_1, ...x_n)$ with $x_i \in \{$ TRUE, FALSE $\}$ such that Φ evaluates to TRUE.

We immediately see that the solubility of this problem depends on the interplay of n and m. Of course, *random k-SAT* is a prime example for a random CSP. Therefore, one can use the non-rigorous predictions for the transitions in the solutions space geometry [64, 73] as a starting point for rigorous research. First of all [73] predicted that as soon as $\alpha = \frac{m}{n}$ passes $\alpha_{\text{SAT}} := 2^k \ln(2) - \frac{1+\ln(2)}{2} + o_k(1)$ the problem has a satisfying configuration for any $\alpha < \alpha_{\text{SAT}}$ and has none for $\alpha > \alpha_{\text{SAT}}$. Although this prediction is based on non-rigorous physics methods, this transition was rigorously justified by [35, 37]. The next step on the way of understanding *random k-SAT* is the prediction of [64, 73] referring to the computational hardness of finding a satisfying assignment. Due to a change in the solution space geometry, they predict that in the intermediate regime $\alpha_{\text{ALG}} < \alpha < \alpha_{\text{SAT}}$ with $\alpha_{\text{ALG}} = 2^k \frac{\ln(k)}{k}$ it is computationally hard to find a satisfying configuration although there exists one. In the special case of random k-SAT physicists predict [77] that this is due the emergence of meta-stable states as soon as one surpasses α_{ALG}. Furthermore, they predict that algorithms, that are not able to observe the entire solution space, might get stuck in these meta-stable states as from the algorithm's perspective the local minimum looks like the optimal solution even though they obviously are not optimal. If this holds for all polynomial-time algorithms, we would immediately get $P \neq NP$. This would be the case because we could take an assignment and verify whether (or not) it is satisfying in polynomial time. But at the same time computing a satisfying assignment in polynomial-time would be hard. Of course, we are not able to prove such a strong result for random k-SAT. There are two

ways to go. First of all, one can propose algorithms that perform well for certain variable/constraint densities α. In a second step, one can check whether certain classes of algorithms are (not) able to surpass the barriers set by the existing algorithms. The state-of-the-art algorithm [25] works all the way up to the predicted 'optimal' barrier α_{ALG}. Since then different types of algorithms tried to surpass that barrier and, so far, no one was able to do so [19, 26, 32, 56]. There are two points we would like to emphasise here upon the strength of the evidence in the case of random k-SAT. First of all, physicists usually believe that certain message-passing types of algorithms are giving the optimal answer to algorithmic problems. While this was often the case for other problems, this seems not to be the case for random k-SAT [26, 32, 56]. Secondly, very recent research relates the separation phenomena with the existence of a, so called, *Overlap Gap Property* (OGP)[7] and show that the performance of a huge class of algorithms (low-degree polynomials) is strongly related to this phenomena. In the case of random k-SAT, [19] show that this class of algorithms also fails to surpass the barrier α_{ALG} set by [25]. Due to recent results on OGP as well as low degree methods [13, 15, 18, 19, 38, 48, 47, 93] it is sometimes conjectured that the class of low-degree-polynomials can be used as proxy for the class of efficient algorithms as it covers most of the a state-of-the-art algorithms. In the case of random k-SAT, this strengthens the believe that the result of A. Coja-Oghlan [25] is indeed optimal. Furthermore, it also strengthens the evidence that there may exist barriers that algorithms cannot surpass efficiently. One can see this remaining gap between α_{ALG} and α_{SAT} as a step towards the P-NP question. The intruding question at this point is the following:

$\Rightarrow$ Are we able to find other toy models and problem instances that strengthen this evidence?

The remainder of this thesis will make steps towards answering this question by assuming several variants of the, so called, *Group Testing Problem*. We will use the next section to motivate the problem set and introduce it on an intuitive level.

1.3. **Group Testing.** The *Group Testing Problem* was introduced in 1943 by R. Dorfman [41]. Due to the high demand at that time the testing capacity for syphilis was limited. Therefore, he raised the following idea to increase the testing capacity:

The Group Testing Problem:
Assume we have a large group of individuals and a very small subset of them suffer from a rare decease. Instead of using one test for one person, is it possible to determine the infection status of each person by conducting pooled tests?

Before turning to the mathematical questions one can tackle within this problem the reader might wonder about an even more fundamental issue here as it seems non-trivial to move from individual to pooled tests on a chemical level. While this procedure is not possible for every virus, there exist notable examples where such procedures are not only possible but even in daily use [72]. The most prominent examples are probably given by COVID-19 [79] and HIV [105]. Furthermore, Group Testing procedures found their way into some real-world applications like DNA sequencing [67, 80] and protein interaction experiments [78, 98]. Furthermore, it appears as an essential tool to face pandemic spread [24].
In his original work R.Dorfman proposed a testing procedure in two rounds. In a first step the individuals are split into groups. One may think of it as collecting the blood or saliva of each participant of a group and mixing it in one pot. If a test returns negative the Dorfman-procedure declares all participants as uninfected. For the participants in

[7] We refer the reader to the recent survey articles [46, 66] for a detailed discussion

positive tests a second round of individual tests follows. The crucial idea of this procedure is two-fold. First of all, a negative test does not contain the virus therefore none of the participants donated it and thereby all are uninfected. While in a positive test at least one participant is responsible for the contamination, we do not know who it was. Therefore, the second round clarifies that.
At that point one might wonder why we so far ignored faulty tests and dilution effects and, of course, this is correct. We cannot always assume the gold standard and we cannot pool an unlimited number of individuals in one test. For instance COVID-19 test are to some extend faulty [106] and due to dilution effects we can only pool a finite number of individuals [79]. While we will address these problems within this thesis (Section 3 and Section 4) by adjusting the model, we start with the general questions one may ask. Beyond n individuals we have k infected individuals and we are allowed to conduct m tests. In this thesis we apply the standard assumption $k \sim n^\theta$ with $\theta \in (0,1)$. While this sub-linear scaling might seem somewhat artificial, it is indeed a natural scaling in early stages of epidemics, due to Heaps' Law [16, 103]. Here again, one can ask similar questions as before. We are interested in the minimum number of tests m such that the test carry some information about the infected subset. Furthermore, we would like to know how many tests we actually need to work with this information efficiently.
Obviously, this problem is another example of a CSP as we can see the test $a_1,...,a_m$ as constraints, our individuals $x_1,...,x_n$ as the variables and the infection status as binary information $\{0,1\}$. The constraints are induced by the individuals participation in a test. We have to find a configuration such that a negative test does not contain a variable $x_i = 1$, thus, does not contain an infected individual. Furthermore, we have to ensure that the configuration does have at least one $x_i = 1$ in each positive test, thus does contain at least one infected individual.
One might notice at this point that this task is a CSP of a different nature. Indeed, it is called *planted CSP*. Until now (for instance random k-SAT) we were interested in finding some solution for the CSP, but we had no preference about which one to choose. In the Group Testing problem this is not sufficient anymore. Here, we would like to recover, approximate or at least detect the set of infected individuals responsible for the test result. A picture one can have in mind is the, so called, *Teacher-Student-Model* (see [107] for an detailed overview). Here, a teacher receives a ground truth $x_1,...,x_n$ and uses $x_1,...,x_n$ together with a model G to generate observable data $a_1,...a_m$. Now, she passes G as well as $a_1,...,a_m$ over to a student. The student now tries to recover $x_1,...,x_n$ from the given information. As this seems quite technical on first glance, everyone of us probably experienced this situation in high school. A teacher reads a fact in a book, writes it on the blackboard and we tried to recover the written facts from the blackboard. Now, we would like to analyse the chances of the students to learn the right facts. We refer the reader to Figure 5 for illustration of a Group Testing instance. In the scope of this thesis, we analyse the chances of the student to recover, approximate and detect the infected subset. We will shed light on the question how model specific constraints influence her chances. The following sections are dedicated to rigorously introduce the models, results and techniques used to obtain the results relevant to this thesis. We will also introduce the state-of-the art before the papers relevant to this thesis [30, 31, 50, 51, 52] have been published/submitted and place the results into perspective.

1.3.1. *Fundamentals, combinatorics and notation.* While the COVID-19 pandemic revealed the importance of the Group Testing problem in its own way, the mathematical analysis of this problem will be the main focus of the thesis at hand (also see [7] for a recent survey). On an explanatory level, we have already seen in Figure 5 that we can

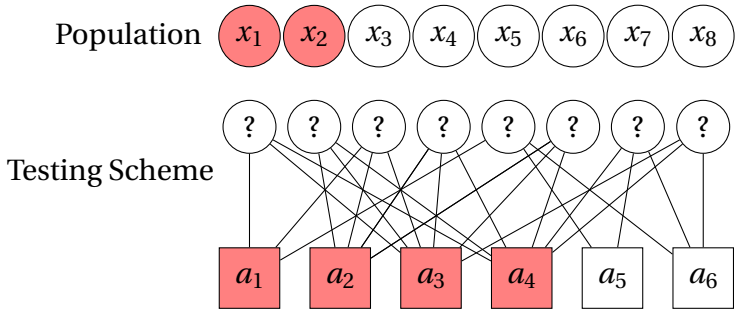

FIGURE 5. A small example of a the Group Testing problem. The top line represents the infection status one would like to find. Circles represent individuals (red infected, white uninfected). Rectangles represent tests (red positive, white negative). The edges between rectangles and circles show which individual participates in which test. The target is to tell which of the question marks are red and white.

translate the problem into a combinatorial framework. Therefore, we introduce the concept of the, so called, factor graph.

Definition 1.4 (Factor Graph, see for instance [74])**.** *Let $\Omega \neq \emptyset$ be a finite set. A Ω-factor graph $G = (V, F, (\partial_a)_{a\in F}, (\psi)_{a\in F})$ consists of*

- *a set V of variable nodes,*
- *a set F of constraint nodes,*
- *a neighborhood $\partial_a \in V$ for each $a \in F$,*
- *and a weight function $\psi_a : \Omega^{|\partial_a|} \to \mathbb{R}$.*

Remark 1.5. *In our Group Testing problem we encode the population's infection status by $\sigma = (\sigma_1, ..., \sigma_n)$ with individual infection status $\sigma_i \in \Omega := \{0,1\}$. The set V contains the individuals and the set F consists of the tests. Our pooling scheme induces ∂_a for all $a \in F$ as we know which individual participates in which test. Note, that in the Group Testing problem the planted assignment $\boldsymbol{\sigma} = (\boldsymbol{\sigma}_1, ..., \boldsymbol{\sigma}_n)$ induced which test is positive and negative. Finally $\psi_a := \psi_a^{(\boldsymbol{\sigma})}$ ensures that a test is positive if at least one infected individual is contained:*

$$\psi_a^{(\boldsymbol{\sigma})} : \{0,1\}^{\partial_a} \to \begin{cases} \mathbb{1}_{\{\sum_{x_i \partial a} \sigma_i > 0\}} & \text{if} \quad \sum_{x_i \partial a} \boldsymbol{\sigma}_i > 0 \\ \mathbb{1}_{\{\sum_{x_i \partial a} \sigma_i = 0\}} & \text{if} \quad \sum_{x_i \partial a} \boldsymbol{\sigma}_i = 0 \end{cases}$$

With this in mind, we can state the basic underlying structures that are responsible for the complexity of certain inference tasks. From now on $\boldsymbol{\sigma}$ denotes a uniform chosen vector $\{0,1\}^n$ with Hamming weight k, encoding the populations underlying infection status. Furthermore, $\hat{\boldsymbol{\sigma}} = \hat{\boldsymbol{\sigma}}(\mathcal{G}, \boldsymbol{\sigma}) \in \{0,1\}^m$ denotes the sequence of (pre-noise) test results, such that $\hat{\boldsymbol{\sigma}}_a = 1$ iff test a contains at least one infected individual, that is

$$\hat{\boldsymbol{\sigma}}_a = \max_{x \in \partial a} \boldsymbol{\sigma}_x.$$

We now introduce the underlying important structures that are important for the noiseless Group Testing. While these do not immediately transfer to the noisy case (but are using adoptions of these ideas) we will introduce the required adoptions in the corresponding Section 4. First of all, we divide our population into two groups (infected, uninfected): Given a pooling scheme $\mathcal{G}$, let

$$V_0(\mathcal{G}) = \{x \in V(\mathcal{G}) : \boldsymbol{\sigma}_x = 0\}$$
$$\text{and} \qquad V_1(\mathcal{G}) = \{x \in V(\mathcal{G}) : \boldsymbol{\sigma}_x = 1\}$$

We find uninfected individuals that are easy to identify. A negative test is a clear indicator for an individual to be uninfected, we call the set of these individuals V_{0-}. Formally,

$$V_{0-}(\mathcal{G}) = \{x \in V_0(\mathcal{G}) : \exists a \in \partial_{\mathcal{G}} x : \hat{\boldsymbol{\sigma}}_a = 0\}. \tag{1.1}$$

Furthermore, there are infected individuals that are easy to identify. If an individual x participates in a positive test a and x remains the only individual in test a after removing all easy uninfected individuals, the remaining x must be the one responsible for the positive test result. Thus, it is easy to identify. Formally, we define them as

$$V_{1--}(\mathcal{G}) = \{x \in V_1(\mathcal{G}) : \exists a \in \partial_{\mathcal{G}} x : (\partial_{\mathcal{G}} a \setminus \{x\}) \subset V_{0-}(\mathcal{G})\}. \tag{1.2}$$

This leaves us with individuals that are less easy to handle. We call an uninfected individual x disguised if it only appears in positive tests. Furthermore, we call an infected individual x disguised if it appears only in positive tests where at least one other infected individual beside x is contained. Formally we define them as,

$$V_{0+}(\mathcal{G}) = \{x \in V_0(\mathcal{G}) : \forall a \in \partial_{\mathcal{G}} x : \hat{\boldsymbol{\sigma}}_a = 1\}. \tag{1.3}$$

$$V_{1+}(\mathcal{G}) = \{x \in V_1(\mathcal{G}) : \forall a \in \partial_{\mathcal{G}} x : (\partial_{\mathcal{G}} a \setminus \{x\}) \cap V_1(\mathcal{G}) \neq \emptyset\}. \tag{1.4}$$

An illustration can be found in Figure 6. The general idea is to analyse the existence of

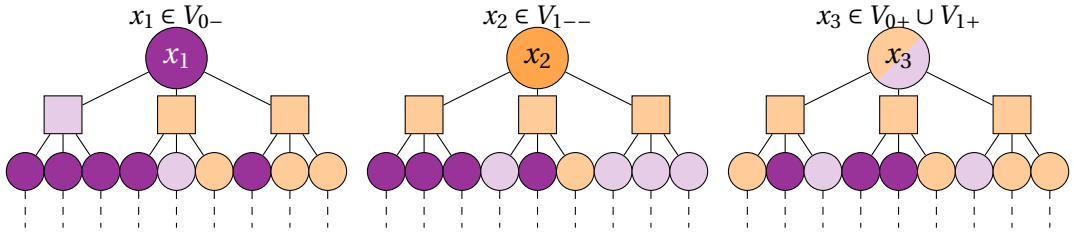

FIGURE 6. Rectangles represent tests and circles individuals. Dark violet individuals are elements of V_{0-} and can be easily identified as uninfected. Light violet individuals are elements of V_{0+}, and even if uninfected themselves, they only appear in positive tests and might be hard to identify. Infected individuals that appear only in tests containing another infected individual are impossible to identify (light orange). Finally, infected individuals of V_{1--} (dark orange) appear in at least one test with only elements of V_{0-} . Thus, after identifying all elements of V_{0-}, they can be identified. The dashed lines represent the fact that the individuals may also participate in other tests; these may include negative tests classifying their participants as uninfected (elements of V_{0-}) even though the particular test displayed is positive. This figure is adopted from [51].

these structures to decide whether the given task on a Group Testing instance is easy, hard or impossible to solve.

Remark 1.6. *So far the individual-types assumed noiseless test outcomes. We will handle noisy tests as well within this thesis (Section 4). Note, that as soon as we change the model to allowing noisy tests a single test is not a clear indicator anymore. While the key idea still is that the influence of the infected and uninfected individuals differ by construction of the Group Testing problem (even with noisy tests), we will have to introduce an appropriate notion of this influence to get started with the noisy case. We will go into detail in Section 4.*

For a pooling scheme $\mathcal{G}$, a underlying population infection status vector $\boldsymbol{\sigma}$ and a test results sequence $\hat{\boldsymbol{\sigma}}$, we define $S_k(\mathcal{G}, \boldsymbol{\sigma})$ as the set of all population infection statuses $\tau \in \{0,1\}^n$ that satisfy the test outcomes $\hat{\boldsymbol{\sigma}}$ (of course, including $\boldsymbol{\sigma}$ itself). Furthermore, we set $Z_k(\mathcal{G}, \boldsymbol{\sigma}) = |S_k(\mathcal{G}, \boldsymbol{\sigma})|$. By Corollary 2.1 of [29] we know that, given the test result and a uniformly sampled planted assignment, all sets in $S_k(\mathcal{G}, \boldsymbol{\sigma})$ are equally likely.

Thus, as soon as multiple satisfying assignments exist one cannot do better than sample one uniformly at random. We will use these underlying structures to analyse the algorithmic and information-theoretic limits of Group Testing. Note that we employ standard Landau-notation $\Theta(*), O(*), o(*)$ and $\omega(*)$ from now on within this thesis.

1.3.2. *On the model variations and the different notions of success.* We consider n individuals and $k \sim n^\theta$ with $\theta \in (0,1)$ of them are infected. We assume $\boldsymbol{\sigma} \in \{0,1\}^n$ to be drawn uniformly at random among all vectors of size n with Hamming weight k. Furthermore, we assume that we conduct m tests. The pooling of the n individuals within the m tests is given by $\boldsymbol{G}(n,m,\theta)$. Let $\hat{\boldsymbol{\sigma}} \in \{0,1\}^m$ denote the test result vector.

Definition 1.7 (The Models for the Group-Testing Problem)**.** *We assume the following variations of $\boldsymbol{G} := \boldsymbol{G}(n,m,\theta)$ within this thesis:*

(1) The standard-model assumes noiseless tests, binary test output and neither tests nor individuals carry any size-constraints. This model is the basis for [30, 31] *and will be addressed in Section 2.*

(2) The noisy-model keeps the binary output and the non-existent size-constraints, but each test carries an error probability. This model is relevant in [52] *and will be addressed in Section 4.*

(3) The size-constrained model keeps the binary output as well as the noiseless tests. The relevant modifications are given by the size-constraints:
   a) *Tests can only contain a limited number of individuals.*
   b) *Individuals can only participate in a limited number of tests.*

   In [51] *we assume this model and we will address the details in Section 3.*

(4) The quantitative-model keeps the noiseless tests and the non-existing size-constraints, but instead of telling us whether or not an infected individual is contained, a test returns the number of infected individuals within the test. This model is relevant in [50] *and will be addressed in Section 5.*

We consider these variants of the model and will introduce them formally in the corresponding sections. From Definition 1.7 we get the problem instance we work on. There are three algorithmic and information-theoretic objectives we handle within this thesis.

Definition 1.8 (Exact Recovery)**.** *An algorithm $\mathscr{A}$ solves* exact recovery *of $\boldsymbol{\sigma}$ for a Group Testing instance, if $\mathscr{A}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}, k)$ correctly identifies the infected individuals with high probability.*

This notion of success is relevant in [30, 50, 51, 52]. The following two claims give us a starting point within the analysis of the limits of exact recovery.

Claim 1.9 (Claim 2.3 of [51])**.** *For any test design, we have $Z_k(\mathscr{G}, \boldsymbol{\sigma}) \geq |V_{1+}(\mathscr{G})||V_{0+}(\mathscr{G})|$. Hence, conditioned on the sets $V_{1+}(\mathscr{G})$ and $V_{0+}(\mathscr{G})$, any inference algorithm fails with probability at least $1 - \frac{1}{|V_{1+}(\mathscr{G})||V_{0+}(\mathscr{G})|}$.*

In other words, exact recovery becomes impossible if we find 'many' disguised infected as well as disguised uninfected individuals. In this case, we could flip the status of some of these individuals and find a second satisfying assignment without being able to distinguish them. For details we refer the reader to [51].

Claim 1.10 (Claim 2.4 of [51])**.** *Exact recovery is easy if $V_1(\mathscr{G}) = V_{1--}(\mathscr{G})$.*

Thus, there exists an algorithm that can solve exact recovery as soon as all infected individuals exhibit the right neighbourhood structure (defined in (1.2)). The, so called, `DD` can use these structures to succeed. Intuitively speaking, it declares all participants

of negative tests as uninfected and checks the remaining individuals for $x \in V_{1--}(\mathcal{G})$. Thus, if one finds all infected individual this way, it succeeds. We will employ the DD to establish the easy regime in multiple models. Therefore, we will come back to this algorithm in Section 2 and Section 3. Again, the details can be found in [51]. A common way to ease the exact recovery criteria is given by weak-recovery as one might be satisfied with a 'good portion' of correctly identified individuals instead of getting the entire set correct.

Definition 1.11 (weak-recovery)**.** *An algorithm $\mathcal{A}$ recovers the ground-truth $\boldsymbol{\sigma}$ of a Group Testing instance δ-weakly, if $\mathcal{A}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}, k)$ outputs an estimate $\tilde{\boldsymbol{\sigma}}$ with Hamming weight k s.t. $\langle \tilde{\boldsymbol{\sigma}}, \boldsymbol{\sigma} \rangle \geq \delta k$.*

We deal with this criteria in [30, 31] and the following claim will be helpful to establish results on weak-recovery.

Claim 1.12 (Section 5.1 of [31])**.** *Fix any constant $\delta > 0$ and let $\boldsymbol{\tau} \in \{0,1\}^n$ be uniformly sampled from $S_k(\mathcal{G}, \boldsymbol{\sigma})$. Then weak-recovery is impossible if*

$$\mathbb{P}\left(\langle \boldsymbol{\sigma}, \boldsymbol{\tau} \rangle \geq \delta k\right) = o(1).$$

In other words, we have to ensure that the probability of two satisfying assignments (with k infected individuals) overlapping too much, is small. Thus, instead of analysing the number of satisfying assignment we have to analyse the relation between different satisfying assignments of the underlying problem set.

An even weaker criteria is called 'detection'. Here, we are not interested in finding the infected subset, but telling whether the underlying pooling scheme is actually coming from a Group Testing instance. Thus, we would like to find a way to distinguish a graph coming from a Group Testing instance from an appropriately chosen random graph model.

Definition 1.13 (Strong Detection [31])**.** *An algorithm $\mathcal{A}$ is said to achieve* strong detection *if, given input $(\boldsymbol{G}, k)$ with $\boldsymbol{G}$ drawn from either a Group Testing instance $\mathbb{P}$ or a random noise model $\mathbb{Q}$ (each chosen with probability $1/2$), it correctly identifies the distribution ($\mathbb{Q}$ or $\mathbb{P}$) with probability $1 - o(1)$.*

The following claim helps us while analysing this success criteria.

Claim 1.14 (Section 7 of [31])**.** *Suppose $\mathbb{P} = \mathbb{P}_n$ and $\mathbb{Q} = \mathbb{Q}_n$ are distributions over $\{0,1\}^p$ for some $p = p_n$. We say that an estimator f strongly separates $\mathbb{P}$ and $\mathbb{Q}$ if*

$$\sqrt{\max\left\{\mathrm{Var}_{\mathbb{P}}[f], \mathrm{Var}_{\mathbb{Q}}[f]\right\}} = o\left(\left|\mathbb{E}_{\mathbb{P}}[f] - \mathbb{E}_{\mathbb{Q}}[f]\right|\right).$$

If there exists a estimator $f = f_n$ that strongly separates $\mathbb{P}$ and $\mathbb{Q}$ then strong detection is possible.

Therefore, as long as we find an estimator that is expected to output two different estimates for both models and they are sufficiently well separated such that it is not too likely that one looks like the other by chance, we can tell the two models apart. In [31] we also deal with this notion of success.

In the remainder of this thesis we will see how these different criteria are addressed for the different models mentioned in Definition 1.7. Section 2 deals with the unconstrained case. We deal with the size-constraints in Section 3. The noisy variant is addressed in Section 4. Finally we address the quantitative output in Section 5.

2. Unconstrained Group Testing

In this section, we handle the standard Group Testing model. We assume our population of size n with an infected subset of size $k \sim n^\theta$. Furthermore, we do not place any size-constraints on neither the tests nor on the individuals. We assume that the tests are error-free. In other words, they return positive if and only if an infected individual is contained. This section is based on [30, 31].

2.1. **Related work.**

As discussed in the previous section, the original idea was proposed by R. Dorfman in 1943. In his work [41] he proposed a first pooling procedure to solve the given task and, of course, the obvious question is whether or not one can improve upon the proposed procedure. Since then, scientists tried to determine the minimum number of tests necessary to reveal information about the infected subset. A trivial counting bound states that exact recovery fails as soon as $m < (1-\varepsilon)m_{\text{count}}$ with $m_{\text{count}} := \frac{k\ln(n/k)}{\ln(2)}$. This follows as any algorithm fails as soon as we find two population infection statuses $\{0,1\}^n$ leading to the same test result. As the output of our m tests is binary and we have to place k infected individuals in our population of size n, the lower bound immediately follows from ensuring

$$2^m \geq \binom{n}{k}. \tag{2.1}$$

With this in mind, one can wonder whether one is able to recover the infected subset by conducting only m_{count} tests. It is known by [8, 12, 58] that this is indeed possible. Thus, exact recovery is easy all the way down to the threshold set by m_{count}. The main weakness of these results comes from the required multi-stage testing (usually referred to as adaptive strategy). Of course, these results were major breakthroughs, but it is not fully satisfactory. This immediately follows from the origin of the Group Testing problem itself. Chemical restrictions [53] and the possible large-scale effects of pandemic spread [24, 71, 79] actually require fast Group Testing schemes for useful practical application. On top of that, one would like to simplify the schemes as much as possible. Therefore, over years scientists tried to reduce the number of stages as much as possible. Thus, the ultimate target is an 1-stage procedure (usually called non-adaptive) that recovers the infected subset with m_{count} tests. Indeed, [87] reduced the number of stages needed to 3. But can we get all the way down to m_{count} with a 2-stage or even an 1-stage procedure for all $\theta \in (0,1)$?
Indeed, [29] showed that exact recovery is information-theoretically possible with an 1-stage procedure by conducting

$$m_{\text{inf}} = m_{\text{inf}}(n,\theta) = \max\left\{\frac{\theta}{\ln^2(2)(1-\theta)}, \frac{1}{\ln(2)}\right\} n^\theta \ln(n/k). \tag{2.2}$$

Note, that this bound matches m_{count} for all $\theta < \frac{\ln(2)}{1+\ln(2)}$. It remained open whether there is also an efficient algorithm that is able to use the provided information as the best-known bound, where an algorithm was able to solve exact recovery efficiently, at this point, required a pooling scheme on m_{DD} tests with

$$m_{\text{DD}} = m_{\text{DD}}(n,\theta) = \max\left\{\frac{\theta}{\ln^2(2)(1-\theta)}, \frac{1}{\ln^2(2)}\right\} n^\theta \ln(n/k). \tag{2.3}$$

Thus, there remained 3 open questions within exact recovery in the Group Testing problem that will be addressed within this thesis:

(1) Can we improve the result of [87] by proposing a 2-stage algorithm that succeeds by conducting m_{count} tests?
(2) Can we propose an efficient 1-stage algorithm that achieves exact recovery with the m_{inf} tests (see (2.2)) proposed by [29]?

(3) Can we find a better 1-stage procedure improving over the best known result of [29] and, thereby, find a way to get all the way down to m_{count} for all θ?

We will answer these question concerned with exact recovery in this thesis.

As we will see in Section 2.2, it will not be possible to solve exact recovery all the way down to m_{count}. Of course, exact recovery is a very strong target to chase. Therefore, in a next step one might wonder, how easing the recovery criteria might change the tests needed to succeed. Two natural candidates for such an analysis are given by weak-recovery and detection. In [101] both criteria were already analysed for one particular 1-stage procedure, the *Bernoulli-Design*. It is known by [29, 89] that there exists a pooling procedure, called *Constant-Column Design*, that information-theoretically improves over the Bernoulli-Design for exact recovery. While in the Bernoulli-Model each individual chooses tests independently with a certain probability, in the Constant-Column Design each individual chooses a fixed number of Δ tests uniformly at random. A natural step for the further understanding of weak-recovery and detection is the analysis of this particular model. Therefore, we will address the weak-recovery as well as the detection task by employing the Constant-Column Design within this thesis.

2.2. **Results.** Recall that we analyse a population of n individuals among $k \sim n^\theta$ are infected. In a first step we analyse the chances of a 1-stage pooling procedure within the exact recovery task.

Theorem 2.1 (Theorem 1.1 and 1.2 of [30])**.** *Let $\varepsilon > 0$ and*

$$m_{\text{inf}} = m_{\text{inf}}(n,\theta) = \max\left\{\frac{\theta}{\ln^2(2)(1-\theta)}, \frac{1}{\ln(2)}\right\} n^\theta \ln(n/k). \tag{2.4}$$

Within 1-stage procedures the following holds:

- *Exact recovery is easy as soon as $m \geq (1+\varepsilon)m_{\text{inf}}$.*
- *Exact recovery is impossible as soon as $m \leq (1-\varepsilon)m_{\text{inf}}$.*

This implies that there does not exist any 1-stage pooling scheme, such that exact recovery is easy all the way down to m_{count} for all $\theta \in (0,1)$. Moreover, we are able to close the gap that was left for efficient algorithms. Thus, it is indeed possible (for certain ranges of θ) to perform exact recovery efficiently all the way down to m_{count} with a suitably chosen 1-stage pooling scheme. A next question is, whether an additional pooling stage might help to achieve exact recovery with m_{count} tests for all $\theta \in (0,1)$. The following theorem confirms this consideration.

Theorem 2.2. *Let $\varepsilon > 0$, then exact recovery is easy as soon as $m \geq (1+\varepsilon)m_{count}$ with a 2-stage pooling procedure.*

We note that for $\theta < \frac{\ln(2)}{1+\ln(2)}$ the bounds for 1-stage and 2-stage procedures match, while for larger values they drift apart. Thus, there appears to be a fundamental gap between multi-stage and 1-stage procedures. The next step of the agenda is given by understanding weak-recovery. The following theorem states our results obtained for this recovery criteria.

Theorem 2.3 (Theorem 1.3 of [30] and Theorem 1 of [31])**.** *Let $\varepsilon > 0$ and*

$$m_{count} = \frac{k\ln(n/k)}{\ln(2)}. \tag{2.5}$$

Within 1-stage procedures the following holds:

- *weak-recovery is easy as soon as $m \geq (1+\varepsilon)m_{count}$.*
- *weak-recovery is information-theoretically possible on a Constant-Column Design as soon as $m \geq (1+\varepsilon)m_{count}$.*
- *weak-recovery is impossible as soon as $m \leq (1-\varepsilon)m_{count}$ if one uses a Constant-Column Design*

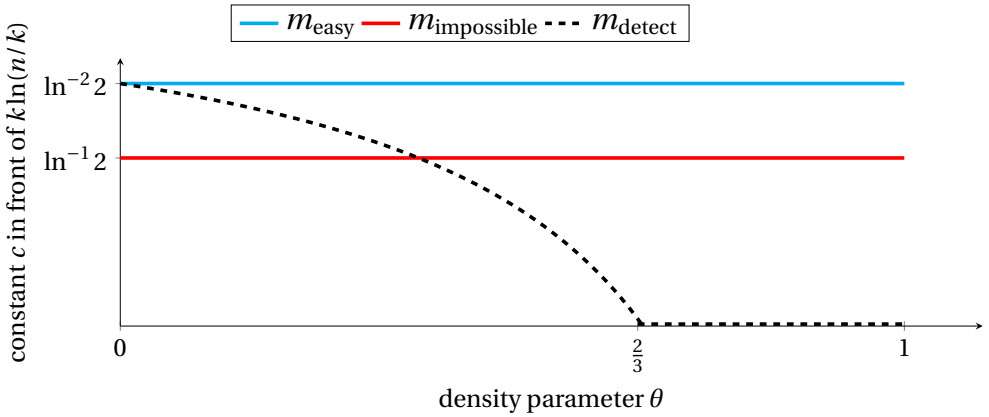

Therefore, weak-recovery is easy all the way down to m_{count}. The theorem only shows the impossibility within the Constant-Column Design. We emphasise here that the Constant-Column Design was proven to be the best-possible test design for exact recovery. Therefore, it is hard to imagine that any 1-stage pooling scheme could improve upon this result as weak-recovery is supposed to be easier than exact recovery. We leave this point as open research question and refer the reader to Section 6. In a final step, we consider the detection problem and the result is summarised in the following theorem.

Theorem 2.4 (Theorem 2 of [31])**.** *Consider the Constant-Column Design (testing variant) with parameters $\theta \in (0,1)$ and $c > 0$. Define*

$$m_{detect} = \max\left\{\left(1 - \frac{\theta}{2(1-\theta)}\right)\frac{1}{\ln^2 2}, 0\right\} k \cdot \ln(n/k). \tag{2.6}$$

Furthermore, let $m^{inf}_{detect} = \min\{m_{detect}, m_{count}\}$

(a) *If $m > m_{detect} > 0$ achieving strong detection is easy.*

(b) *If $m > m^{inf}_{detect} > 0$ achieving strong detection is information-theoretically possible.*

Note, that there exists an algorithm that performs strong detection even below m_{count} for certain ranges of large θ. For smaller θ the algorithm is not able to achieve detection all the way down to m_{count}.

Remark 2.5. *While the result is not due to the author of this thesis, we emphasise that* [31] *indicates that detection is low-degree hard below m_{detect}. As already discussed in Section 1.2.2 this carries evidence that the result obtained in Theorem 2.4 a) is best possible for the Constant-Column Design. Furthermore, combining this low-degree hardness result with Theorem 2.4 b) indicates that there may exist a 'information-theoretically possible, but computationally hard' phase for detection in the Group Testing problem. Furthermore, detection is supposed to be easier than weak as well as exact recovery. Therefore, this might hint to a computational hard phase within these tasks as well. We leave the details to* [31] *and the analysis of this phenomena as open research question (see Section 6).*

2.3. **Proof strategy for exact recovery.** In [29] it was shown, that it is information-theoretically possible to perform exact recovery as long as one conducts $m > (1+\varepsilon)m_{\text{inf}}$ tests (with $\varepsilon > 0$). The first part of Theorem 2.1 says that there is an efficient algorithm that succeeds as soon as we provide this information. Therefore, to prove this part we propose a pooling scheme as well as an efficient algorithm such that exact recovery becomes easy.

We propose a 1-stage pooling scheme on m_{inf} tests that is inspired from various successes of the *spatial-coupling method* known from coding theory (see for instance

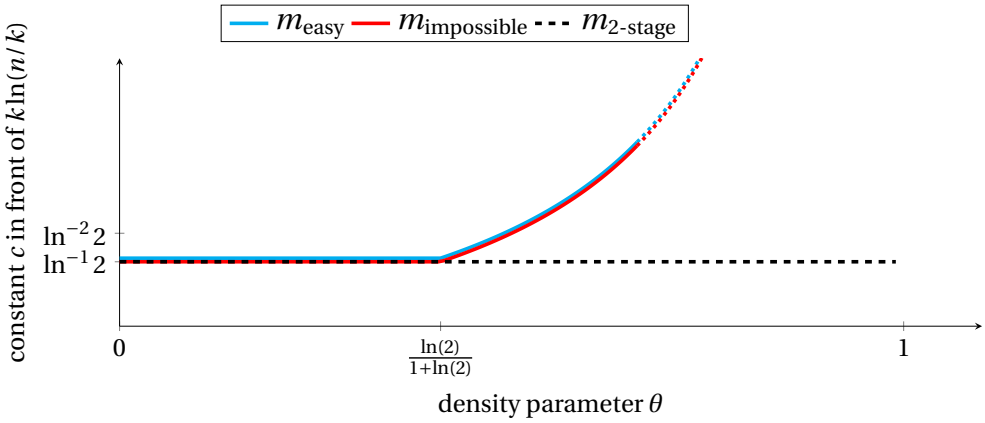

FIGURE 7. The figure deals with exact recovery. We see that there exists an efficient algorithm that succeeds all the way down to the impossible threshold for 1-stage pooling schemes (blue/red line). Furthermore, there exists a 2-stage pooling such that exact recovery becomes easy all the way down to m_{count}. Note, that there is a fundamental gap between 1-stage and 2-stage pooling schemes. The blue and red line continue to increase for all $\theta \in (0,1)$, but we omit it here due to clarity.

[65]).
We divide our n individuals as well as the m_{inf} tests into compartments $V[1],..V[T]$ and $F[1],..F[T]$ with $T = \Theta(\sqrt{n})$. Now, we take some of these compartments and mark them. These compartments will be employed as a 'seed of additional information'. Now each individual chooses $\Delta = \Theta(\ln(n))$ tests uniformly at random from its own compartment and the s subsequent compartments with $s = \Theta(\ln(\ln(n)))$. We order the compartments as a ring such that the latter individual compartments start choosing from the first ones, as soon as this becomes necessary. See Figure 8 for an illustration. Although the pooling is randomised, we introduced a certain geometry within the pooling procedure that helps us inferring the individuals' infection status. Now, we take the seed compartments and add sufficiently many test $F[0]$ such that we can infer all individuals in these seed compartments. Luckily, these additional tests needed are of order $o(m_{\text{inf}})$ and therefore vanish in the total number of tests. For details, we refer the reader to Appendix A or [30]. Thus, beside the randomised geometry within the ring, another crucial feature of our pooling scheme is given by the fact that we will be able to determine the status of the individuals contained in the seed and that these individuals are placed in other tests along the ring as well. Note, that each test contains approximately $\Gamma \sim n^{1-\theta}$ individuals.

Having a new pooling scheme is not sufficient for making exact recovery easy. Therefore, we also introduce an (efficient) algorithm that works appropriately on the ring induced by our new pooling procedure.
Our algorithm works in three stages. We emphasise, that all three stages are conducted on the same 1-stage pooling scheme and that no re-pooling will happen between the stages. In a first step, we employ the fact that we added sufficiently many tests such that Claim 1.10 holds for the seed compartments. As discussed, we can employ the well-known DD algorithm to infer the infection status correctly.
Now, the algorithm moves along the ring one compartment after the other. Of course, we can declare all individuals appearing in negative tests as uninfected (by definition). We are left with the infected individuals and disguised uninfected individuals. Along the ring we count currently 'unexplained' positive tests. Furthermore, we weight these

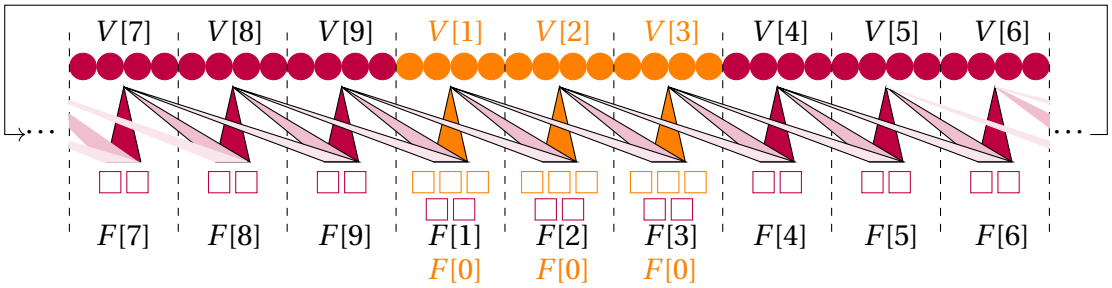

FIGURE 8. The spatially coupled test design with $n = 36, \ell = 9, s = 3$. The individuals in the seed groups $V[1] \cup \cdots \cup V[s]$ (orange) are equipped with additional test $F[0]$ (orange rectangles). The purple rectangles represent the tests $F[1] \cup \cdots \cup F[\ell]$. This is adopted from [30].

unexplained tests according to the relative position between the individuals' compartment and the test's compartment. The closer the unexplained test's position is located to the individual's position in the ring, the higher its weight. This gives us a first estimate of the infection status vector.
In a final step, we use this estimate and combine it with the fact that for $m > m_{\text{inf}}$ each infected individual is the only infected individual in a certain amount of its tests.

The proof is based on a careful analysis of evolution of the distributions along the ring. In the end, we check that the distributions induced for an infected individual and an uninfected individual are well-separated for the estimates of the algorithm. Thus, we apply a threshold approach to separate the two individuals types. We established that exact recovery is easy as soon as we choose $m > m_{\text{inf}}$ as we proposed a combination of pooling scheme and efficient algorithm that succeeds to perform exact recovery.

The second part of Theorem 2.1 says that there cannot be any better 1-stage pooling scheme for successful inference.
We use Claim 1.9 and show that as soon as we move below m_{inf} we have many disguised individuals (both infected and uninfected) for any pooling scheme. For $\theta < \frac{\ln(2)}{1+\ln(2)}$ this immediately follows from the counting bound argument (2.1). For $\theta > \frac{\ln(2)}{1+\ln(2)}$ we carefully analyse an arbitrary pooling scheme. We realise that the event of being disguised in a test is positively correlated. For θ close to 1 we can apply the well-known FKG-inequality [43] to show that in these cases we find disguised individuals as soon as we drop below m_{inf}. Now, we extend this result by moving from these large θ cases all the way down to $\theta = \frac{\ln(2)}{1+\ln(2)}$. The key idea is to carefully dilute the population by adding uninfected 'dummies' to lower the infection density. The impossibility result follows from combining the existence of many disguised individuals within the dilution process. We refer the reader to Appendix A or [30] for the details.

The 2-stage result of Theorem 2.2 directly follows from the observation that if we run the pooling scheme as well as the algorithm with $m = m_{\text{count}}$ tests, we will find only $o(k)$ many individuals that do not carry sufficient evidence to be either classified as infected or uninfected. Therefore, these are not yet classified. We now can use a second round of individual testing to get them right as well. As this additional number of tests is $o(m_{\text{count}})$, we do not increase the number of tests needed. Thus, a 2-stage algorithm succeeds by combining our optimal 1-stage pooling with individual testing.

2.4. **Proof strategy for weak-recovery.** We start with the easy part of Theorem 2.3. Here again, the observation that our pooling scheme employed for exact recovery together with the given algorithm only leaves $o(k)$ individuals unclassified while applied

on m_{count} tests shows that it is easy to solve weak-recovery, as there exists a pooling scheme as well as an algorithm that solves the weak-recovery task efficiently. Namely, the one discussed in the previous section.
As the spatially-coupled design is a special case of the Constant-Column Design, this implies that a Constant-Column Design with m_{count} tests contains sufficient information such that it is information-theoretically possible to solve the weak-recovery task. We remind ourselves that in the Constant- Column Design each individual chooses exactly Δ tests uniformly at random.

For the impossible part we have to show that Claim 1.12 holds as soon as the number of tests drops below m_{count}.
Of course, by the trivial counting bound (2.1) we know that as soon as the number of tests drops below m_{count}, we have exponentially many satisfying assignments beside the true assignment $\boldsymbol{\sigma}$. Here, we have to show that the probability that a $\tau \neq \boldsymbol{\sigma}$ has a constant overlap with the ground truth $\boldsymbol{\sigma}$ is small.
We start with removing the easy to identify uninfected individuals as well as the negative tests used for their classification. For the Constant-Column Design, this leaves us with a pooling scheme on N individuals participating in $\Delta = \Theta(\ln(n))$ of the total number of M positive tests. Let $\mathcal{N}$ denote the event that the N, M and the tests degree sequence $\Gamma = (\Gamma_1, ..., \Gamma_M)$ behave like expected. Thus, we find

$$M = \left(1 \pm n^{-\Omega(1)}\right) \frac{k\Delta}{2\ln(2)} \qquad \text{and} \qquad N = \left(1 \pm n^{-\Omega(1)}\right) n^{1-(1-\theta)c\ln^2(2)}. \tag{2.7}$$

and

$$\frac{N\Delta}{M} - \ln^2(N)\sqrt{\frac{N\Delta}{M}} \leq \min_j \mathbf{\Gamma}_j \leq \max_j \mathbf{\Gamma}_j \leq \frac{N\Delta}{M} + \ln^2(N)\sqrt{\frac{N\Delta}{M}}. \tag{2.8}$$

An immediate observation is that the tests are not independent as they share individuals. So we have to get a handle on these dependencies. Therefore, we introduce an auxiliary probability space that handles the tests as independent. Afterwards, we have to ensure that we can transfer the results obtained in the modified setting back to the original model. We will provide a road map here and refer the reader to Appendix B or [31] for more details.

Let $Z(G)$ denote the number of solution to the reduced Group Testing instance. Thus, the number of vectors τ with Hamming weight k such that each of the M tests contains at least one $\tau_i = 1$. Furthermore, let $Z^G_{\boldsymbol{\sigma}}(\alpha)$ denote the number of solutions τ' having overlap $\langle \tau', \boldsymbol{\sigma} \rangle = \lfloor \alpha k \rfloor$.
For Claim 1.11 we have to ensure

$$\sum_{\delta k \leq \ell \leq k} Z^G_{\boldsymbol{\sigma}}(\ell / k) = o\left(Z(G)\right)$$

hold with probability $1 - o(1)$. By Markov's inequality it suffices to work with $\mathbb{E}\left(Z^G_{\boldsymbol{\sigma}}(\ell/k) \,\middle|\, \mathcal{N}\right)$. For the planted Group Testing instance $\mathbb{P}_\Delta$ (where indeed a infected subset of size k is planted) this calculation turns out to be very challenging. Therefore, we introduce a first auxiliary distribution $\mathbb{Q}_\Delta$ that handles the N individuals as equals. All N individuals draw Δ tests from the M available tests.
Applying the *planting trick* from [2] to $(\mathbb{Q}_\Delta, \mathbb{P}_\Delta)$ ensures that it suffices to show

$$\sum_{\delta k \leq \ell \leq k} \mathbb{E}_{\mathbb{Q}_\Delta}\left(\boldsymbol{Z}^G(\ell/k) \mid \mathcal{N}\right) = o\left(\mathbb{E}_{\mathbb{Q}_\Delta}\left(\boldsymbol{Z}^G \mid \mathcal{N}\right)^2\right). \tag{2.9}$$

We realise that the elements within the test degree sequence fluctuate under $\mathbb{Q}_\Delta$. Thus, we introduce a regularised null model $\mathbb{Q}_{\Delta,\Gamma}$ that fixes the test degree to exactly $\Gamma = \frac{N\Delta}{M}$. We show that one does not loose 'too many' solutions through this regularisation step.

In particular, we show that

$$\mathbb{E}_{\mathbb{Q}_{\Delta,\Gamma}}\left[\boldsymbol{Z}^G\right] \leq \mathbb{E}_{\mathbb{Q}_\Delta}\left[\boldsymbol{Z}^G \mid \mathcal{N}\right]\exp\left(\delta k\Delta\right) \quad \text{and} \tag{2.10}$$

$$\mathbb{E}_{\mathbb{Q}_{\Delta,\Gamma}}\left[\boldsymbol{Z}^G(\alpha)\right] \geq \mathbb{E}_{\mathbb{Q}_\Delta}\left[\boldsymbol{Z}(G)(\alpha) \mid \mathcal{N}\right]\exp\left(-\delta k\Delta\right). \tag{2.11}$$

Therefore, it suffices to show

$$\frac{\mathbb{E}_{\mathbb{Q}_{\Delta,\Gamma}}\left(\boldsymbol{Z}^G(\alpha)\right)}{\mathbb{E}_{\mathbb{Q}_{\Delta,\Gamma}}\left(\boldsymbol{Z}^G\right)^2} \leq \exp(-\varepsilon k\Delta).$$

Thus, we transferred Claim 1.11 from the original planted model that was difficult to handle to a more accessible model. Here, we can calculate the required first and second moment bounds and the results can be transferred to the original model.

For the first moment, we count the expected number of solutions of Hamming weight k. Instead of performing the calculation for the k subset, we analyse $k\Delta$ independent half-edges. To simplify the calculation we infect each of the $N\Delta$ half-edges with probability q. Obviously, we have to ensure that all tests receive at least one infected half-edge and that we do neither over- nor undershoot the number of infected half-edges in a Group Testing instance given by $k\Delta$. Thus, if we choose q such that it satisfies

$$\frac{q}{1-(1-q)^\Gamma} = \frac{\Delta k}{\Gamma M}. \tag{2.12}$$

we find with Bayes-Theorem

$$\mathbb{E}[\boldsymbol{Z}_0^{(\Gamma)}] = N^{-O(1)}\binom{N}{k}\frac{(1-(1-q)^\Gamma)^M}{\binom{\Gamma M}{\Delta k}q^{\Delta k}(1-q)^{\Gamma M-\Delta k}}. \tag{2.13}$$

In a similar way, we handle the bound on $\mathbb{E}\left[\boldsymbol{Z}^G(\alpha)\right]$. Instead of analysing two assignments with fixed infected subsets, we again want to infect the half-edges independently with a certain probability. Of course, we again have to ensure that under both assignment each test receives at least one infected half-edge and that this independent infection process does not over -or undershoot the $k\Delta$ infected half-edges induced by the original Group Testing instance. Here, we get an additional restriction induced by α as the overlap of the two assignments matters as well. Therefore, we use $q_{11}, q_{01}, q_{10}, q_{00}$ such that we can control whether τ_i is infected in both, only in one of the two or uninfected in both. Thus, if we choose $(q_{00}, q_{01}, q_{10}, q_{11}) \in [0,1]^4$ as the solution to the system

$$q_{00}+q_{01}+q_{10}+q_{11} = 1 \qquad q_{01} = q_{10} \tag{2.14}$$

$$\frac{q_{11}}{1-2(1-q_{10}-q_{11})^\Gamma+q_{00}^\Gamma} = \alpha\frac{k\Delta}{\Gamma M} \qquad \frac{q_{01}\left(1-(q_{00}+q_{10})^{\Gamma-1}\right)}{1-2(1-q_{01}-q_{11})^\Gamma+q_{00}^\Gamma} = (1-\alpha)\frac{k\Delta}{\Gamma M}, \tag{2.15}$$

then we get with Bayes-Theorem

$$\mathbb{E}[\boldsymbol{Z}_0^{(\Gamma)}(\alpha)] = N^{-O(1)}\binom{N}{\alpha k,\,(1-\alpha)k,\,(1-\alpha)k}$$

$$\cdot\frac{\left(1-2(1-q_{01}-q_{11})^\Gamma+q_{00}^\Gamma\right)^M}{\binom{N\Delta}{\alpha k\Delta,\,(1-\alpha)k\Delta,\,(1-\alpha)k\Delta,\,(N-2k+\alpha k)\Delta}q_{11}^{\alpha k\Delta}q_{10}^{2(k-\alpha k)\Delta}q_{00}^{N\Delta-2k\Delta+\alpha k\Delta}}. \tag{2.16}$$

Now, we analyse both bounds to see whether we can ensure Claim 1.11. We start with the first moment bound and obtain

$$\mathbb{E}[\boldsymbol{Z}_0^{(\Gamma)}] = \exp\left(o\left(k\Delta\right)\right)\exp\left(k\Delta\frac{1-c\ln(2)}{c\ln(2)}\right). \tag{2.17}$$

In a second step, we reformulate (2.16) as[8]

$$G(\alpha, q_{01}, q_{11}) = o(\Delta k) + \ln\Bigg(\binom{N}{\alpha k, (1-\alpha)k, (1-\alpha)k} \cdot \frac{\left(1-2(1-q_{01}-q_{11})^{\Gamma}+q_{00}^{\Gamma}\right)^{M}}{\binom{N\Delta}{\alpha k\Delta, (1-\alpha)k\Delta, (1-\alpha)k\Delta, (N-2k+\alpha k)\Delta} q_{11}^{\alpha k\Delta} q_{10}^{2(k-\alpha k)\Delta} q_{00}^{N\Delta-2k\Delta+\alpha k\Delta}}\Bigg). \tag{2.18}$$

We use this expression to establish (2.10). Therefore, we analyse (2.18) to ensure that there exists $\varepsilon > 0$ such that for all $\hat{\alpha} \in (0,1]$,

$$G\Big(\hat{\alpha}, q_{01}(\hat{\alpha}), q_{11}(\hat{\alpha})\Big) < (1-\varepsilon)k\Delta \frac{2(1-c\ln(2))}{c\ln(2)}.$$

It is hard to analyse $G(\alpha, q_{01}, q_{11})$ as obtaining a closed form expression turns out to be difficult. Fortunately, we are only interested in an upper-bound and for any $\alpha \in (0,1]$ and any $(q_{00}, q_{01}, q_{10}, q_{11}) \in [0,1]^4$ we find,

$$\mathbb{E}[\boldsymbol{Z}_0^{(\Gamma)}(\alpha)] \leq \binom{N}{\alpha k, (1-\alpha)k, (1-\alpha)k} \cdot \frac{\left(1-2(1-q_{01}-q_{11})^{\Gamma}+q_{00}^{\Gamma}\right)^{M}}{\binom{N\Delta}{\alpha k\Delta, (1-\alpha)k\Delta, (1-\alpha)k\Delta, (N-2k+\alpha k)\Delta} q_{11}^{\alpha k\Delta} q_{10}^{2(k-\alpha k)\Delta} q_{00}^{N\Delta-2k\Delta+\alpha k\Delta}}. \tag{2.19}$$

Thus, it suffices to choose $(q_{00}, q_{01}, q_{10}, q_{11})$ in an appropriate way to construct an upper-bound that we can analyse. For

$$q_{01} = x_0 \frac{k}{N} \qquad \text{and} \qquad q_{11} = x_1 \frac{k}{N}$$

we partition the interval $(0,1)$ as

$$I_1 = \left(0, \frac{1}{4}\right], \; I_2 = \left(\frac{1}{4}, \frac{85}{100}\right), \; I_3 = \left[\frac{85}{100}, 1\right). \tag{2.20}$$

We use the following choice of x_0 and x_1, and define

$$x_0(\alpha) = \mathbf{1}_{\{\alpha \in I_1\}} \cdot \left(-\frac{3}{5}\alpha + \frac{1}{2}\right) + \mathbf{1}_{\{\alpha \in I_2\}} \cdot \left(\frac{1}{2} - \frac{3}{10\ln 2}\alpha\right) + \mathbf{1}_{\{\alpha \in I_3\}} \cdot (1-\alpha), \tag{2.21}$$

$$x_1(\alpha) = \mathbf{1}_{\{\alpha \in I_1\}} \cdot \frac{\alpha}{5} + \mathbf{1}_{\{\alpha \in I_2\}} \cdot \frac{\alpha}{5\ln 2} - \mathbf{1}_{\{\alpha \in I_3\}} \cdot \frac{16\alpha - 11}{10}. \tag{2.22}$$

By case distinction we analyse each interval separately and see that indeed for all $\alpha \in (0,1)$ we get

$$\frac{1}{k\Delta} G(\alpha, q_{01}, q_{11}) < (1-\varepsilon)\frac{2(1-c\ln 2)}{c\ln 2}$$

Indeed, the intermediate overlap contributions are smaller than the first moment squared. Now we are able to show that Claim 1.12 holds in our auxiliary model.
As we already pointed out, we can transfer this result back to the original model and Claim 1.12 still holds. This makes weak-recovery impossible in the Constant-Column Design. For details we refer the reader to [31] or Appendix B.

2.5. **Proof strategy for detection.** We have to show that detection is easy as soon as we choose the number of tests large enough. We again employ the Constant-Column Design as pooling scheme. In a first step, we remove the easy to identify uninfected individuals (participants of negative tests). We remove these individuals as well as the negative tests from the graph.
Now, we have a remainder graph on M positive tests with N remaining individuals (see (2.7) for the sizes) and each of the individuals is (by construction) contained in Δ tests.

[8] Note, that q_{01} and q_{11} determine q_{10} and q_{00} by construction.

In the detection problem, we are interested in whether or not there is a way to distinguish a Group Testing instance (drawn from $\mathbb{P}$) from a random noise instance (drawn from an appropriately chosen null-distribution $\mathbb{Q}$). Under $\mathbb{P}$ there exists a subset of size k with the constraint that each test must contain at least one of them. Thus, an infected subset that ensures all positive tests to be positive. In the null model $\mathbb{Q}$ the underlying pooling graph is just produced as a random experiment of N individuals choosing Δ of the M tests without replacement. Thus, there is no underlying infected structure in the null model.

As strong separation implies strong detection we have to find an estimator f such that Claim 1.14 holds. Our choice is

$$\mathbb{V}(\Gamma_1,\dots,\Gamma_M) = \sum_{j=1}^{M}\left(\Gamma_j - \frac{N\Delta}{M}\right)^2. \tag{2.23}$$

We observe that $\left|\mathbb{E}_{\mathbb{P}}[f] - \mathbb{E}_{\mathbb{Q}}[f]\right| = \Theta(k\Delta)$.
Furthermore, we show that the two estimations for $\mathbb{P}$ and $\mathbb{Q}$ are well-separated. By applying [104] we find that the following condition must hold to ensure the deviation of the estimator from its expectation to be small enough:

$$\sqrt{\frac{N^2}{k}} = o(k\Delta) \quad \Leftrightarrow \quad c > \left(1 - \frac{\theta}{2(1-\theta)}\right)\frac{1}{\ln^2(2)}$$

The theorem is an immediate consequence as the estimator of choice (2.23) works when applied to a Constant-Column Design with sufficiently many tests. Thus, the problem becomes easy. For details we refer the reader to [31] or Appendix B.

3. Sparsity-Constrained Group-Testing

In this section we shed light on ways to deal with an obvious artificial assumption of the standard Group Testing model. In Section 2 we have seen that in the optimal setting an individual participates in $\Delta = \Theta(\ln(n))$ tests and a test contains approximately $\Gamma = \Theta\left(n^{1-\theta}\right)$ individuals. In real world application one faces two essential difficulties that, indeed, contradict with the values of Δ and Γ necessary for the optimal pooling obtained in [30]. First of all, dilution effects appear such that the chemical signal (e.g. concentration of molecules) of the virus might get too weak. Prominent examples for such effects are given by HIV [105] and COVID-19 [79]. Moreover, the finite sample size (e.g. blood) per individual causes that an individual can only be tested a certain number of times. We again want to analyse the information-theoretic thresholds as well as the algorithmic-thresholds in this constrained model. This section is based on [51].

3.1. Related work.

The most relevant prior work is given by [49]. The authors analysed the case $\Delta = o(\ln(n)$ as well as $\Gamma = o\left(n^{1-\theta}\right)$ via the COMP algorithm where all individuals in negative tests are declared uninfected and all remaining individuals as infected. Informally their bounds read as follows:

- Δ-constrained model:
  - **(Converse)** For $\Delta = o(\ln n)$, any test design (with tests conducted in parallel) with error probability at most ξ requires $m \geq \Delta k\left(\frac{n}{k}\right)^{\frac{1-5\xi}{\Delta}}$, for sufficiently small ξ and sufficiently large n . (Theorem 4.1 in [49])
  - **(Achievability)** Under a suitably-chosen random test design and the COMP algorithm, the error probability is at most ξ provided that $m \geq \lceil e\Delta k\left(\frac{n}{\xi}\right)^{\frac{1}{\Delta}}\rceil$. (Theorem 4.2 in [49])

- Γ-constrained model:
  - **(Converse)** For $\Gamma = \Theta\big(\big(\frac{n}{k}\big)^{\beta}\big)$ with $\beta \in [0,1)$, any test design (with tests conducted in parallel) with error probability at most ξ requires $m \geq \frac{1-6\xi}{1-\beta} \cdot \frac{n}{\Gamma}$, for sufficiently large n. (Theorem 4.5 in [49])
  - **(Achievability)** Under a suitably-chosen random test design and COMP recovery, for $\Gamma = \Theta\big(\big(\frac{n}{k}\big)^{\beta}\big)$ with $\beta \in [0,1)$ and $\xi = n^{-\zeta}$ with $\zeta > 0$, the error probability is at most ξ when $m \geq \lceil \frac{1+\zeta}{(1-\theta)(1-\beta)} \rceil \cdot \lceil \frac{n}{\Gamma} \rceil$. (Theorem 4.6 in [49])

While their bounds leave a significant gap, our analysis will improve upon their results and will almost entirely close these remaining gaps.

3.2. **Results.** Here, we obtain information-theoretic as well as algorithmic thresholds for the size-constrained models. In this section, we focus on 1-stage procedures only. We consider the case $\Delta = o(\ln(n)$ as well as $\Gamma = \Theta(1)$. Therefore, we consider the following two adoption of the Group Testing model defined in Definition 1.7:

- The Γ-constrained model: We have n individuals upon which k are infected. We are allowed to conduct m tests in parallel. The output is binary and the tests are error-free. We impose the additional constraint that a test may only contain $\Gamma = \Theta(1)$ individuals.
- The Δ-constrained model: We have n individuals upon which k are infected. We are allowed to conduct m tests in parallel. The output is binary and the tests are error-free. We impose the additional constraint that an individual may only participate in $\Delta = o(\ln(n))$ tests.

For both models we are interested in exact recovery (Definition 1.8).

We pinpoint algorithmic as well as information-theoretic thresholds in these two models. The proof is based on combining a careful analysis of the combinatorial properties induced by the underlying pooling schemes (size-constraints) with the structural properties induced by the Group Testing instance itself. We will discuss the proof strategy of the Δ-constrained model in Section 3.3 and the Γ-constrained model in Section 3.4.

Theorem 3.1 (Δ-constrained model, Theorem 3.1, 3.2 and 3.3 of [51])**.** *Let $\varepsilon > 0$, $\theta \in (0,1)$ and define*

$$m_{\Delta,inf} = \max\left\{e^{-1}\Delta k^{1+\frac{1-\theta}{\Delta\theta}}, \Delta k^{1+\frac{1}{\Delta}}\right\}$$

and

$$m_{\Delta,alg} = \max\left\{\Delta k^{1+\frac{1-\theta}{\Delta\theta}}, \Delta k^{1+\frac{1}{\Delta}}\right\}.$$

Within the Δ-constrained model the following holds

- *Exact recovery is easy as soon as $m \geq (1+\varepsilon)m_{\Delta,alg}$.*
- *Exact recovery is impossible as soon as $m \leq (1-\varepsilon)m_{\Delta,inf}$.*

Here, we note that our algorithmic as well as information-theoretic bound almost match. While the asymptotic order of tests matches for all $\theta \in (0,1)$, we see, that there remains a constant e^{-1}-factor gap between the two bounds for $\theta < 1/2$.

Theorem 3.2 (Γ-constrained model, Theorem 4.1, 4.10 and 4.18 of [51])**.** *Let $\varepsilon > 0$, $\theta \in (0,1)$ and define*

$$m_{\Gamma} = \max\left\{\left(1 + \left\lfloor \frac{\theta}{1-\theta} \right\rfloor\right)\frac{n}{\Gamma}, \frac{2n}{\Gamma+1}\right\}.$$

Within the Γ-constrained model the following holds

- *Exact recovery is easy as soon as $m \geq (1+\varepsilon)m_{\Gamma}$.*
- *Exact recovery is impossible as soon as $m \leq (1-\varepsilon)m_{\Gamma}$.*

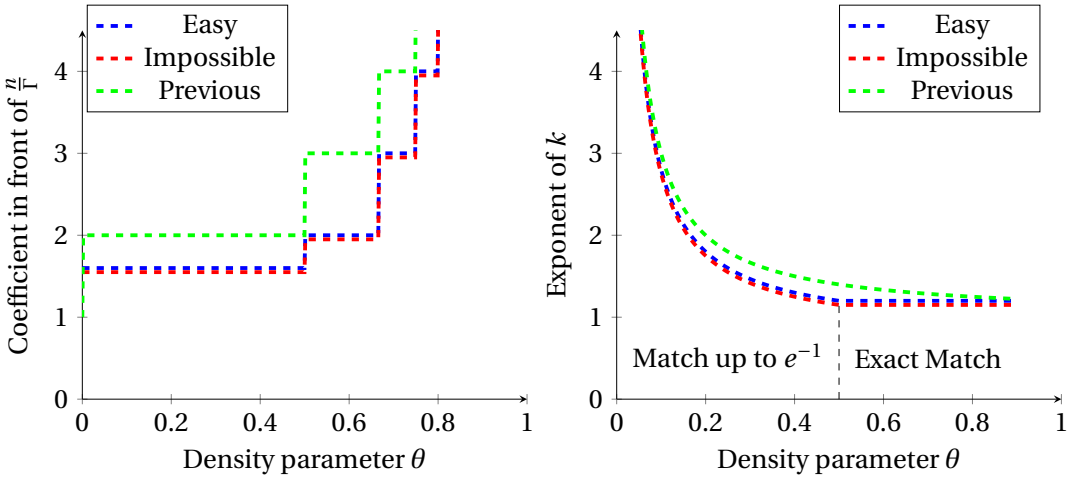

FIGURE 9. The information-theoretic as well as algorithmic threshold for exact recovery in the Δ-constrained model (right-hand side) and Γ-constrained model (left-hand side). For $\Delta = 5$ and $\Gamma = 4$. For the Γ-constrained model the C of $m = (C + o(1))\frac{n}{\Gamma}$ is plotted for $\theta \in (0,1)$. For the Δ-constrained model the η of $m = C\Delta k^{\eta}(1 + o(1))$ is plotted for $\theta \in (0,1)$.

Note, that within the Γ-constrained model we do find an easy-impossible transition as the algorithmic as well as the information-theoretic bound match. We will propose the optimal pooling scheme together with the suitable optimal algorithm in Section 3.4. For both models, the main ideas are based on adopting Claim 1.9 and Claim 1.10 to the constrained models. To satisfy Claim 1.10 we employ well-known algorithms (stated in Algorithm 1). An illustration of the obtained bounds for both

1 Declare every individual x that appears in a negative test as uninfected; remove all such individuals.
2 Declare all individuals that are now the sole individual in a (positive) test as infected.
3 Proceed as follows depending on the algorithm:
- For DD, declare all remaining individuals as uninfected.
- For SCOMP, repeat the following step until no unexplained[9]positive tests remain: Declare as infected the (previously undeclared) individual in the largest number of unexplained positive tests.

Algorithm 1: The DD and SCOMP algorithms as defined by [6].

models can be found in Figure 9.

3.3. **Proof strategy for the Δ-constrained model.**

3.3.1. *On the easy-threshold.* Here, we want to establish that it is, indeed, possible to infer the infected set given the pooling scheme and the test results. Therefore, we use Claim 1.10 to address it. Thus, we have to find a pooling scheme as well as an efficient algorithm that succeeds with high probability. Our choice is a combination of the Constant-Column Design and Algorithm 1. While [29, 30] showed that this is the best possible pooling in the unconstrained model, we employ this pooling in the size-constrained model as well. For this pooling procedure each individual chooses its tests uniformly at random with replacement. Now, we analyse the properties of the underlying pooling graph and are able to show that certain properties hold with high probability. In a next step, we analyse the individual types to see whether or not the

DD-Algorithm succeeds or not. This is the case as soon as Claim 1.10 is satisfied. The result is obtained by a careful large deviation analysis of the underlying distributions on the pooling scheme. Thereby, we carefully check, which graph structures occur and how they help/harm the algorithmic performance. In particular, we analyse the occurrence of easy to identify individuals (1.1),(1.2) and disguised uninfected individuals (1.3) under a Δ-regular pooling. We see that Claim 1.10 holds as soon as we surpass the number of tests claimed in the theorem. Thus, the DD-algorithm is able to infer the infected set efficiently when applied to the Constant-Column Design. Details can be found in III-E of [51].

3.3.2. *On the impossible-threshold.* The main idea for deriving an impossible regime for exact recovery in Group Testing is the use of Claim 1.9. We have to argue that there is not sufficient information contained to infer the correct infection status. The first part of our lower bound is based on an universal information-theoretic lower-bound. We upper-bound the success probability of an algorithm that is working on a pooling graph $\mathcal{G}$ with k infected individuals. We use the fact that the best-possible inference algorithm cannot do better than drawing an uniform sample from $S_k(\mathcal{G})$ (Corollary 2.1 of [29]), and we are able to obtain an upper-bound depending on n, k and m. Setting this upper-bound strictly smaller than 1 and solving for m leads to a lower-bound on m^* such that any algorithm applied to any pooling scheme with less than m^* test will have a non-trivial error probability.
The second part is based on a careful analysis of an arbitrary pooling graph. Instead of directly analysing the original model with an underlying infection vector of Hamming weight k, we pass over to an auxiliary model where an individual is infected independently with a certain probability. The target is to show that one finds many disguised individuals in this auxiliary model and transfer this result back to the original model. We apply the two-round exposure technique in the auxiliary model to ensure that we find many disguised individuals as soon as we choose the number of tests below a certain threshold, the threshold claimed in the theorem. In a first step, we build a set of infected individuals $\mathcal{K}$ by marking them as infected independently. In a second step, we turn to the second neighbourhood of $\mathcal{K}$ and mark these individuals as infected with a certain probability. This procedure enables us to lower-bound the average probability of being disguised in the auxiliary model. By translating this probability from the auxiliary to the original model we see, that with non-trivial probability there are also disguised individuals in the original model. Therefore, combining the existence of many disguised individuals with Claim 1.9, we see that exact recovery is impossible as soon as we cross the threshold claimed in the theorem. The details can be found in Section III-C and III-D of [51]

3.4. **Proof strategy for the Γ-constrained model.**

3.4.1. *On the easy-threshold.* Again we want to find a pooling scheme as well as an efficient algorithm to solve exact recovery. In this model the choice is more delicate than in the previous model, as we have to ensure that each test contains at most $\Gamma = \Theta(1)$ individuals. Our choice is the following:

$$\tilde{\mathcal{G}_\Gamma}(\theta) = \begin{cases} \mathcal{G}_\Gamma & \text{if } \theta \geq 1/2 \\ \mathcal{G}_\Gamma^* & \text{otherwise} \end{cases} \tag{3.1}$$

To obtain $\mathcal{G}_\Gamma$ we employ the configuration model. We clone each individual Γ times and each test $\Delta = \frac{m\Gamma}{n}$ times. Afterwards, we build a perfect matching upon these clones. Furthermore, we obtain $\mathcal{G}_\Gamma^*$ in a three-step procedure. First, we select a set of γ individuals uniformly at random. Now, we again build a regular graph (individuals with degree 2 and tests with degree $\Gamma - 1$) via the configuration model. In a final step, we match γ individuals with the graph obtained in step two. The result is the pooling

graph we will work on. We again analyse the properties of the underlying graph. In particular, the $\Gamma = \Theta(1)$ is causing the analysis to be more delicate. We carefully check how the different individual types are influencing the graph. We see that DD succeeds on the regular part of our pooling scheme, but for certain sparsity levels it is necessary to add the additional greedy steps of SCOMP to succeed on the entire pooling graph with high probability. This is due to the fact that adding the γ individuals in the final step of the pooling procedure may confuse DD, while SCOMP works properly. The performance guarantee of the proposed pooling with the stated algorithm matches the information-theoretic lower-bound. Thus, leading to an optimal choice. The details can be found in Sections IV-D and IV-E of [51].

3.4.2. *On the impossible-threshold.* The argument is two-fold. For the first part of the statement, we consider a large infection spread (dense case). The second part is due to sparse infection spread (sparse case).

In the dense case, we again consider an auxiliary model that is easier to handle. Therefore, we argue that we find many disguised individuals in this model. In the end, we transfer these findings back to the original model. We again employ an auxiliary model, where we infect individuals independently with a certain probability. The proof hinges on two main observations. First of all, we see that the property of being disguised is a local property. Thus, two individuals, that do not share any test as well as no fellow individuals in their tests (distance at least 6), happen to be disguised independently. Secondly, we use the fact that additional infected individuals increase the probability of being disguised. Thus, we employ the FKG-inequality to lower-bound the probability of being disguised. In the end, we use these two observation to argue that we find many disguised individuals in the auxiliary model as soon as we surpass the number of tests claimed in the theorem. Now, we argue that the occurrence of many disguised individuals transfers from one to the other model.

In the sparse case, the crucial observation used is also two-fold. First of all, it is immediately clear that tests with only one individual are just individual tests and directly reveal the infection status of the contained individual. Second of all, as soon as a positive test contains more than one individual of degree one, inference of these individuals is impossible. This follows as we cannot tell which of the individuals is responsible for the positive test, but as they are only contained in one test there is no further information available about these individuals. Therefore, the best we can do in that case is guessing.
We show that as soon as the number of tests drops below the threshold claimed in the theorem, we find many individuals with degree 1. Furthermore, we show that a successful inference algorithm requires the number of tests, containing multiple degree 1 individuals, to be small.
Combining these statements yields the theorem. The details can be found in Section IV-B of [51].

4. Noisy Group-Testing

We adjust the standard model by assuming that the obtained test results may not meet the gold standard, thus, being error-prone. See Figure 10 for illustration. Note, that the content of this section is based on [52].

4.1. **Related work.** We dealt with the standard Group Testing model in one of the previous sections (Section 2). Some of the assumptions seem quite artificial. One of them is the fact that 'a test returns positive if and only if an infected individual is contained'.

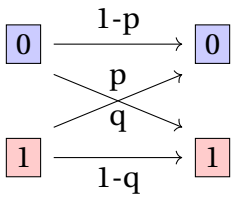

FIGURE 10. The p-q-noise model: Each result of the standard noiseless group test is noisy and gets potentially flipped. The status of positive as well as negative tests are flipped or kept with a certain probability.

Obviously, this does generally not hold in medical application. Thus, we cannot assume the gold standard within our testing procedure. A prominent example where the tests are not error-free is given by COVID-19 [106]. The two important types of noise in the medical scope are sensitivity (positive correct) and specificity (negative correct). It is important to realise that values for sensitivity and specificity are usually not equal [69]. Therefore, one step towards understanding Group Testing in a more realistic setup is the analysis of noisy Group Testing. Due to the medical requirements an analysis of a general $p-q$-model is necessary (see Figure 10). The idea of extending the model to noisy measurements raised a lot of attention [22, 87, 88, 89, 90, 92]. We are interested in the number of tests necessary such that the task of exact recovery becomes easy. Two well-known algorithms in the Group Testing literature are COMP and DD. In [22, 92] noisy variants were established (see Algorithm 2 and Algorithm 3). So far, the only algorithmic performance guarantees for noisy Group Testing were obtained for the Bernoulli-model, where each individual chooses to participate in a test independently with probability p [22, 92, 90]. In this setting the noisy COMP was already analysed by [22].

1 Declare every individual that appears in $\alpha\Delta$ or more displayed negative tests as healthy.
2 Declare all remaining individuals as infected.

Algorithm 2: The noisy COMP algorithm

The noisy DD was already analysed by [92] for the Bernoulli-model.

1 Declare every individual that appears in $\alpha\Delta$ or more displayed negative tests as healthy and remove such individual from every assigned test.
2 Declare every yet unclassified individual who is now the only unclassified individual in $\beta\Delta$ or more displayed positive tests as infected.
3 Declare all remaining individuals as healthy.

Algorithm 3: The noisy DD algorithm [92]

As [30] (compare Section 2) showed that the Constant-Column Design is the optimal choice in the noiseless case we take first steps to transfer this result to the noisy variant of the Group Testing problem by analysing the performance of the two most common algorithms.

4.2. **Results.** In the noisy model we derive the thresholds such that the two algorithms succeed performing exact recovery within the Constant-Column Design. The formal statement for noisy COMP reads as follows:

Theorem 4.1 (Noisy COMP,Theorem 2.1 in [52])**.** *Let $p,q \geq 0$, $p+q<1, d \in (0,\infty), \alpha \in (q, e^{-d}(1-p)+\left(1-e^{-d}\right)q)$. Suppose that $0<\theta<1$ and let*

$$m_{COMP} = m_{COMP}(n,\theta,p,q) = \min_{\alpha,d} \max\{b_1(\alpha,d), b_2(\alpha,d)\}\, k\ln(n/k)$$

$$\text{where} \qquad b_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|q\right)}$$

$$\text{and} \qquad b_2(\alpha,d) = \frac{1}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}.$$

If $m \geq (1+\varepsilon)m_{COMP}$ for some $\varepsilon>0$, exact recovery is easy via noisy COMP.

For noisy DD we obtain the following statement:

Theorem 4.2 (Noisy DD,Theorem 2.2 in [52])**.** *Let $p,q \geq 0$, $p+q<1, d \in (0,\infty), \alpha \in (q, e^{-d}(1-p)+\left(1-e^{-d}\right)q)$ and $\beta \in (0, e^{-d}(1-q))$ and define $w = e^{-d}p+(1-e^{-d})(1-q)$. Suppose that $0<\theta<1$ and let*

$$m_{DD} = m_{DD}(n,\theta,p,q) = \min_{\alpha,\beta,d} \max\left\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\alpha,\beta,d)\right\} k\ln(n/k)$$

$$\text{where} \qquad c_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|q\right)}$$

$$\text{and} \qquad c_2(\alpha,d) = \frac{1}{dD_{\mathrm{KL}}\left(\alpha\|1-w\right)}$$

$$\text{and} \qquad c_3(\beta,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\beta\|(1-q)e^{-d}\right)}$$

$$\text{and} \qquad c_4(\alpha,\beta,d) = \max_{1-\alpha\leq z\leq 1}\left\{\frac{1}{1-\theta}\frac{1}{d\left(D_{\mathrm{KL}}\left(z\|w\right)+\mathbf{1}\left\{\beta>\frac{ze^{-d}p}{w}\right\}zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right)}\right\}.$$

If $m \geq (1+\varepsilon)m_{DD}$ for some $\varepsilon>0$, then exact recovery is easy via noisy DD.

We derive conditions under which exact recovery becomes easy in Theorem 4.1 and Theorem 4.2. Furthermore, we rigorously prove the improvement upon the results of [22, 92] for a large set of parameter and, thereby, set the strongest performance guarantees rigorously proved for (efficient) exact recovery in the general $p-q$-model. Furthermore, we conduct a Shannon-Capacity analysis for the $p-q$-noise model to put our results into a channel-perspective. We omit these results here and refer the reader to [52]. In the end, we apply our generalised results to the standard channels. To make the bounds more accessible, we illustrate the obtained bounds in Figure 11 for both algorithms applied at different noise levels. All details can be found in Appendix D or [52]. In the following we provide a proof idea. We will discuss further directions in the scope of the noisy Group Testing model in Section 6.

4.3. **Proof strategy for the noisy model.** We follow the standard target within the analysis of the easy regime of algorithms. Here, we have to find a property that separates infected and uninfected individuals. Furthermore, the algorithms we analyse have to be able to use this property efficiently. As already discussed, we have two choices to make in Group Testing:

- The pooling scheme: We choose the Constant-Column Design.
- The inference algorithm: We choose COMP and DD.

Therefore, we carefully analyse the influence of infected individuals and uninfected individuals on the pooling scheme. Thereby, we realise that we, indeed, can distinguish their neighbourhood structure.
For the noisy COMP algorithm we use the fact that in the pre-noise setting each infected

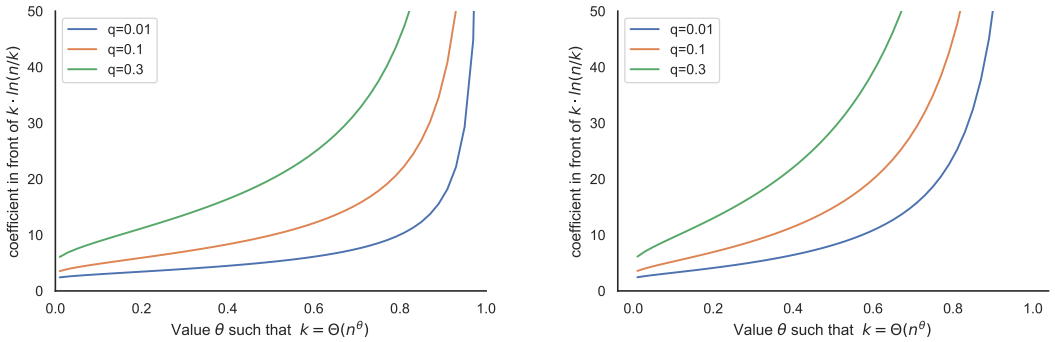

FIGURE 11. Evolution of the bounds for noisy DD (left) and noisy COMP (right) in the Z-channel for different noise level. Exact recovery is easy for the corresponding algorithm above the line and hard below. Note, that a similar evolution can be observed in the reverse Z, the Binary symmetric channel and general choices of q and q, but these are omitted here.

individual will cause all its neighbouring tests to be positive. Of course, in the post-noise setting some of the tests may be displayed negative. Now, we carefully derive the underlying distributions and apply Chernoff-Bounds to reveal the number of tests ensuring the number of displayed negative tests (DN) to be well-separated for all infected and uninfected individuals with high probability (illustrated on the left hand side of Figure 12). The analysis of the noisy DD is more delicate. While the first round is quite similar to the noisy COMP analysis, we ease the requirements for the initial estimation. Instead of exact recovery in the first step, we only need a first estimate and can leave some of the uninfected individuals $\big($a set of size $o\,(n^\eta)$ for some $\eta \in (0,1)\big)$ as undeclared. Thus, this estimation requires less tests in the first round and we need to ensure that the algorithm is able to fix this in the final round (conducted on the same pooling instance, no re-pooling). In the second step we realise that we can distinguish infected and so-far undeclared uninfected individuals by the following criteria, which is illustrated on the right hand side in Figure 12:

- Type Displayed-Positive-Single (DP-S): Displayed positive tests in which all other individuals are already declared as uninfected.
- Type Displayed-Positive-Multiple (DP-M): Displayed positive tests with at least one other individual that is not contained in the estimated set of uninfected individuals.

The crucial observation is the fact that in the pre-noise setting a uninfected individual needs another infected individual in the test to render a test positive, while an infected individual does not carry this requirement. We transfer this observation by carefully applying Chernoff-Bounds to the noisy case. We derive conditions for the number of tests m such that infected and uninfected individuals are well separated. The comparison with previous results as well as the application to the standard channels (Z, reverse Z, binary symmetric) is conducted by applying Theorem 4.1 and Theorem 4.2 in the corresponding settings.

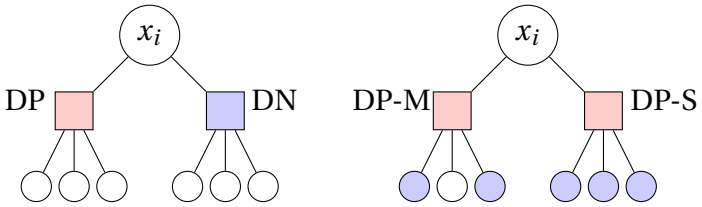

FIGURE 12. The relevant neighbourhood structures for the analysis of the algorithms, on the left for the first stage and on the right for the second step. Rectangles represent tests (displayed positive in red, displayed negative in blue). Blue circles represent individuals that have been classified as healthy in the first step of DD (or by COMP). White circles represent individuals that are unclassified in the current stage. We refer to displayed negative tests as Type DN, displayed positive tests as Type DP, displayed positive with a single unclassified individual as Type DP-S and displayed positive with a multiple unclassified individual as Type DP-M. This figure is adopted from [52].

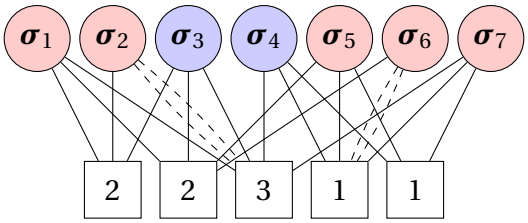

FIGURE 13. A small example with population infection $\boldsymbol{\sigma} = (1,1,0,0,1,0,0) \in \{0,1\}^7$ at the top and tests $a_1,\dots,a_5$ at the bottom. The edges of the bipartite (multi-) graph $\boldsymbol{G}$ indicate which individual is contained in which test. The dashed lines highlight the occurrence of multi-edges. The goal is to reconstruct $\boldsymbol{\sigma}$ given only $\boldsymbol{G}$ and the test results $(2,2,3,1,1)$.

5. QUANTITATIVE GROUP TESTING

Here we adopt the quantitative model of Definition 1.7. This section is based on [50][10]. While we have reliable tests and do not require to satisfy any size-constraints, we do not have a binary output anymore. In this section, we assume that each test returns the number of infected individuals contained in the test. We refer the reader to Figure 13 for a small-size example.

5.1. **Related work in the quantitative group testing model.** Until now, we considered tests that return whether an infected individual is contained in the test or not. At this point, we might wonder how receiving additional information might simplify the problem. One way to gain additional information is given by applying the quantitative model. In this case, we know exactly how many infected individuals are contained in a test. While this assumption might seem artificial at first glance, such testing procedures found their way into many real-world applications such as PCR tests in a bio-medical context [14], biological processes such as DNA screening [20, 94], or deep neural network on a GPU [68].
Starting with the early works of Djackov [39], and Shapiro [95], the quantitative variant

[10]Note, that the paper handles the reconstruction of a signal vector with Hamming weight k from additive queries. To keep it close to the other contributions contained in this thesis, we think of it as a population with k infected individuals and the tests return the number of infected individuals contained in the test.

of the Group Testing model raised some attention over the years. Again, we are interested in the information-theoretic as well as the algorithmic thresholds of this problem set.

Over the years, some steps towards the answer of this question were taken. While a simple counting bound shows that at least

$$m_{\text{count}}^{quant} > \frac{\ln\left(\frac{n}{k}\right)}{\ln(k)}k$$

tests are necessary for exact recovery, [39] showed that exact recovery is impossible as soon as the number of tests drops below $m_{\text{Imp}} = 2m_{\text{count}}^{quant}$. The remaining question is whether or not we can (not) achieve the bound information-theoretically or even efficiently. Thereby, Grebrinski and Kucherov [54] showed that exact recovery is information-theoretically possible by conducting
$m_{\text{GreKu}} = 4m_{\text{count}}^{quant}$. So far, the results do not depend on the scaling of k. In the case of $k = \Theta(n)$, it is known that as soon as $m > (1+\varepsilon)m_{\text{imp}}$, it is information-theoretically possible to solve exact recovery [5, 91]. Their results do not extend to our scaling $k = n^\theta$ with $\theta \in (0,1)$ and it remained open whether or not a similar transition occurs in the sublinear scaling regime. Therefore, we address (and are indeed able to close) the remaining gap between the results of Djackov and Grebinski/Kucherov in Theorem 5.1. Please note, that [42] achieved the same improvement by applying different methods.

Now, we turn to efficient algorithms. Different sophisticated algorithms were proposed to solve exact recovery in the setting at hand [40, 42, 44, 61, 62, 85]. All these algorithms (including the one proposed in [50]) required $m = \Theta(k\ln(n))$ and thereby leave a gap of $\Theta(\ln(n))$ to the information-theoretic possible threshold. Note, that Hahn-Klimroth and Müller [55] were able to close this gap up to a constant factor and thereby improve over the results of [50].

5.2. **Results.** We consider the noiseless, unconstrained quantitative model of Definition 1.7. Furthermore, we assume the infection spread to scale as $k \sim n^\theta$ within the population with $\theta \in (0,1)$. We are interested in exact recovery and pinpoint the threshold beyond which exact recovery is information-theoretically possible. Furthermore, we propose a pooling scheme as well as a greedy-type algorithm that solves exact recovery efficiently. We pinpoint the number of tests necessary for the algorithm to succeed.

Theorem 5.1 (Theorem 1 and Theorem 2 of [50])**.** *Let* $0 < \theta < 1$, $k = n^\theta$ *and* $\varepsilon > 0$ *and let*

$$m_{Pos}(n,\theta) = 2\frac{k\ln(n/k)}{\ln(k)} = 2\frac{1-\theta}{\theta}k,$$

$$m_{Easy}(n,\theta) = 4\left(1-\frac{1}{\sqrt{e}}\right)\frac{1+\sqrt{\theta}}{1-\sqrt{\theta}}k\ln(n/k).$$

Then exact recovery is possible as soon as $m > (1+\varepsilon)m_{Pos}$ *and easy as soon as* $m > (1+\varepsilon)m_{Easy}(n,\theta)$.

While we derive a possible and an easy regime for exact recovery in the model at hand, we note that the order of the thresholds obtained in Theorem 5.1 do not match.

5.3. **Proof strategy.** For both of our bounds we assume a Constant-Column Design $\boldsymbol{G}$ where each test draws $n/2$ individuals uniformly at random with replacement. Now,

by combining $\boldsymbol{G}$ with the ground truth $\boldsymbol{\sigma}$ we obtain our test result vector

$$\hat{\boldsymbol{\sigma}}_j = \sum_{x_i \in \partial a_j} \boldsymbol{\sigma}_i$$

5.3.1. *The possible-threshold.* The main idea is based on counting alternative assignments $\tau \in \{0,1\}^n$ with Hamming weight k that lead to the same sequence of test results as induced by the ground truth $\boldsymbol{\sigma}$. We start by counting alternative assignments with a certain overlap. Let $Z_{k,\ell}$ denote the number of assignments with Hamming weight k that are consistent with the test result $\hat{\boldsymbol{\sigma}}$ and have overlap ℓ with $\boldsymbol{\sigma}$. Formally, we analyse

$$Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) = \left|\left\{\sigma \in S_k(\boldsymbol{G}, \boldsymbol{y}) : \sigma \neq \boldsymbol{\sigma}, \langle \boldsymbol{\sigma}, \sigma \rangle = \ell\right\}\right|.$$

It is sufficient to show that $\sum_{\ell=0}^{k-1} Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) = 0$ for $m \geq (1+\varepsilon) m_{\mathrm{Pos}}$ w.h.p.. We achieve this by a two-fold argument. In a first step, we combine a first moment calculation with Markov's inequality to show that there does not exist an alternative assignment with small overlap. We obtain this result by realising that the probability of arriving at an alternative assignment can be approximated via a returning random walk, which simplifies the first moment calculation. Within our first moment analysis, we have two opposite powers at work. On the one side, the entropy is increasing in n. On the other side, the probability term is decreasing in n. Now, we choose the number of tests m large enough, such that the probability is strong enough to drag the entropy down and ensure that the first moment tends to zero. Thus, as soon as the first moment tends to zero we employ Markov's inequality to ensure that no alternative assignment with small overlap exists with high probability.
For the large overlap case, we employ a standard coupon-collector argument. In the end, our analysis reveals that as long as $m > (1+\varepsilon) m_{\mathrm{Pos}}$ we have only one satisfying assignment, the ground truth. Thus, there is, indeed, sufficient information contained to infer $\boldsymbol{\sigma}$ and m_{Pos} matches m_{Imp}. Details can be found in Section IV of [50].

5.3.2. *The easy-threshold.* To prove that the problem is easy beyond a certain point, it suffices to show that there exist a polynomial-time algorithm which returns the correct underlying $\boldsymbol{\sigma}$ with high probability. Therefore, we introduce Algorithm 4. This algorithm takes the pooling graph as well as the test result as an input. Then, it calculates the sum of all test results in the neighbourhood ∂x of each individual x and centralises the result by its mean. In the end, the algorithm declares the k individuals with the largest neighbourhood sum as infected. The key idea is that, of course, the neighbour-

Input: m, k, querying method `query`
Output: estimation $\tilde{\boldsymbol{\sigma}}$ for $\boldsymbol{\sigma}$.

for $i = 1$ **to** m **do in parallel**
  sample a multiset a_i of size Γ from $[n]$
  compute $\boldsymbol{y}_i \leftarrow \mathtt{query}(a_i)$
  // The query method guarantees that $\boldsymbol{y}_i = \sum_{j \in a_i} \boldsymbol{\sigma}(j)$.
for $i = 1$ **to** n **do**
  calculate $\Psi_i \leftarrow \sum_{j=1}^{m} \mathbf{1}\{i \in a_j\} \cdot \boldsymbol{y}_j$
  calculate $\boldsymbol{\Delta}_i^{\star} \leftarrow \sum_{j=1}^{m} \mathbf{1}\{i \in a_j\}$
sort coordinates of $\tilde{\boldsymbol{\sigma}}$ in decreasing order by $\Psi_i - \boldsymbol{\Delta}_i^{\star} \frac{k}{2}$
set $\tilde{\boldsymbol{\sigma}}$ to 1 for the first k (sorted) coordinates
set $\tilde{\boldsymbol{\sigma}}$ to 0 for the remaining $n-k$ (sorted) coordinates

Algorithm 4: The Maximum Neighbourhood Algorithm

hood sum between infected and uninfected individuals differ as each infected individual contributed its weight (its infection status) to all of its tests, while an uninfected individual did not do that. To prove the success of the algorithm, we employ a careful large deviation analysis ensuring that as soon as $m > (1+\varepsilon)m_{\text{Easy}}$, we find that the values for infected and uninfected individuals are well separated with high probability. The details can be found in Section III of [50].

6. Conclusion

Obviously the COVID-19 crisis showed in its own way, why understanding the fundamentals of the Group Testing problem is of major interest. In this thesis, we resolved some of the major open problems within the Group Testing community [7]. The pandemic showed that we are not able to immediately transfer our results on the standard Group Testing problem to the real world. Therefore, our results on the size-constrained as well as the noisy Group Testing problem can be seen as first steps towards a more applicable testing procedure. From a Group Testing perspective we can formulate the following open research questions:

- Can we transfer our spatially-coupled test design as well as an adaption of the algorithm to the noisy variant of the Group Testing problem?
- Can we close the remaining easy vs. information-theoretically possible vs. low-degree-hard gaps within Constant-Column Design? How does the low-degree hardness bound for detection transfer to weak and exact recovery?
- How do our results for the noisy and size-constrained Group Testing problem transfer to the linear infection spread $k = \Theta(n)$?
- Can we close the remaining gaps that are left between information-theoretic and algorithmic thresholds?
- Do we find other problem sets where one finds such fundamental gaps between multi-stage and 1-stage procedures?

As we discussed in the introduction, we are interested in the hardness of various problem sets. We worked with a particular problem set and revealed structures that may be helpful or disruptive while working with a data set. We believe that this understanding might be transferred to other problems as well. For the general context, we want to highlight that the result of [31] shows that for certain choices of parameters the detection task is low-degree hard, while [30] indicates that it is information-theoretically possible. As mentioned above, the low-degree framework is often conjectured as a proxy for efficient algorithms. Therefore, it would be interesting to understand this information-theoretically possible but low-degree-hard regime [30, 31] in more detail. Furthermore, we have seen that the spatially-coupled pooling scheme can even solve exact recovery there. Therefore, a deeper understanding of the spatial-coupling as a general approach might offer a fruitful direction for understanding the fundamentals of the initial $P-NP$ question and, thereby, contribute to the general understanding of algorithms.

7. Author's Contribution

This thesis is based on [30, 31, 52, 51, 50] and the author's contributions are stated in this section.

- Optimal Group Testing,
  co-written with A. Coja-Oghlan (ACO), M. Hahn-Klimroth (MHK) and P. Loick (PL),
  Journal version:
  Combinatorics, Probability and Computing, 30(6), 811-848 (2021).
  Conference version:
  Proceedings of 33rd Conference on Learning Theory (COLT'20), PMLR 125:1374-1388.

Together with ACO, MHK and PL the author of this thesis contributed to both parts of Theorem 2.1. The authors jointly developed the pooling scheme and conducted the analysis together (part 1 of the theorem). The basic ideas behind the algorithm were developed jointly. While the unweighted algorithm analysis conducted by MHK, PL and the author left a gap, ACO was able to close this gap by improving the analysis. The converse bound (part 2 of the algorithm) is based on ideas of ACO. The author contributed to this part of the theorem by working out the formal derivation of this theorem together with the other authors. The result in Theorem 2.2 is due to PL and MHK.

- Statistical and Computational Phase Transitions in Group Testing,
  co-written with A. Coja-Oghlan, M. Hahn-Klimroth, A. Wein (AW) and I. Zadik (IZ),
  Submitted and currently under review

The author contributed to the first and second moment analysis in the auxiliary model. Based on ideas of ACO, the author derived the bounds needed for the analysis together with MHK. The author contributed to the transfer of the auxiliary model to the original model together with ACO,MHK and IZ. This closes part 3 of Theorem 2.3. Furthermore, the author developed the detection analysis as well as the performance guarantees together with ACO, MHK and AW.

- Near optimal sparsity-constrained group testing: improved bounds and algorithms,
  co-written with M. Hahn-Klimroth, O. Parczyk (OP), M. Penschuck, M. Rolvien, J. Scarlett (JS) and N. Tan (NT),
  journal version:
  Accepted and to appear in IEEE Transactions on Information Theory

The author contributed to the results for both models. Together with MHK, he worked on the easy-threshold. While the ideas behind the converse bound were mainly developed by MHK and OP, the author contributed within the development of the converse bounds within the write-up and review process. Together with MHK, OP, JS and NT the author contributed to the detailed write-up of the results and the paper.

- Improved Bounds for Noisy Group Testing With Constant Tests per Item,
  co-written with O. Johnson (OJ), P. Loick and M. Rolvien (MR),
  journal version:
  IEEE Transactions on Information Theory, vol. 68, no. 4, pp. 2604-2621 (2022)

The author derived the ideas for the bounds for the DD and the COMP algorithm together with PL. The application to the standard channels as well as the derivation of the conditions necessary for the improvement upon previous results are due to the author. Furthermore, he contributed to the write-up and dealt with the formalisation of the proofs within the review process together with OJ. The simulation results are due to MR.

- On the Parallel Reconstruction from Pooled Data,
  co-written with M. Hahn-Klimroth, P. Loick and D. Kaaser (DK),
  conference version:
  Accepted and to appear in Proceedings of 36th IEEE International Parallel and Distributed Processing Symposium (IPDPS'22)

This paper is a joint product of all four authors. While the ideas were developed by MHK, PL and DK, the author contributed in formalising the results. Together with MHK and DK he developed the final version published at IPDPS'20. He contributed to both results of Theorem 5.

8. Deutsche Zusammenfassung

8.1. **Einführung.** Diese Doktorarbeit basiert auf [30, 31, 50, 51, 52]. Im Folgenden werden die Resultate dieser Arbeiten zusammengefasst und in den Kontext eingeordnet. In unserer heutigen Welt spielen Algorithmen und Datenstrukturen eine zentrale Rolle und fast jeder Teil unseres Lebens wird von solchen beeinflusst. Hierbei drängen sich zwei Fragen unmittelbar auf:

- Warum funktionieren Algorithmen?
- Wo sind deren Grenzen?

Obwohl wir alle mit der Nutzung von Algorithmen vertraut sind, stellt sich deren mathematische Analyse oft als schwierig heraus. Da unsere moderne Welt ein sich schnell entwickelndes, komplexes System darstellt, ist das Beantworten der oben angeführten Fragen in diesem Kontext so gut wie aussichtslos. Deshalb haben Wissenschaftler begonnen, Modelle zu entwickeln, die einzelne Aspekte der realen Welt abbilden, aber durch vereinfachende Annahmen leicht zu analysieren bleiben. Nun stellt man schnell fest, dass einige Probleme zwar auf den ersten Blick leicht wirken, aber am Ende doch schwerer zu lösen sind als gedacht. Daraus ergibt sich das berühmte P-vs.-NP Problem (eines der Millennium-Probleme der Mathematik). Man unterscheidet zwischen Problemen, für die eine gegebene Lösung leicht als Lösung zu verifizieren ist (NP) und Problemen, für die das Finden einer Lösung leicht ist (P). Die fundamentale Frage ist nun, ob diese Mengen gleich sind oder nicht. Da noch kein Beweis in die eine oder andere Richtung vorliegt, versuchen Wissenschaftler Hinweise für $P = NP$ oder $P \neq NP$ zu finden. Das Interesse beschränkt sich nicht nur auf Mathematiker und Informatiker, sondern auch Physiker, probieren den Erfolg oder das Scheitern von Algorithmen auf naturgegebene Phänomene zurückzuführen. Über die Jahre hat sich gezeigt, dass sich die Vorhersagen der Physik für interagierende Atome auf einige Algorithmen und Datenstrukturen übertragen lassen. Diese Vorhersagen bilden oft einen Startpunkt für mathematische Forschung und einige der Vorhersagen sind inzwischen bewiesen. Obwohl diese Vorhersagen Hinweise im Bezug auf die P-vs.-NP Frage bereithalten, ist die Frage weiterhin offen und Wissenschaftler versuchen weiterhin die Beweislage zu verdichten.

In dieser Thesis tragen wir unseren Teil dazu bei, Problemstellungen zu finden, bei denen die Lösung leicht, schwer oder unmöglich ist. Wir nehmen uns ein spezielles Problem heraus und analysieren, ob und warum dieses Problem algorithmisch lösbar ist oder nicht. Das Problem unserer Wahl ist das *Group Testing Problem*
(Gruppentest-Problem). Im Jahre 1943 wurde das Problem von R. Dorfman erstmals eingeführt. Angenommen wir haben eine Gruppe von n Personen und k von ihnen sind erkrankt. Anstatt jeden einzeln zu testen, können wir Gruppen testen. Dies kann zum Beispiel durch das Vermischen von mehreren Speichel-Proben erreicht werden. Ein solcher Gruppen-Test ist positiv genau dann, wenn mindestens ein Infizierter in der Gruppe ist und negativ, wenn alle Teilnehmerinnen gesund sind. Die Frage ist nun, was die minimale Anzahl an Tests ist, die nötig ist, sodass das Problem algorithmisch

leicht, schwer oder unmöglich zu lösen ist. In dieser Arbeit betrachten wir verschiedene Abwandlungen dieses Problems. Zunächst betrachten wir das Standard-Modell, in dem alles erlaubt ist, um die Infizierten zu finden. Offensichtlich stößt dieses Modell in der Realität sehr schnell an seine Grenzen. Ein prominentes Beispiel für diese Methode ist das Virus COVID-19. Es wurde schnell klar, dass in der Realität einige Restriktionen zu beachten sind. Auf der einen Seite sind die Tests nicht fehlerfrei. Auf der anderen Seite kann man nicht unbegrenzt viele Proben in einem Test mischen. Wir haben unser Modell an diese Anforderungen angepasst und ebenfalls Resultate erarbeitet. Die verschiedenen Modelle werden wir kurz vorstellen und die Ergebnisse zusammenfassen. Natürlich kann man sich auch fragen, ob man wirklich alle Infizierten finden will oder, ob man auch mit einer approximativen Menge zufrieden ist. Dieses Kriterium des Erfolgs kann sogar noch weiter abgeschwächt werden und es stellt sich die Frage, ab wann der wahrnehmbare Einfluss der Infizierten auf die Teststruktur verschwindet. Zu den angepassten Kriterien des Erfolgs sind auch Ergebnisse in dieser Thesis enthalten. Die Zuweisung der Personen in die entsprechenden Tests kann als Graph-Struktur aufgefasst werden. In Abbildung 5 ist ein Beispiel zu finden. Nun stellt man fest, dass Infizierte und Gesunde einen unterschiedlichen Einfluss auf die zugrundeliegende Struktur haben. Es ist zu erkennen, dass es Strukturen gibt, die es sehr leicht machen, Infizierte und Gesunde direkt zu erkennen (1.1), (1.2). Außerdem erkennt man, dass es Strukturen gibt, die es unmöglich machen, Infizierte und Gesunde zu unterscheiden (1.3),(1.4). Zur Veranschaulichung kann Abbildung 6 herangezogen werden. Wir nutzen diese Beobachtungen, um unsere Resultate zum exakten Finden der infizierten Menge herzuleiten.

- Das exakte Finden der Infizierten Menge ist leicht, wenn ein Algorithmus diese eindeutig erkennt.
- Sobald es Infizierte und Gesunde gibt, die wir nicht unterscheiden können, ist das exakte Finden der infizierten Menge unmöglich.

Für Resultate zum approximativen Finden der Menge oder zum Erkennen des Einflusses der infizierten Menge ist diese Beobachtung nicht mehr ausreichend. Für Resultate zum approximativen Finden der Menge müssen wir zeigen, dass bei zwei Mengen, die das Testergebnis erklären, nur wenige infizierte Personen in beiden Mengen infiziert sind. Um den Einfluss der infizierten Menge zu erkennen, müssen wir einen Schätzer finden, der in Erwartung unterschiedliche Werte ausgibt, wenn man ihn für ein tatsächliches Gruppen Test Modell und eine zufällig erzeugte Graph-Struktur anwendet. Zusätzlich darf die Schwankung des Schätzers nicht zu groß sein, sodass die Ergebnisse voneinander getrennt bleiben und sich nicht durch Zufall überlappen können.

8.2. **Ergebnisse.**

8.2.1. *Das Standard-Modell.* Im Jahre 1943 hatte R. Dorfman die Idee, dass man durch das Testen von Gruppen statt Einzelpersonen die Kapazität an durchführbaren Tests erhöhen kann. Er schlug ein zweistufiges Verfahren vor, in dem zunächst Gruppen getestet werden. Negativ getestete Gruppen werden als negativ deklariert. Für positiv getestete Gruppen folgt ein individueller Test, um festzustellen, wer am Ende für das positive Testergebnis verantwortlich war. Nun stellt sich die Frage, was die minimale Anzahl an Tests ist, die nötig ist, um die Infizierten zu finden. Da man den Überblick verliert, solbald mehrere infizierte Mengen das selbe Testergebnis liefern, folgt mit

$$2^m \geq \binom{n}{k}$$

eine untere Schranke $m_{\text{count}} := \frac{k\ln(n/k)}{\ln(2)}$. Nun stellt sich die Frage, ob es tatsächlich möglich ist, die infizierte Menge zu finden, wenn man nur m_{count} Tests durchführt. In dieser Thesis haben wir gezeigt, dass es leicht ist, die infizierte Menge mit einem 2-stufigen

Test-Verfahren auf m_{count} Tests zu finden. Außerdem zeigen wir, dass dies unter bestimmten Bedingung sogar in einem 1-stufigen Test-Verfahren möglich ist. Wir zeigen, dass bei 1-stufigen Verfahren

$$m_{\text{inf}} = m_{\text{inf}}(n,\theta) = \max\left\{\frac{\theta}{\ln^2(2)(1-\theta)}, \frac{1}{\ln(2)}\right\} n^\theta \ln(n/k).$$

Tests ausreichen. Einerseits geben wir ein Test-Verfahren und einen Algorithmus an, der die infizierte Menge tatsächlich findet, sobald $m > m_{\text{inf}}$. Andererseits zeigen wir, dass es keine Möglichkeit gibt mit weniger Tests auszukommen. Die folgenden Theoreme fassen unsere Resultate zum exakten Finden der Menge zusammen:

Theorem 8.1 (Theorem 1.1 und 1.2 aus [30])**.** *Sei $\varepsilon > 0$ und*

$$m_{\text{inf}} = m_{\text{inf}}(n,\theta) = \max\left\{\frac{\theta}{\ln^2(2)(1-\theta)}, \frac{1}{\ln(2)}\right\} n^\theta \ln(n/k).$$

Für 1-stufige Verfahren gilt:

- *Das exakte Finden der infizierten Menge ist leicht, wenn man $m \geq (1+\varepsilon) m_{\text{inf}}$ Tests verwendet.*
- *Das exakte Finden der infizierten Menge ist unmöglich, wenn man $m \leq (1-\varepsilon) m_{\text{inf}}$ Tests verwendet.*

Theorem 8.2. *Sei $\varepsilon > 0$ und*

$$m_{count} := \frac{k\ln(n/k)}{\ln(2)}.$$

Für 2-stufige Verfahren gilt:

- *Das exakte Finden der infizierten Menge ist leicht, wenn man $m \geq (1+\varepsilon) m_{count}$ Tests verwendet.*

Der Beweis, dass es tatsächlich leicht ist, die infizierte Menge mit einem 1-stufigen Verfahren zu finden, basiert darauf, dass wir ein Testverfahren entwickeln und einen Algorithmus angeben, der die Infizierten und Gesunden anhand der vorhandenen Informationen auseinanderhalten kann. Die Idee der Struktur ist in Abbildung 8 zu finden. Außerdem zeigen wir, dass dies von keinen Testverfahren mit $m \leq (1-\varepsilon) m_{\text{inf}}$ erreicht werden kann, da viele Infizierte und Gesunde dann nicht mehr unterscheidbar sind. Für die Details des Beweises verweisen wir den Leser auf Appendix A oder [30]. Für den 2-stufigen Fall zeigen wir, dass unser Algorithmus bei unserem Testverfahren mit lediglich m_{count} Tests nur sehr wenige Individuen nicht sicher erkennt. Diese können in einer zweiten Runde durch individuelles Testen dann nachträglich klassifiziert werden, ohne die Gesamtanzahl der Tests zu beeinflussen.

In einem nächsten Schritt analysieren wir die abgeschwächten Versionen des Erfolgs, das approximative Finden und das Erkennen des Einflusses der infizierten Menge.

Theorem 8.3 (Theorem 1.2 aus [30] und Theorem 1 aus [31])**.** *Sei $\varepsilon > 0$ und*

$$m_{count} = \frac{k\ln(n/k)}{\ln(2)}. \tag{8.1}$$

Für 1-stufige Verfahren gilt:

- *Das approximative Finden der infizierten Menge ist leicht, sobald man $m \geq (1+\varepsilon) m_{count}$ Tests verwendet.*
- *Das approximative Finden der infizierten Menge ist informationstheoretisch möglich für das Konstante-Spalten Test Verfahren, sobald man $m \geq (1+\varepsilon) m_{count}$ Tests verwendet.*
- *Das approximative Finden der infizierten Menge ist unmöglich für das Konstante-Spalten Test Verfahren, sobald man $m \leq (1-\varepsilon) m_{count}$ Tests verwendet.*

Im Konstante-Spalten Testverfahren wählt jedes Individuum exakt Δ Tests aus. Die ersten beiden Teile des Theorems folgen direkt aus unseren Resultaten zum exakten Finden der Menge (da unser optimales Testverfahren mit m_{count} approximativ funktioniert und ein Spezialfall des Konstante-Spalten Verfahrens ist). Für den dritten Teil führen wir eine Moment Berechnung durch und zeigen, dass es mit hoher Wahrscheinlichkeit keine Lösungen mit konstanter Überlappung gibt.

Nun wollen wir noch sehen, ab wann man den Einfluss der infizierten Mengen (effizient) erkennen kann.

Theorem 8.4 (Theorem 2 aus [31])**.** *Im Konstante-Spalten Verfahren mit $\theta \in (0,1)$ und $c > 0$ sei*

$$m_{detect} = \max\left\{\frac{1}{\ln^2 2}\left(1 - \frac{\theta}{2(1-\theta)}\right), 0\right\} k \cdot \ln(n/k). \tag{8.2}$$

Außerdem definieren wir $m^{inf}_{detect} = \min(m_{detect}, m_{count})$. Dann gilt das Folgende:

(a) *Mit $m > m_{detect} > 0$ ist es leicht den Einfluss der infizierten Menge zu erkennen.*
(b) *Mit $m > m^{inf}_{detect}$ ist informationstheoretisch möglich den Einfluss der infizierten Menge zu erkennen.*

Zum Beweis verwenden wir den Schätzer

$$\mathbb{V}(\Gamma_1, \ldots, \Gamma_M) = \sum_{j=1}^{M}\left(\Gamma_j - \frac{N\Delta}{M}\right)^2. \tag{8.3}$$

und zeigen, dass der Abstand der Erwartungswerte für das Gruppen-Test Modell und das zufällige Modell weit genug auseinander liegen, um mit der resultierenden Standardabweichung umgehen zu können, wenn die Anzahl an Tests groß genug (größer als m_{detect}) ist. Außerdem zeigen wir für Teil b) des Theorems, dass ein Algorithmus zum approximativen Finden der Menge zum Erkennen des Einflusses verwendet werden kann.

8.2.2. *Das Größen-Beschränkte Modell.* Im Standard-Modell haben wir gezeigt, wo die Grenzen für das Gruppentest-Problem liegen, wenn alles erlaubt ist und keine weiteren Einschränkungen vorgegeben sind. Es ist offensichtlich, dass diese Freiheit bei realen Anwendungen nicht vorliegen wird. Für einige Viren, wie zum Beispiel HIV [105] oder COVID-19 [79], ist bekannt, dass die Funktionsfähigkeit der Tests nur bei Anwendung auf begrenzt viele Individuen gewährleistet ist. In diesem Abschnitt gehen wir darauf ein, wie sich diese zusätzliche Einschränkung auf die Grenzen des Gruppentest-Problems auswirken. Wir nehmen wieder an, dass sich in der Population der Größe n eine infizierte Untergruppe der Größe $k \sim n^\theta$ befindet. Erste Schritte in dieser Hinsicht wurden bereits durch [49] unternommen. Es wurden bereits erste obere und untere Schranken für den Erfolg und das Scheitern im Größen-Beschränkten Modell analysiert. In [49] wird angenommen, dass ein Individuum nur an $\Delta = o(\ln(n))$ Tests teilnehmen darf und jeder Test nur maximal $o\left(n^{1-\theta}\right)$. Die Schranken von [49] lassen sich wie folgt zusammenfassen:

- Das Δ-Beschränkte Modell:
  - Für $\Delta = o(\ln n)$, jedes Test-Verfahren (mit parallelen Tests) und Fehler-Wahrscheinlichkeit ξ benötigt $m \geq \Delta k\left(\frac{n}{k}\right)^{\frac{1-5\xi}{\Delta}}$, für ausreichend kleines ξ und ausreichend großes n . (Theorem 4.1 in [49])
  - Für ein angemessen gewähltes zufälliges Test-Verfahren mit $m \geq \lceil e\Delta k\left(\frac{n}{\xi}\right)^{\frac{1}{\Delta}}\rceil$ Tests hat der COMP Algorithmus eine Fehler-Wahrscheinlichkeit von höchsten ξ. (Theorem 4.2 in [49])

- Das Γ-Beschränkte Modell:
  - Für $\Gamma = \Theta\left(\left(\frac{n}{k}\right)^{\beta}\right)$ mit $\beta \in [0,1)$, jedes Testverfahren (mit parallelen Tests) mit Fehler-Wahrscheinlichkeit höchstens ξ benötigt $m \geq \frac{1-6\xi}{1-\beta} \cdot \frac{n}{\Gamma}$, für ausreichend großes n. (Theorem 4.5 in [49])
  - Für ein entsprechend gewähltes zufälliges Test-Verfahren und den COMP Algorithmus, mit $\Gamma = \Theta\left(\left(\frac{n}{k}\right)^{\beta}\right)$ für $\beta \in [0,1)$ und $\xi = n^{-\zeta}$ mit $\zeta > 0$, dann ist die Fehler-Wahrscheinlichkeit höchstens ξ wenn $m \geq \lceil \frac{1+\zeta}{(1-\theta)(1-\beta)} \rceil \cdot \lceil \frac{n}{\Gamma} \rceil$. (Theorem 4.6 in [49])

In unserer Arbeit verbessern wir für einige dieser Modelle sowohl die oberen als auch die unteren Schranken.

Theorem 8.5 (Δ-Beschränktes-Modell, Theorem 3.1, 3.2 und 3.3 in [51])**.** *Sei* $\varepsilon > 0$, $\theta \in (0,1)$ *und definiere*

$$m_{\Delta,inf} = \max\left\{ e^{-1} \Delta k^{1+\frac{1-\theta}{\Delta\theta}}, \Delta k^{1+\frac{1}{\Delta}} \right\}$$

$$\textit{und} \qquad m_{\Delta,alg} = \max\left\{ \Delta k^{1+\lfloor \frac{\theta}{1-\theta} \rfloor}, \Delta k^{1+\frac{1}{\Delta}} \right\}.$$

Im Δ*-Beschränkten-Modell hält das Folgende:*

- *Das exakte Finden der infizierten Menge ist leicht, sobald* $m \geq (1+\varepsilon) m_{\Delta,alg}$.
- *Das exakte Finden der infizierten Menge ist unmöglich, sobald* $m \leq (1-\varepsilon) m_{\Delta,inf}$.

Wir halten fest, dass unsere Schranken fast übereinstimmen. Obwohl die Ordung für alle $\theta \in (0,1)$ übereinstimmt, sehen wir eine Konstante e^{-1} Lücke für $\theta < 1/2$.

Theorem 8.6 (Γ-Beschränktes-Modell, Theorem 4.1, 4.10 und 4.18 in [51])**.** *Sei* $\varepsilon > 0$, $\theta \in (0,1)$ *und definiere*

$$m_{\Gamma} = \max\left\{ \left(1 + \left\lfloor \frac{\theta}{1-\theta} \right\rfloor\right) \frac{n}{\Gamma}, \frac{2n}{\Gamma+1} \right\}.$$

Im Γ*-Beschränkten-Modell hält das folgende:*

- *Das exakte Finden der infizierten Menge ist leicht, sobald* $m \geq (1+\varepsilon) m_{\Gamma}$.
- *Das exakte Finden der infizierten Menge ist, unmöglich sobald* $m \leq (1-\varepsilon) m_{\Gamma}$.

Wir sehen, dass unsere Schranken übereinstimmen und wir einen direkten Übergang von 'leicht zu lösen' zu 'unmöglich zu lösen' finden.

Wir werden hier nur die Beweis-Idee vorstellen und verweisen den interessierten Leser zu Appendix C oder [51]. Um zu zeigen, dass das Finden der infizierten Menge leicht bzw. unmöglich ist, analysieren wir die Nachbarschaften der Individuen. Es ist unmöglich, die infizierte Menge zu finden, wenn es sowohl infizierte als auch gesunde Individuen gibt, die ununterscheidbar sind. In diesem Fall können wir den jeweiligen Status der Individuen austauschen und würden die Änderung nicht bemerken. In beiden Modellen zeigen wir, dass diese verbotenen Strukturen in allen Test-Verfahren auftreten sobald die Anzahl an Tests zu gering wird. Um zu zeigen, dass das Finden der Menge leicht ist, benötigen wir ein Test-Verfahren und einen Algorithmus, um die infizierte Menge finden. Wir geben in beiden Modellen ein entsprechendes Test-Verfahren an und zeigen, dass der so genannte DD-Algorithmus ausreicht, um alle Infizierten zu finden, sobald wir die Anzahl an Tests groß genug wählen. Eine Veranschaulichung der Schranken kann in Abbildung 9 gefunden werden.

8.2.3. *Das Fehlerhafte-Test Modell.* Eine weitere Annahme des Standard-Modells, die in der Realität zu Problemen führen kann, sind die fehlerfreien Tests.
Am Beispiel COVID-19 [106] sieht man, dass man diese Annahme nicht in der Realität finden wird. Deshalb passen wir unser Modell in diesem Abschnitt an und gehen

davon aus, dass jeder Test mit einer bestimmten Wahrscheinlichkeit sein Testergebnis ändert. Zur Verdeutlichung der Idee kann Abbildung 10 herangezogen werden. Dieses Modell hat bereits einige Aufmerksamkeit genossen [22, 87, 88, 89, 90, 92]. Wir sind nun daran interessiert, inwieweit wir Algorithmen und Testverfahren angeben können, die die infizierte Menge finden, obwohl die Tests nicht fehlerfrei sind. Hierzu schauen wir uns das Konstante-Spalten-Verfahren an (jedes Individuum zieht sich Δ Test uniform zufällig aus den m verfügbaren Tests mit Zurücklegen). Nun analysieren wir die Erfolgschancen von Algorithmus 2 und Algorithmus 3. Wir erhalten die folgenden Schranken für Algorithmus2:

Theorem 8.7 (Fehlerbehafteter COMP,Theorem 2.1 in [52])**.** *Seien* $p,q \geq 0$, $p+q<1, d \in (0,\infty), \alpha \in \left(q, e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)$. *Angenommen* $0<\theta<1$ *und sei*

$$m_{COMP} = m_{COMP}(n,\theta,p,q) = \min_{\alpha,d} \max\{b_1(\alpha,d), b_2(\alpha,d)\}\, k\ln(n/k)$$

$$\begin{aligned} &mit \qquad b_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|q\right)} \\ &und \qquad b_2(\alpha,d) = \frac{1}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}. \end{aligned}$$

Für $m \geq (1+\varepsilon)m_{COMP}$ *und ein* $\varepsilon>0$ *ist das exakte Finden der infizierten Menge mit dem fehlerbehafteten* COMP *Algorithmus leicht.*

Für Algorithmus 3 erhalten wir:

Theorem 8.8 (Fehlerbehafteter DD,Theorem 2.2 in [52])**.** *Seien* $p,q \geq 0$, $p+q<1, d \in (0,\infty), \alpha \in \left(q, e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)$ *und* $\beta \in (0, e^{-d}(1-q))$ *und definiere* $w = e^{-d}p + (1-e^{-d})(1-q)$. *Angenommen* $0<\theta<1$ *und sei*

$$m_{DD} = m_{DD}(n,\theta,p,q) = \min_{\alpha,\beta,d} \max\left\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\alpha,\beta,d)\right\} k\ln(n/k)$$

$$\begin{aligned} &mit \qquad c_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|q\right)} \\ &und \qquad c_2(\alpha,d) = \frac{1}{dD_{\mathrm{KL}}\left(\alpha\|1-w\right)} \\ &und \qquad c_3(\beta,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\beta\|(1-q)e^{-d}\right)} \\ &und \qquad c_4(\alpha,\beta,d) = \max_{1-\alpha\leq z\leq 1}\left\{\frac{1}{1-\theta}\frac{1}{d\left(D_{\mathrm{KL}}\left(z\|w\right)+\mathbf{1}\left\{\beta>\frac{ze^{-d}p}{w}\right\}zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right)}\right\}. \end{aligned}$$

Für $m \geq (1+\varepsilon)m_{DD}$ *und* $\varepsilon>0$ *ist das Finden der infizierten Menge mit dem fehlerbehafteten* DD *Algorithmus leicht.*

Da diese Schranken auf den ersten Blick kompliziert wirken, kann der interessierte Leser in Abbildung 4 die Realisierung der Schranken für verschiedene Fehler-Wahrscheinlichkeiten einsehen.

Für detaillierte Beweise verweisen wir den Leser auf Appendix D oder [52]. Wir gehen jedoch kurz auf die Beweisidee ein. Wir verwenden das Konstante-Spalten Verfahren und analysieren, wie sich die Fehler-Wahrscheinlichkeiten auf die Nachbarschaften von infizierten und gesunden Individuen auswirken. Man erkennt, dass die beiden Arten von Individuen unterschiedliche Eigenschaften in der zugrundeliegenden Graph-Struktur aufweisen. Dies liegt daran, dass der Unterschied bereits im fehlerfreien Modell auftritt und sich in einer bestimmten Art und Weise auf das fehlerbehaftete Modell übertragen lässt. Sowohl Algorithmus 2 als auch Algorithmus 3 nutzen einige

dieser unterschiedlichen Strukturen aus. Unser Beweis basiert im Grunde darauf zu beweisen, dass sich die entsprechenden Strukturen für Infizierte und Gesunde genug unterscheiden, um sie mit den entsprechenden Algorithmen zu nutzen. Dies ist möglich, sobald die Anzahl an Tests groß genug gewählt wird. Wir berechnen die benötigten Anzahlen und sie führen zu den Schranken, die in den Theoremen angegeben sind.

8.2.4. *Das Quantitative Modell.* Eine weitere Abwandlung des Modells beschäftigt sich mit der Art der Ausgabe der Testergebnisse. Im Standard-Modell wird eine binäre Ausgabe angenommen. Einige Anwendungen von Gruppen Test Verfahren [14, 20, 68, 94] arbeiten jedoch mit nicht binären Ausgaben. Deshalb passen wir das Modell entsprechend an und beschäftigen uns mit den Schranken, sodass das exakte Finden der infizierten Menge möglich ist. Außerdem wollen wir wissen, wann genügend Information durch das Testverfahren bereit gestellt wird, um das exakte Finden der Menge zu ermöglichen. Wir nehmen wieder an, dass sich in einer Population von n Individuen eine Menge von $k \sim n^\theta$ Kranken befindet. Die Tests hingegen ändern jetzt ihre Ausgabe. Anstatt einen binären Ausgang anzusetzen, gibt ein Test die Anzahl der im Test enthaltenen Infizierten aus. In diesem Modell kann wieder eine simple Zähl-Schranke

$$m_{\text{count}}^{\text{quant}} > \frac{\ln\left(\frac{n}{k}\right)}{\ln(k)} k$$

angegeben werden. Es wurde in [39] gezeigt, dass das exakte Finden der infizierten Menge unmöglich ist, sobald die Anzahl der Tests unter $2m_{\text{count}}^{\text{quant}}$ fällt. Die Frage ist nun, ob es mit $2m_{\text{count}}^{\text{quant}}$ Tests überhaupt möglich ist, die infizierte Menge zu finden. Die beste Schranke in dieser Hinsicht wurde von [54] mit $m_{\text{GeKu}} = 4m_{\text{count}}^{\text{quant}}$ angegeben. Wir schließen diese Lücke und zeigen, dass $2m_{\text{count}}^{\text{quant}}$ Tests ausreichen. Die nächste Frage ist, ob ein Algorithmus existiert, der diese Informationen nutzen kann und die infizierte Menge findet. Es gibt bereits verschiedene Algorithmen, die für die Lösung des Problems vorgestellt wurden wie z.B. [40, 42, 44, 61, 62, 85]. Alle diese Algorithmem benötigen $m = \Theta(k\ln(n))$ Tests.

Theorem 8.9 (Theorem 1 und Theorem 2 in [50])**.** *Sei* $0 < \theta < 1$, $k = n^\theta$ *und* $\varepsilon > 0$ *und sei*

$$m_{pos}(n,\theta) = 2\frac{k\ln(n/k)}{\ln(k)} = 2\frac{1-\theta}{\theta}k,$$

$$m_{Easy}(n,\theta) = 4\left(1-\frac{1}{\sqrt{e}}\right)\frac{1+\sqrt{\theta}}{1-\sqrt{\theta}}k\ln(n/k).$$

Das exakte Finden der infizierten Menge ist möglich mit $m > (1+\varepsilon)m_{pos}$ *und leicht mit* $m > (1+\varepsilon)m_{Easy}(n,\theta)$ *Tests.*

Wir halten fest, dass unsere Schranken eine Lücke der Ordnung $\ln(k)$ aufweist. Diese Lücke konnte im Nachhinein von [55] geschlossen werden. Für die Details der Beweise unserer Schranken verweisen wir den Leser zu Appendix E oder [50]. Wir geben wieder die Beweisidee an. Auf der einen Seite zeigen wir, dass es für ein richtig gewähltes einstufiges Test-Verfahren mit $m > m_{\text{pos}}$ genau eine Konfiguration mit k Kranken gibt, die das Testergebnis erklärt. Für diese Aussage zeigen wir, dass es weder eine weitere Konfiguration mit großer noch eine weitere mit kleiner Überlappung gibt. Wir analysieren die Möglichkeit eine weitere erklärende Konfiguration mit bestimmter Überlappung zu erzeugen. Sobald wir die Anzahl der Tests groß genug wählen, ist dies mit hoher Wahrscheinlichkeit nicht mehr möglich. Auf der anderen Seite geben wir einen Algorithmus an, der die infizierte Menge findet. Für diesen Schritt sehen wir ein, dass sich der Einfluss der infizierten Menge auf das Testverfahren vom Einfluss der gesunden Menge unterscheidet. Die Eigenschaft, die wir für jedes Individuum nutzen, ist die Summe der Testergebnisse in der direkten Nachbarschaft der Individuen. Hier zeigen

wir nun, dass diese Summen gut separiert sind für die beiden Gruppen. Von daher ist der Algorithmus erfolgreich, sobald die Anzahl an Tests groß genug ($m > m_{\text{Easy}}$) ist.

8.2.5. *Fazit.* Die Covid-19 Krise hat auf eindrucksvolle Art und Weise gezeigt, warum das grundsätzliches Verständnis des Gruppentest-Problems von großer Bedeutung ist. In dieser Arbeit wurden einige der dringendsten offene Frage der Gruppentest-Forschung für das Standard-Modell beantwortet [7]. Außerdem haben wir Fortschritte darin gemacht, diese Resultate auf Modelle zu übertragen, die näher an der Realität sind als das Standard-Modell.
Wir haben Strukturen herausgearbeitet, die für den Erfolg und das Scheitern von Algorithmen verantwortlich sind. Wir glauben, dass diese Einblicke auch in anderen verwandten Problemen genutzt werden können. Es stellen sich natürlich einige Fragen, die in weiterer Forschung beleuchtet werden sollten:

- Kann der Spatial-Coupling Ansatz auf das fehlerhafte Testmodell übertragen werden?
- Können die Lücken zwischen leicht vs. möglich vs. hart im Konstante-Spalten Test-Verfahren geschlossen werden?
- Wie übertragen sich unsere Ergebnisse auf eine linearen Infektionsdichte $k \sim \Theta(n)$?
- Lassen sich die verbleibenden Lücken zwischen 'genug Information vorhanden' und 'Information algorithmisch leicht nutzbar' für das Gruppen Test Modell schließen?
- Können die beobachteten Vorteile durch mehrstufige Verfahren im Vergleich zu einstufigen Verfahren auf andere algorithmische Aufgaben übertragen werden?

Wir gehen davon aus, dass die Einblicke, die wir durch unsere Ergebnisse im Gruppentest-Problem erhalten haben, hilfreich sein werden.

References

[1] E. Abbe. Community detection and stochastic block models: recent developments. *Journal of Machine Learning Research*, 18:177:1–177:86, 2017.

[2] D. Achlioptas and A. Coja-Oghlan. Algorithmic barriers from phase transitions. *Proceedings of 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS'08)*, page 793–802, 2008.

[3] D. Achlioptas, A. Coja-Oghlan, and F. Ricci-Tersenghi. On the solution space geometry of random formulas. *Random Structures and Algorithms*, 38:251–268, 2011.

[4] D. Achlioptas, A. Naor, and Y. Peres. Rigorous location of phase transitions in hard optimization problems. *Nature*, 435:759–764, 2005.

[5] A. Alaoui, A. Ramdas, F. Krzakala, L. Zdeborová, and M. Jordan. Decoding from pooled data: Phase transitions of message passing. *IEEE Transactions on Information Theory*, 65(1):572–585, 2019.

[6] M. Aldridge, L. Baldassini, and O. Johnson. Group testing algorithms: bounds and simulations. *IEEE Transactions on Information Theory*, 60:3671–3687, 2014.

[7] M. Aldridge, O. Johnson, and J. Scarlett. Group testing: an information theory perspective. *Foundations and Trends in Communications and Information Theory*, 15(3–4):196–392, 2019.

[8] A. Allemann. An efficient algorithm for combinatorial group testing. *Information Theory and Combinatorics and Search Theory. Lecture Notes in Computer Science*, 7777:569–596, 2013.

[9] N. Alon, M. Krivelevich, and B. Sudakov. Finding a large hidden clique in a random graph. *Random Structures and Algorithms*, 13(3-4):457–466, 1998.

[10] D. Applegate, R. Bixby, V. Chvátal, and W. Cook. The traveling salesman problem: A computational study. *Princeton University Press*, 2011.

[11] A. Auffinger, W. Chen, and Q. Zeng. The sk model is infinite step replica symmetry breaking at zero temperature. *Communication on pure and applied mathematics*, 73:921–943, 2020.

[12] L. Baldassini, O. Johnson, and M. Aldridge. The capacity of adaptive group testing. *Proceedings of 2013 IEEE International Symposium on Information Theory (ISIT'13)*, pages 2676–2680, 2013.

[13] A. Bandeira, J. Banks, D. Kunisky, C. Moore, and A. Wein. Spectral planting and the hardness of refuting cuts, colorability, and communities in random graphs. *Proceedings of the 34th Annual Conference on Learning Theory (COLT'21)*, pages 1–64, 2021.

[14] R. Ben-Ami and A. Klochendler et.al. Large-scale implementation of pooled rna extraction and rt-pcr for sars-cov-2 detection. *Clinical Microbiology and Infection*, 26(9):1248–1253, 2020.

[15] G. Ben-Arous, A. Wein, and I. Zadik. Free energy wells and overlap gap property in sparse pca. *Proceedings of the 33rd Annual Conference on Learning Theory (COLT'20)*, pages 479–482, 2020.

[16] R. Benz, S. Swamidass, and P. Baldi. Discovery of power-laws in chemical space. *Journal of Chemical Information and Modeling*, 48:1138–1151, 2008.

[17] E. Bolthausen. An iterative construction of solutions of the tap equations for the sherrington–kirkpatrick model. *Communications in Mathematical Physics*, 325:333—-366, 2014.

[18] M. Brennan, G. Bresler, J. Li, and T. Schramm. Statistical query algorithms and low-degree tests are almost equivalent. *Proceedings of 34th Annual Conference on Learning Theory (COLT'21)*, 2021.

[19] G. Bresler and B. Huang. The algorithmic phase transition of random k-sat for low degree polynomials. *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS'21)*, 2021.

[20] C. Cao, C. Li, and X. Sun. Quantitative group testing-based overlapping pool sequencing to identify rare variant carriers. *BMC Bioinformatics*, 15:195, 2014.

[21] T. Castellani and A. Cavagna. Spin-glass theory for pedestrians. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(5):P05012, 2005.

[22] C. Chan, P. Che, S. Jaggi, and V. Saligrama. Non-adaptive probabilistic group testing with noisy measurements: near-optimal bounds with efficient algorithms. *Proceedings of 49th Annual Allerton Conference on Communication, Control, and Computing*, 1:1832–1839, 2011.

[23] W. Chen. On the almeida-thouless transition line in the sherrington-kirkpatrick model with centered gaussian external field. *Electronic Communications in Probability*, 26.65, 2021.

[24] I. Cheong. The experience of south korea with covid-19. *Mitigating the COVID Economic Crisis: Act Fast and Do Whatever It Takes (CEPR Press)*, pages 113–120, 2020.

[25] A. Coja-Oghlan. A better algorithm for random k-sat. *SIAM Journal on Computing*, 39(7):2823–2864, 2010.

[26] A. Coja-Oghlan. Belief propagation guided decimation fails on random formulas. *Journal of the ACM*, 63:1–55, 2017.

[27] A. Coja-Oghlan and C. Efthymiou. On independent sets in random graphs. *Random Structures and Algorithms*, 47(3):436–486, 2014.

[28] A. Coja-Oghlan, C. Efthymiou, N. Jaafari, M. Kang, and T. Kapetanopoulos. Charting the replica symmetric phase. *Communications in Mathematical Physics*, 359(2):603–698, 2018.

[29] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick. Information-theoretic and algorithmic thresholds for group testing. *IEEE Transactions on Information Theory*, 66(12):7911–7928, 2020.

[30] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick. Optimal group testing. *Combinatorics, Probability and Computing*, pages 1–38, 2020.

[31] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, A. Wein, and I. Zadik. Statistical and computational phase transitions in group testing. *Under Review*, 2022.

[32] A. Coja-Oghlan, A. Haqshenas, and S. Hetterich. Walksat stalls well below satisfiability. *SIAM Journal on Discrete Mathematics*, 31:1160–1173, 2017.

[33] A. Coja-Oghlan, T. Kapetanopoulos, and N. Müller. The replica symmetric phase of random constraint satisfaction problems. *Combinatorics, Probability and Computing*, 29(3):346–422, 2019.

[34] A. Coja-Oghlan, F. Krzakala, W. Perkins, and L. Zdeborová. Information-theoretic thresholds from the cavity method. *Advances in Mathematics*, 333:694–795, 2018.

[35] A. Coja-Oghlan and K. Panagiotou. The asymptotic k-sat threshold. *Advances in Mathematics*, 288:985–1068, 2016.

[36] J. de Almeida and D. Thouless. Stability of the sherrington-kirkpatrick model of spin glasses. *Journal of Physics A: Mathematical and General*, pages 983–990, 1978.

[37] J. Ding, A. Sly, and N. Sun. Proof of the satisfiability conjecture for large k. *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC'15)*, page 59–68, 2015.

[38] Y. Ding, D. Kunisky, A. Wein, and A. Bandeira. The average-case time complexity of certifying the restricted isometry property. *IEEE Transactions on Information Theory*, 67:7355–7361, 2021.

[39] A. G. Djackov. On a search model of false coins. *Topics in Information Theory. Hungarian Academy of Science*, 16:163–170, 1975.

[40] D. Donoho and J. Tanner. Thresholds for the recovery of sparse solutions via l1 minimisation. *Proceedings of the 40th Annual Conference on Information Sciences and Systems (CISS'06)*, pages 202–206, 2006.

[41] R. Dorfman. The detection of defective members of large populations. *Annals of Mathematical Statistics*, 14:436–440, 1943.

[42] U. Feige and A. Lellouche. Quantitative group testing and the rank of random matrices. *ArXiv-Preprint*, 2020.

[43] C. Fortuin, P. Kasteleyn, and J. Ginibre. Correlation inequalities on some partially ordered sets. *Communications in Mathematical Physics*, 22(2):89–103, 1971.

[44] S. Foucart and H. Rauhut. An invitation to compressive sensing. *Applied and Numerical Harmonic Analysis*, pages 1–39, 2013.

[45] M. Gabrié, V. Dani, G. Semerjian, and L. Zdeborová. Phase transitions in the q-coloring of random hypergraphs. *Journal of Physics A: Mathematical and Theoretical*, 50(50):505002, 2017.

[46] D. Gamarnik. The overlap gap property: A topological barrier to optimizing over random structures. *Perspective, National Academy of Sciences*, 2021.

[47] D. Gamarnik, A. Jagannath, and A. Wein. Low-degree hardness of random optimization problems. *Proceedings of 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS'20)*, 2020.

[48] D. Gamarnik and I. Zadik. The landscape of the planted clique problem: Dense subgraphs and the overlap gap property. *ArXiv-Preprint*, 2019.

[49] V. Gandikota, E. Grigorescu, S. Jaggi, and S. Zhou. Nearly optimal sparse group testing. *IEEE Transactions on Information Theory*, 65(5):2760–2773, 2019.

[50] O. Gebhard, M. Hahn-Klimroth, P. Loick, and D. Kaaser. On the parallel reconstruction from pooled data. *to appear in proceedings of 36th IEEE International Parallel and Distributed Processing Symposium (IPDPS'22)*, 2022.

[51] O. Gebhard, M. Hahn-Klimroth, O. Parczyk, M. Penschuck, M. Rolvien, J. Scarlett, and N. Tan. Near optimal sparsity-constrained group testing: improved bounds and algorithms. *to appear in IEEE Transactions on Information Theory*, 2021.

[52] O. Gebhard, O. Johnson, P. Loick, and M. Rolvien. Improved bounds for noisy group testing with constant tests per item. *IEEE Transactions on Information Theory*, 68(4):2604–2621, 2020.

[53] E. Gould. Methods for long-term virus preservation. *Mol Biotechnol*, 13:57–66, 1999.

[54] V. Grebinski and G. Kucherov. Optimal reconstruction of graphs under the additive model. *Algorithmica*, 28(1):104–124, 2000.

[55] M. Hahn-Klimroth and N. Müller. Near optimal efficient decoding from pooled data. *ArXiv-Preprint*, 2022.

[56] S. Hetterich. Analysing survey propagation guided decimation on random formulas. *Proceedings of the 43rd International Colloquium on Automata, Languages and Programming (ICALP'16)*, 2016.

[57] S. Homer and A. Selman. Computability and complexity theory. *Springer Publishing Company, Incorporated*, 2011.

[58] F. Hwang. A method for detecting all defective members in a population by group testing. *Journal of the American Statistical Association*, 67(339):605–608, 1972.

[59] A. Jagannath and I. Tobasco. Some properties of the phase diagram for mixed p-spin glasses. *Probability Theory and Related Fields*, 167(3–4):615–672, 2016.

[60] I. Jolliffe and J. Cadima. Principal component analysis: a review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374, 2016.

[61] E. Karimi, F. Kazemi, A. Heidarzadeh, K. Narayanan, and A. Sprintson. Non-adaptive quantitative group testing using irregular sparse graph codes. *Proceedings of 58th Annual Allerton Conference on Communication, Control, and Computing*, pages 608–614, 2019.

[62] E. Karimi, F. Kazemi, A. Heidarzadeh, K. Narayanan, and A. Sprintson. Sparse graph codes for non-adaptive quantitative group testing. *IEEE Information Theory Workshop (ITW)*, pages 1–5, 2019.

[63] R. Karp. Reducibility among combinatorial problems. *Complexity of Computer Computations*, 1972.

[64] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proceedings of the National Academy of Sciences*, 104(25):10318–10323, 2007.

[65] S. Kudekar, T. Richardson, and R. Urbanke. Threshold saturation via spatial coupling: Why convolutional LDPC ensembles perform so well over the BEC. *IEEE Transactions on Information Theory*, 57(2):803–834, 2011.

[66] D. Kunisky, A. Wein, and A. Bandeira. Notes on computational hardness of hypothesis testing: Predictions using the low-degree likelihood ratio. *ArXiv-Preprint*, 2019.

[67] H. Kwang-Ming and D. Ding-Zhu. Pooling designs and nonadaptive group testing: important tools for DNA sequencing. *World Scientific*, 2006.

[68] W. Liang and J. Zou. Neural group testing to accelerate deep learning. In *Proceedings of 2021 IEEE International Symposium on Information Theory (ISIT'21)*, pages 958–963, 2021.

[69] S. Long, C. Prober, and M. Fischer. Principles and practice of pediatric infectious diseases. *Elsevier*, 2018.

[70] G. Parisi M. Mézard. A replica analysis of the travelling salesman problem. *Journal de Physique*, 47(8):1285–1296, 1986.

[71] N. Madhav, B. Oppenheim, M. Gallivan, P. Mulembakani, E. Rubin, and N. Wolfe. Pandemics: Risks, impacts and mitigation. *The World Bank:Disease control priorities*, 9:315–345, 2017.

[72] C. McMahan, J. Tebbs, and C. Bilder. Informative dorfman screening. *Journal of the International Biometric Society*, 68:287–296, 2012.

[73] M. Mézard, G.Parisi, and R.Zecchina. "analytic and algorithmic solution of random satisfiability problem. *Science*, 297:812, 2002.

[74] M. Mézard and A. Montanari. *Information, Physics, and Computation*. Oxford University Press, 2009.

[75] M. Mézard and G. Parisi. The cavity method at zero temperature. *Journal of Statistical Physics*, 111:1–34, 2003.

[76] M. Mézard, G. Parisi, M. A. Virasoro, and D. Thouless. Spin glass theory and beyond. *World Scientific*, 1987.

[77] M. Mézard and R. Zecchina. Randomk-satisfiability problem: From an analytic solution to an efficient algorithm. *Physical Review E*, 66(5), 2002.

[78] R. Mourad, Z. Dawy, and F. Morcos. Designing pooling systems for noisy high-throughput protein-protein interaction experiments using boolean compressed sensing. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 10:1478–1490, 2013.

[79] L. Mutesa and P. et al. Ndishimye. A strategy for finding people infected with SARS-CoV-2: optimizing pooled testing at low prevalence. *Nature*, 589:276–280, 2021.

[80] H. Ngo and D. Du. A survey on combinatorial group testing algorithms with applications to dna library screening. *Discrete Mathematical Problems with Medical Applications*, 7:171–182, 2000.

[81] H. Nishimori. Exact results and critical properties of the ising model with competing interactions. *Journal of Physics C: Solid State Physics*, 13(21), 1980.

[82] H. Nishimori. Internal energy, specific heat and correlation function of the bond-random ising model. *Progress of Theoretical Physics*, 66(4):1169–1181, 1981.

[83] D. Panchenko. The sherrington-kirkpatrick model: An overview. *Journal of Statistical Physics*, 149(2):362—383, 2012.

[84] D. Panchenko. The parisi ultrametricity conjecture. *Annals of Mathematics*, 177:383–393, 2013.

[85] Y. Pati, R. Rezaiifar, and P. Krishnaprasad. Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. *Proceedings of 27th Asilomar Conference on Signals, Systems and Computers (ACSSC'93)*, 1:40–44, 1993.

[86] R. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21:120–126, 1978.

[87] J. Scarlett. Noisy adaptive group testing: Bounds and algorithms. *IEEE Transactions on Information Theory*, 65:3646–3661, 2018.

[88] J. Scarlett. An efficient algorithm for capacity-approaching noisy adaptive group testing. *Proceedings of 2019 IEEE International Symposium on Information Theory (ISIT'19)*, pages 2679–2683, 2019.

[89] J. Scarlett and V. Cevher. Phase transitions in group testing. *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms(SODA'16)*, 1:40–53, 2016.

[90] J. Scarlett and V. Cevher. Near-optimal noisy group testing via separate decoding of items. *IEEE Journal of Selected Topics in Signal Processing*, 12(5):902–915, 2017.

[91] J. Scarlett and V. Cevher. Phase transitions in the pooled data problem. *Proceedings of the 30th Conference on Neural Information Processing Systems (NeurIPS'17)*, pages 376–384, 2017.

[92] J. Scarlett and O. Johnson. Noisy non-adaptive group testing: A (near-)definite defectives, approach. *IEEE Transactions on Information Theory*, 66(6):3775–3797, 2020.

[93] T. Schramm and A. Wein. Computational barriers to estimation from low-degree polynomials. *to appear in Annals of Statistics*, 2020.

[94] P. Sham, J. S. Bader, and I. Craig et al. Dna pooling: a tool for large-scale association studies. *Nature Reviews Genetics*, 3:862–871, 2002.

[95] H. S. Shapiro. Problem e 1399. *Amer. Math. Monthly*, 67:82, 1960.

[96] M. Sipser. Introduction to the theory of computation. *PWS Publishing Company*, 1996.

[97] M. Talagrand. Mean-field models for spin glasses. *Springer,Series of Modern Surveys in Mathematics*, 54, 2011.

[98] N. Thierry-Mieg. A new pooling strategy for high-throughput screening: the shifted transversal design. *BMC Bioinformatics*, 7:28, 2006.

[99] D. Thouless, P. Anderson, and R. Palmer. Solution of solvable model in spin glasses. *Philosophical Magazin*, 35:593–601, 1977.

[100] F. Toninelli. About the almeida-thouless transition line in the sherrington-kirkpatrick mean-field spin glass model. *Europhysics Letters*, 60:764–767, 2020.

[101] L. Truong, M. Aldridge, and J. Scarlett. On the all-or-nothing behavior of bernoulli group testing. *IEEE Journal on Selected Areas in Information Theory*, 1(3):669–680, 2020.

[102] Unknown. Der handlungsreisende wie er sein soll und was er zu thun hat, um aufträge zu erhalten und eines glücklichen erfolgs in seinen geschäften gewiß zu sein. *Unknown*, 1832.

[103] L. Wang, X. Li, Y. Zhang, and K. Zhang. Evolution of scaling emergence in large-scale spatial epidemic spreading. *Public Library of Science ONE*, 6, 2011.

[104] L. Warnke. On the method of typical bounded differences. *Combinatorics, Probability and Computing*, 25(2), 2016.

[105] L. Wein and S. Zenios. Pooled testing for HIV screening: Capturing the dilution effect. *Operations Research*, 44:543–569, 1996.

[106] S. Woloshin, N. Patel, and A. Kesselheim. False negative tests for SARS-CoV-2 infection — challenges and implications. *New England Journal of Medicine*, 2020.

[107] L. Zdeborová and F. Krzakala. Statistical physics of inference: thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.

APPENDIX A. OPTIMAL GROUP TESTING

OPTIMAL GROUP TESTING

AMIN COJA-OGHLAN, OLIVER GEBHARD, MAX HAHN-KLIMROTH, PHILIPP LOICK

ABSTRACT. In the group testing problem the aim is to identify a small set of $k \sim n^\theta$ infected individuals out of a population size n, $0 < \theta < 1$. We avail ourselves of a test procedure capable of testing groups of individuals, with the test returning a positive result iff at least one individual in the group is infected. The aim is to devise a test design with as few tests as possible so that the set of infected individuals can be identified correctly with high probability. We establish an explicit sharp information-theoretic/algorithmic phase transition m_{inf} for non-adaptive group testing, where all tests are conducted in parallel. Thus, with more than m_{inf} tests the infected individuals can be identified in polynomial time w.h.p., while learning the set of infected individuals is information-theoretically impossible with fewer tests. In addition, we develop an optimal adaptive scheme where the tests are conducted in two stages. *MSc: 05C80, 60B20, 68P30*

1. INTRODUCTION

1.1. **Background and motivation.** Various intriguing combinatorial problems come as inference tasks where we are to learn a hidden ground truth by means of indirect queries. The goal is to get by with as small a number of queries as possible. The ultimate solution to such a problem should consist of a positive algorithmic result showing that a certain number of queries suffice to learn the ground truth efficiently, complemented by a matching information-theoretic lower bound showing that with fewer queries the problem is insoluble, regardless of computational resources.

Group testing is a prime example of such an inference problem [6]. The objective is to identify within a large population of size n a subset of k individuals infected with a rare disease. We presume that the number of infected individuals scales as a power $k = \lceil n^\theta \rceil$ of the population size with an exponent $\theta \in (0,1)$, a parametrisation suited to modelling the pivotal early stages of an epidemic [36]. Indeed, since early on in an epidemic test kits might be in short supply, it is vital to get the most diagnostic power out the least number of tests. To this end we assume that the test gear is capable of not merely testing a single individual but an entire group. The test comes back positive if any one individual in the group is infected and negative otherwise. While in *non-adaptive* group testing all tests are conducted in parallel, in *adaptive* group testing test are conducted in several stages. In either case we are free to allocate individuals to test groups as we please. Randomisation is allowed. What is the least number of tests required so that the set of infected individuals can be inferred from the test results with high probability? Furthermore, in adaptive group testing, what is the smallest depth of test stages required?

Closing the considerable gaps that the best prior bounds left, the main results of this paper furnish matching algorithmic and information-theoretic bounds for both adaptive and non-adaptive group testing. Specifically, the best prior information-theoretic lower bound derives from the following folklore observation. Suppose that we conduct m tests that each return either 'positive' or 'negative'. Then to correctly identify the set of infected individuals we need the total number 2^m of conceivable test results to asymptotically exceed the number $\binom{n}{k}$ of possible sets of infected individuals. Hence, $2^m \geq (1+o(1))\binom{n}{k}$. Thus, Stirling's formula yields the lower bound

$$m_{\mathrm{ad}} = \frac{1-\theta}{\ln 2} n^\theta \ln n, \tag{1.1}$$

which applies to both adaptive and non-adaptive testing. On the positive side, a randomised non-adaptive test design with

$$m_{\mathrm{DD}} \sim \frac{\max\{\theta, 1-\theta\}}{\ln^2 2} n^\theta \ln n \tag{1.2}$$

Supported by DFG CO 646/3 and Stiftung Polytechnische Gesellschaft. An extended abstract version of this work has been submitted to the COLT 2020 conference.

tests exists from which a greedy algorithm called DD correctly infers the set of infected individuals w.h.p. [22]. Clearly, $m_{\mathrm{ad}} < m_{\mathrm{DD}}$ for all infection densities θ and $m_{\mathrm{DD}}/m_{\mathrm{ad}} \to \infty$ as $\theta \to 1$. In addition, there is an efficient adaptive three-stage group testing scheme that asymptotically matches the lower bound m_{ad} [33].

We proceed to state the main results of the paper. First, improving both the information-theoretic and the algorithmic bounds, we present optimal results for non-adaptive group testing. Subsequently we show how the non-adaptive result can be harnessed to perform adaptive group testing with the least possible number $(1+o(1))m_{\mathrm{ad}}$ of tests in only two stages.

1.2. **Non-adaptive group testing.** A *non-adaptive test design* is a bipartite graph $G = (V \cup F, E)$ with one vertex class $V = V_n = \{x_1, \ldots, x_n\}$ representing individuals and the other class $F = F_m = \{a_1, \ldots, a_m\}$ representing tests. For a vertex v of G denote by $\partial v = \partial_G v$ the set of neighbours of v. Thus, an individual x_j takes part in a test a_i iff $x_j \in \partial a_i$. Since we can shuffle the individuals randomly, we may safely assume that the vector $\boldsymbol{\sigma} \in \{0,1\}^V$ whose 1-entries mark the infected individuals is a uniformly random vector of Hamming weight k. Furthermore, the test results induced by $\boldsymbol{\sigma}$ read

$$\hat{\boldsymbol{\sigma}}_{a_i} = \hat{\boldsymbol{\sigma}}_{G,a_i} = \max_{x \in \partial a_i} \boldsymbol{\sigma}_x.$$

Hence, given $\hat{\boldsymbol{\sigma}} = \hat{\boldsymbol{\sigma}}_G = (\hat{\boldsymbol{\sigma}}_{G,a})_{a \in F}$ and G we aim to infer $\boldsymbol{\sigma}$. Thus, we can represent an inference procedure by a function $\mathscr{A}_G : \{0,1\}^m \to \{0,1\}^n$. The following theorem improves the lower bound on the number of tests required for successful inference. Let

$$m_{\mathrm{inf}} = m_{\mathrm{inf}}(n,\theta) = \max\left\{\frac{\theta}{\ln^2 2}, \frac{1-\theta}{\ln 2}\right\} n^\theta \ln n. \tag{1.3}$$

Theorem 1.1. *For any $0 < \theta < 1$, $\varepsilon > 0$ there exists $n_0 = n_0(\theta,\varepsilon)$ such that for all $n > n_0$, all test designs G with $m \le (1-\varepsilon)m_{\mathrm{inf}}$ tests and for every function $\mathscr{A}_G : \{0,1\}^m \to \{0,1\}^n$ we have*

$$\mathbb{P}\left[\mathscr{A}_G(\hat{\boldsymbol{\sigma}}_G) = \boldsymbol{\sigma}\right] < \varepsilon. \tag{1.4}$$

Theorem 1.1 rules out both deterministic and randomised test designs and inference procedures because (1.4) holds uniformly for all G and all $\mathscr{A}_G$. Thus, no test design, randomised or not, with fewer than m_{inf} tests allows to infer the set of infected individuals with a non-vanishing probability. Since m_{inf} matches m_{DD} from (1.2) for $\theta \ge 1/2$, Theorem 1.1 shows that the positive result from [22] is optimal in this regime. The following theorem closes the remaining gap by furnishing an optimal positive result for all θ.

Theorem 1.2. *For any $0 < \theta < 1$, $\varepsilon > 0$ there is $n_0 = n_0(\theta,\varepsilon)$ such that for every $n > n_0$ there exist a randomised test design $\boldsymbol{G}$ comprising $m \le (1+\varepsilon)m_{\mathrm{inf}}$ tests and a polynomial time algorithm* SPIV *that given $\boldsymbol{G}$ and the test results $\hat{\boldsymbol{\sigma}}_{\boldsymbol{G}}$ outputs $\boldsymbol{\sigma}$ w.h.p.*

An obvious candidate for an optimal test design appears be a plain random bipartite graph. In fact, prior to the present work the best known test design consisted of a uniformly random bipartite graph where all vertices in V_n have the same degree Δ. In other words, every individual independently joins Δ random test groups. Applied to this random Δ-out test design the DD algorithm correctly recovers the set of infected individuals in polynomial time provided that the number of tests exceeds m_{DD} from (1.2). However, m_{DD} strictly exceeds m_{inf} for $\theta < 1/2$. While the random Δ-out test design with $(1+o(1))m_{\mathrm{inf}}$ tests is known to admit an exponential time algorithm that successfully infers the set of infected individuals w.h.p. [11], we do not know of a polynomial time that solves this inference problem. Instead, to facilitate the new efficient inference algorithm SPIV the test design for Theorem 1.2 relies on a blend of a geometric and a random construction that is inspired by recent advances in coding theory known as spatially coupled low-density parity check codes [18, 26].

Finally, for

$$\theta \le \frac{\ln 2}{1+\ln 2} \approx 0.41 \tag{1.5}$$

the number m_{inf} of tests required by Theorem 1.2 matches the folklore lower bound m_{ad} from (1.2) that applies to both adaptive and non-adaptive group testing. Hence, in this regime adaptivity confers no advantage. By contrast, for $\theta > \ln(2)/(1+\ln 2)$ the adaptive bound m_{ad} is strictly smaller than m_{inf}. Consequently, in this regime at least two test stages are necessary to match the lower bound. Indeed, the next theorem shows that two stages suffice.

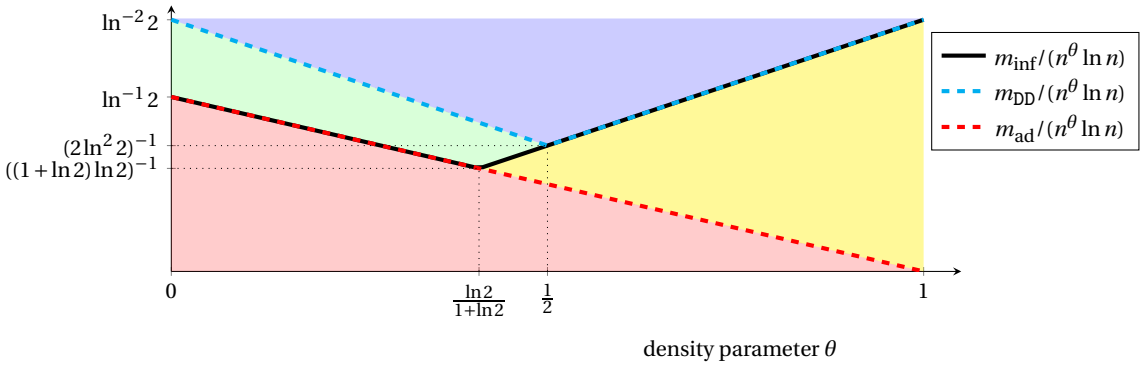

FIGURE 1. The phase transitions in group testing. The best previously known algorithm `DD` succeeds in the blue but not in the green region. The new algorithm `SPIV` succeeds in both the blue and the green region. The black line indicates the non-adaptive information-theoretic threshold m_{inf}, below which non-adaptive group testing is impossible. In the red area even (multi-stage) adaptive inference is impossible. Finally, the two-stage adaptive group testing algorithm from Theorem 1.3 succeeds in the yellow region.

1.3. **Adaptive group testing.** A *two-stage test design* consists of a bipartite graph $G = (V, F)$ along with a second bipartite graph $G' = G'(G, \hat{\boldsymbol{\sigma}}_G) = (V', F')$ with $V' \subset V$ that may depend on the tests results $\hat{\boldsymbol{\sigma}}_G$ of the first test design G. Hence, the task is to learn $\boldsymbol{\sigma}$ correctly w.h.p. from $G, \hat{\boldsymbol{\sigma}}_G, G'$ and the test results $\hat{\boldsymbol{\sigma}}_{G'}$ from the second stage while minimising the total number $|F| + |F'|$ of tests. The following theorem shows that a two-stage test design and an efficient inference algorithm exist that meet the multi-stage adaptive lower bound (1.1).

Theorem 1.3. *For any $0 < \theta < 1$, $\varepsilon > 0$ there is $n_0 = n_0(\theta, \varepsilon)$ such that for every $n > n_0$ there exist a two-stage test design with no more than $(1+\varepsilon)m_{\mathrm{ad}}$ tests in total and a polynomial time inference algorithm that outputs $\boldsymbol{\sigma}$ with high probability.*

Theorem 1.3 improves over [33] by reducing the number of stages from three to two, thus potentially significantly reducing the overall time required to complete the test procedure [10, 28]. The proof of Theorem 1.3 combines the test design and efficient algorithm from Theorem 1.2 with ideas from [32].

The question of whether an 'adaptivity gap' exists for group testing, i.e., if the number of tests can be reduced by allowing multiple stages, has been raised prominently [6]. Theorems 1.1–1.3 answer this question comprehensively. While for $\theta \le \ln(2)/(1+\ln(2)) \approx 0.41$ adaptivity confers no advantage, Theorem 1.1 shows that for $\theta > \ln(2)/(1+\ln(2))$ there is a widening gap between m_{ad} and the number m_{inf} of tests required by the optimal non-adaptive test design. Further, Theorem 1.3 demonstrates that this gap can be closed by allowing merely two stages. Figure 1 illustrates the thresholds from Theorems 1.1–1.3.

1.4. **Discussion.** The group testing problem was first raised in 1943, when Dorfman [15] proposed a two-stage adaptive test design to test the US Army for syphilis: in a first stage disjoint groups of equal size are tested. All members of negative test groups are definitely uninfected. Then, in the second stage the members of positive test groups get tested individually. Of course, this test design is far from optimal, but Dorfman's contribution triggered attempts at devising improved test schemes.

At first combinatorial group testing, where the aim is to construct a test design that is guaranteed to succeed on *all* vectors $\boldsymbol{\sigma}$, attracted significant attention. This version of the problem was studied, among others, by Erdős and Rényi [17], D'yachkov and Rykov [16] and Kautz and Singleton [23]. Hwang [20] was the first to propose an adaptive test design that asymptotically meets the information-theoretic lower bound m_{ad} from (1.1) for all $\theta \in [0,1]$. However, this test design requires an unbounded number of stages. Conversely, D'yachkov and Rykov [16] showed that m_{ad} tests do not suffice for non-adaptive group testing. Indeed, $m \ge \min\{\Omega(k^2), n\}$ tests are required non-adaptively, making individual testing optimal for $\theta > 1/2$. For an excellent survey of combinatorial group testing see [6].

Since the early 2000s attention has shifted to probabilistic group testing, which we study here as well. Thus, instead of asking for test designs and algorithms that are guaranteed to work for *all* $\boldsymbol{\sigma}$, we are content with recovering $\boldsymbol{\sigma}$ with high probability. Berger and Levenshtein [8] presented a two-stage probabilistic group testing design and

algorithm requiring

$$m_{\mathrm{BL,ad}} \sim 4n^{\theta}\ln n$$

tests in expectation. Their test design, known as the Bernoulli design, is based on a random bipartite graph where each individual joins every test independently with a carefully chosen edge probability. For a fixed θ the number $m_{\mathrm{BL,ad}}$ of tests is within a bounded factor of the information-theoretic lower bound m_{ad} from (1.1), although the gap $m_{\mathrm{ad}}/m_{\mathrm{BL,ad}}$ diverges as $\theta \to 1$. Unsurprisingly, the work of Berger and Levenshtein spurred efforts at closing the gap. Mézard, Tarzia and Toninelli proposed a different two-stage test design whose first stage consists of a random bipartite graph called the constant weight design [29]. Here each individual independently joins an equal number of random tests. For their two-stage design they obtained an inference algorithm that gets by with about

$$m_{\mathrm{MTT,ad}} \sim \frac{1-\theta}{\ln^2 2}\, n^{\theta}\ln n. \tag{1.6}$$

tests, a factor of $1/\ln 2$ above the elementary bound m_{ad}. Conversely, Mézard, Tarzia and Toninelli showed by means of the FKG inequality and positive correlation arguments that two-stage test algorithms from a certain restricted class cannot beat the bound (1.6). Furthermore, Aldridge, Johnson and Scarlett analysed non-adaptive test designs and inference algorithms [4, 22]. For the Bernoulli test design their best efficient algorithm DD requires

$$m_{\mathrm{DD,Be}} \sim \mathrm{e}\cdot\max\{\theta, 1-\theta\}\, n^{\theta}\ln n.$$

tests. For the constant weight design they obtained the bound m_{DD} from (1.2). In addition, in a previous article [11] we showed that on the constant weight design an exponential time algorithm correctly identifies the set of infected individuals w.h.p. if the number of tests exceeds m_{inf} from (1.3). Furthermore, Scarlett [33] discovered the aforementioned three-stage test design and polynomial time algorithm that matches the universal lower bound m_{ad} from (1.1). Finally, concerning lower bounds, in the case of a linear number $k = \Theta(n)$ infected individuals Aldridge [5] showed via arguments similar to [29] that individual testing is optimal in the non-adaptive case, while Ungar [35] proved that individual testing is optimal even adaptively once $k \geq (3-\sqrt{5})n/2$.

A further variant of group testing is known as the quantitative group testing or the coin weighing problem. In this problem tests are assumed to not merely indicate the presence of at least one infected individual but to return the number of infected individuals. Thus, the tests are significantly more powerful. For quantitative group testing with k infected individuals Alaoui, Ramdas, Krzakala, Zdeborová and Jordan [3] presented a test design with

$$m_{\mathrm{QGT}} \sim 2\left(1+\frac{(n-k)\ln(1-k/n)}{k\ln(k/n)}\right)\frac{k\ln(n/k)}{\ln(k)} \qquad \text{for} \qquad k=\Theta(n)$$

tests from which the set of infected individuals can be inferred in exponential time; the paper actually deals with the slightly more general pooled data problem. However, no efficient algorithm is known to come within a constant factor of m_{QGT}. Indeed, the best efficient algorithm, due to the same authors [2], requires $\Omega(k\ln(n/k))$ tests.

More broadly, the idea of harnessing random graphs to tackle inference problems has been gaining momentum. One important success has been the development of capacity achieving linear codes called spatially coupled low-density parity check ('LDPC') codes [26, 27]. The Tanner graphs of these codes, which represent their check matrices, consist of a linear sequence of sparse random bipartite graphs with one class of vertices corresponding to the bits of the codeword and the other class corresponding to the parity checks. The bits and the checks are divided equitably into a number of compartments, which are arranged along a line. Each bit of the codeword takes part in random checks in a small number of preceding and subsequent compartments of checks along the line. This combination of a spatial arrangement and randomness facilitates efficient decoding by means of the Belief Propagation message passing algorithm. Furthermore, the general design idea of combining a linear spatial structure with a random graph has been extended to other inference problems. Perhaps the most prominent example is compressed sensing, i.e., solving an underdetermined linear system subject to a sparsity constraint [13, 14, 24, 25], where a variant of Belief Propagation called Approximate Message Passing matches an information-theoretic lower bound from [37].

While in some inference problems such as LDPC decoding or compressed sensing the number of queries required to enable an efficient inference algorithm matches the information-theoretic lower bound, in many other problems gaps remain. A prominent example is the stochastic block model [1, 12, 30], an extreme case of which is the notorious planted clique problem [7]. For both these models the existence of a genuine computationally

intractable phase where the problem can be solved in exponential but not in polynomial time appears to be an intriguing possibility. Further examples include code division multiple access [34, 38], quantitative group testing [2], sparse principal component analysis [9] and sparse high-dimensional regression [31]. The problem of solving the group testing inference problem on the test design from [22] could be added to the list. Indeed, while an exponential time algorithm (that reduces the problem to minimum hypergraph vertex cover) infers the set of infected individuals w.h.p. with only $(1+\varepsilon)m_{\inf}$ tests, the best known polynomial algorithm requires $(1+\varepsilon)m_{\mathrm{DD}}$ tests.

Instead of developing a better algorithm for the test design from [22], here we exercise the discretion of constructing a different test design that the group testing problem affords. The new design is tailored to enable an efficient algorithm SPIV for Theorem 1.2 that gets by with $(1+\varepsilon)m_{\inf}$ tests. While prior applications of the idea of spatial coupling such as coding and compressed sensing required sophisticated message passing algorithms [18, 26, 27], the SPIV algorithm is purely combinatorial and extremely transparent. The main step of the algorithm merely computes a weighted sum to discriminate between infected individuals and 'disguised' healthy individuals. Furthermore, the analysis of the algorithm is based on a technically subtle but conceptually clean large deviations analysis. This technique of blending combinatorial ideas and large deviations methods with spatial coupling promises to be an exciting route for future research. Applications might include noisy versions of group testing, quantitative group testing or the coin weighing problem [2]. Beyond these immediate extensions, it would be most interesting to see if the SPIV strategy extends to other inference problems for sparse data.

1.5. **Organisation.** After collecting some preliminaries and introducing notation in Section 2, we prove Theorem 1.1 in Section 3. Section 4 then deals with the test design and the inference algorithm for Theorem 1.2. Finally, in Section 5 we prove Theorem 1.3.

2. Preliminaries

As we saw in Section 1.2 a non-adaptive test design can be represented by a bipartite graph $G = (V \cup F, E)$ with one vertex class V representing the individuals and the other class F representing the tests. We refer to the number $|V|$ of individuals as the *order* of the test design and to the number $|F|$ of tests as its *size*. For a vertex v of G we denote by $\partial_G v$ the set of neighbours. Where G is apparent from the notation we just write ∂v. Furthermore, for an integer $k \le |V|$ we denote by $\boldsymbol{\sigma}_{G,k} = (\boldsymbol{\sigma}_{G,k,x})_{x\in V} \in \{0,1\}^V$ a random vector of Hamming weight k. Additionally, we let

$$\hat{\boldsymbol{\sigma}}_{G,k} = (\hat{\boldsymbol{\sigma}}_{G,k,a})_{a\in F} \in \{0,1\}^F \qquad \text{with} \qquad \hat{\boldsymbol{\sigma}}_{G,k,a} = \max_{x\in\partial_G a} \boldsymbol{\sigma}_{G,k,x} \tag{2.1}$$

be the associated vector of test results. Where G and/or k are apparent from the context, we drop them from the notation. More generally, for a given vector $\tau \in \{0,1\}^V$ we introduce a vector $\hat{\tau}_G = (\hat{\tau}_{G,a})_{a\in F}$ by letting $\hat{\tau}_{G,a} = \max_{x\in\partial_G a}\tau_x$, just as in (2.1). Furthermore, for a given $\tau \in \{0,1\}^V$ we let

$$V_0(G,\tau) = \{x \in V : \tau_x = 0\}, \quad V_1(G,\tau) = \{x \in V : \tau_x = 1\}, \quad F_0(G,\tau) = \{a \in F : \hat{\tau}_{G,a} = 0\}, \quad F_1(G,\tau) = \{a \in F : \hat{\tau}_{G,a} = 1\}.$$

The *Kullback-Leibler divergence* of $p, q \in (0,1)$ is denoted by

$$D_{\mathrm{KL}}(q\|p) = q\ln\left(\frac{q}{p}\right) + (1-q)\ln\left(\frac{1-q}{1-p}\right).$$

We will occasionally apply the following Chernoff bound.

Lemma 2.1 ([21])**.** *Let $\boldsymbol{X}$ be a binomial random variable with parameters N, p. Then*

$$\mathbb{P}[X \ge qN] \le \exp\left(-ND_{\mathrm{KL}}(q\|p)\right) \quad \text{for } p < q < 1, \tag{2.2}$$

$$\mathbb{P}[X \le qN] \le \exp\left(-ND_{\mathrm{KL}}(q\|p)\right) \quad \text{for } 0 < q < p. \tag{2.3}$$

In addition, we recall that the *hypergeometric distribution* $\mathrm{Hyp}(L,M,N)$ is defined by

$$\mathbb{P}[\mathrm{Hyp}(L,M,N) = k] = \binom{M}{k}\binom{L-M}{N-k}\binom{L}{N}^{-1}. \qquad (k \in \{0,1,\dots,M\wedge N\}).$$

Hence, out of a total of L items of which M are special we draw N items without replacement and count the number of special items in the draw. The mean of the hypergeometric distribution equals MN/L. It is well known that the Chernoff bound extends to the hypergeometric distribution.

Lemma 2.2 ([19])**.** *For a hypergeometric variable $\boldsymbol{X} \sim \mathrm{Hyp}(L,M,N)$ the bounds* (2.2)–(2.3) *hold with $p = M/L$.*

Throughout the paper we use asymptotic notation $o(\cdot),\omega(\cdot),O(\cdot),\Omega(\cdot),\Theta(\cdot)$ to refer to limit $n\to\infty$. It is understood that the constants hidden in, e.g., a $O(\cdot)$-term may depend on the density parameter θ or other parameters.

3. The information theoretic lower bound

In this section we prove Theorem 1.1. The proof combines techniques based on the FKG inequality and positive correlation that were developed in [6, 29] with new combinatorial ideas. Throughout this section we fix a number $\theta\in(0,1)$ and we let $k=\lceil n^\theta\rceil$.

3.1. **Outline.** The starting point is a simple and well known observation. Namely, for a test design $G=G_{n,m}=(V_n,F_m)$ and a vector $\tau\in\{0,1\}^{F_m}$ of test results let

$$\mathscr{S}_k(G,\tau)=\left\{\sigma\in\{0,1\}^{V_n}:\sum_{x\in V_n}\sigma_x=k,\ \hat\sigma_G=\tau\right\}$$

be the set of all possible vectors σ of Hamming weight k that give rise to the test results τ. Further, let $Z_k(G,\tau)=|\mathscr{S}_k(G,\tau)|$ be the number of such vectors σ. Also recall that $\boldsymbol\sigma=\boldsymbol\sigma_{G,k}\in\{0,1\}^{V_n}$ is a random vector of Hamming weight k and that $\hat{\boldsymbol\sigma}=\hat{\boldsymbol\sigma}_{G,k}$ comprises the test results that $\boldsymbol\sigma$ renders under the test design G. We observe that the posterior of $\boldsymbol\sigma$ given $\hat{\boldsymbol\sigma}$ is the uniform distribution on $\mathscr{S}_k(G,\hat{\boldsymbol\sigma})$.

Fact 3.1. *For any G, $\sigma\in\{0,1\}^{V_n}$ we have $\mathbb{P}[\boldsymbol\sigma=\sigma\mid\hat{\boldsymbol\sigma}]=\mathbf{1}\{\sigma\in\mathscr{S}_k(G,\hat{\boldsymbol\sigma})\}/Z_k(G,\hat{\boldsymbol\sigma})$.*

As an immediate consequence of Fact 3.1, the success probability of any inference scheme $\mathscr{A}_G:\{0,1\}^{F_m}\to\{0,1\}^{V_n}$ is bounded by $1/Z_k(G,\hat{\boldsymbol\sigma})$. Indeed, an optimal inference algorithm is to simply return a uniform sample from $\mathscr{S}_k(G,\hat{\boldsymbol\sigma})$.

Fact 3.2. *For any test design G and for any $\mathscr{A}_G:\{0,1\}^{F_m}\to\{0,1\}^{V_n}$ we have $\mathbb{P}[\mathscr{A}_G(\hat{\boldsymbol\sigma})=\boldsymbol\sigma\mid\hat{\boldsymbol\sigma}]\le 1/Z_k(G,\hat{\boldsymbol\sigma})$.*

Hence, in order to prove Theorem 1.1 we just need to show that $Z_k(G,\hat{\boldsymbol\sigma})$ is large for any test design G with $m<(1-\varepsilon)m_{\inf}$ tests. In other words, we need to show that w.h.p. there are many vectors $\sigma\in\mathscr{S}_k(G,\hat{\boldsymbol\sigma})$ that give rise to the test results $\hat{\boldsymbol\sigma}$.

We obtain these σ by making diligent local changes to $\boldsymbol\sigma$. More precisely, we identify two sets $V_{0+}=V_{0+}(G,\boldsymbol\sigma)$, $V_{1+}=V_{1+}(G,\boldsymbol\sigma)$ of individuals whose infection status can be flipped without altering the test results. Specifically, following [5] we call an individual $x\in V_n$ *disguised* if every test $a\in\partial_G x$ contains another individual $y\in\partial_G a\setminus\{x\}$ with $\boldsymbol\sigma_y=1$. Let $V_+=V_+(G,\boldsymbol\sigma)$ be the set of all disguised individuals. Moreover, let

$$V_{0+}=V_{0+}(G,\boldsymbol\sigma)=\{x\in V_+:\boldsymbol\sigma_x=0\},\qquad V_{1+}=V_{1+}(G,\boldsymbol\sigma)=\{x\in V_+:\boldsymbol\sigma_x=1\}. \tag{3.1}$$

Hence, V_{0+} is the set of all healthy disguised individuals while V_{1+} contains all infected disguised individuals.

Fact 3.3. *We have $Z_k(G,\hat{\boldsymbol\sigma})\ge|V_{0+}(G,\boldsymbol\sigma)|\cdot|V_{1+}(G,\boldsymbol\sigma)|$.*

Proof. For a pair $(x,y)\in V_{0+}(G,\boldsymbol\sigma)\times V_{1+}(G,\boldsymbol\sigma)$ obtain $\boldsymbol\tau$ from $\boldsymbol\sigma$ by letting $\boldsymbol\tau_x=1,\boldsymbol\tau_y=0$ and $\boldsymbol\tau_z=\boldsymbol\sigma_z$ for all $z\neq x,y$. Then $\boldsymbol\tau$ has Hamming weight k and $\hat{\boldsymbol\tau}_G=\hat{\boldsymbol\sigma}$. Thus, $\boldsymbol\tau\in\mathscr{S}_k(G,\hat{\boldsymbol\sigma})$. □

Hence, an obvious proof strategy for Theorem 1.1 is to exhibit a large number of disguised individuals. A similar strategy has been pursued in the proof of the conditional lower bound of Mézard, Tarzia and Toninelli [29] and the proof of Aldridge's lower bound for the linear case $k=\Theta(n)$ [5]. Both [5, 29] exhibit disguised individuals via positive correlation and the FKG inequality. However, we do not see how to stretch such arguments to obtain the desired lower bound for all $\theta\in(0,1)$. Yet for θ *extremely* close to one it is possible to combine the positive correlation argument with new combinatorial ideas to obtain the following.

Proposition 3.4. *For any $\varepsilon>0$ there exists $\theta_0=\theta_0(\varepsilon)<1$ such that for every $\theta\in(\theta_0,1)$ there exists $n_0=n_0(\theta,\varepsilon)$ such that for all $n>n_0$ and all test designs $G=G_{n,m}$ with $m\le(1-\varepsilon)m_{\inf}$ we have*

$$\mathbb{P}[|V_{0+}(G,\boldsymbol\sigma)|\wedge|V_{1+}(G,\boldsymbol\sigma)|\ge\ln n]>1-\varepsilon.$$

The proof of Proposition 3.4 can be found in Section 3.2.

The second step towards Theorem 1.1 is a reduction from larger to smaller values of θ. Suppose we wish to apply a test scheme designed for an infection density $\theta\in(0,1)$ to a larger infection density $\theta'\in(\theta,1)$. Then we could dilute the larger infection density by adding a large number of healthy 'dummy' individuals. A careful analysis of this dilution process yields the following result. Due to the elementary lower bound (1.1) we need not worry about $\theta\le\ln(2)/(1+\ln 2)$.

Proposition 3.5. *For any* $\ln(2)/(1+\ln(2)) < \theta < \theta' < 1$, $t > 0$ *there exists* $n_0 = n_0(\theta,\theta',t) > 0$ *such that for every* $n > n_0$ *and for every test design* G *of order* n *there exist an integer* n' *such that*

$$k = \lceil n^{\theta} \rceil = \lceil n'^{\theta'} \rceil$$

and a test design G' *of order* n' *with the same number of tests as* G *such that the following is true. Let* $\boldsymbol{\tau} \in \{0,1\}^{V_{n'}}$ *be a random vector of Hamming weight* k *and let* $\hat{\boldsymbol{\tau}}_a = \max_{x \in \partial_{G'} a} \boldsymbol{\tau}_x$ *comprise the tests results of* G'. *Then*

$$\mathbb{P}\left[Z_k(G,\hat{\boldsymbol{\sigma}}) \le t\right] \le \mathbb{P}\left[Z_k(G',\hat{\boldsymbol{\tau}}) \le t\right].$$

Hence, if a test design exists for $\theta < \theta'$ that beats $m_{\inf}(n,\theta)$, then there is a test design for infection density θ' that beats $m_{\inf}(n',\theta')$. We prove Proposition 3.4 in Section 3.2. Theorem 1.1 is an easy consequence of Propositions 3.4 and 3.5.

Proof of Theorem 1.1. For $\theta \le \ln(2)/(1+\ln(2))$ the assertion follows from the elementary lower bound (1.1). Hence, fix $\varepsilon > 0$ and assume for contradiction that some $\theta \in (\ln(2)/(1+\ln(2)),1)$ for infinitely many n admits a test design G of order n and size $m \le (1-\varepsilon)m_{\inf}(n,\theta)$ such that $\mathbb{P}[Z_k(G,\hat{\boldsymbol{\sigma}}_G) \le t] \ge \varepsilon$. Then Proposition 3.5 shows that for $\theta' > \theta$ arbitrarily close to one for an integer n' with $k = \lceil n'^{\theta'} \rceil$ a test design $G' = G_{n',m}$ exists such that

$$\mathbb{P}\left[Z_k(G',\hat{\boldsymbol{\tau}}) \le 1/\varepsilon\right] \ge \varepsilon. \tag{3.2}$$

Furthermore, (1.3) shows that for large n,

$$m_{\inf}(n',\theta') = \frac{\theta'}{\ln^2 2} n'^{\theta'} \ln n' = \frac{\theta + o(1)}{\ln^2 2} n^{\theta} \ln n = (1+o(1)) m_{\inf}(n,\theta).$$

Hence, the number m of tests of G' satisfies $m \le (1-\varepsilon+o(1))m_{\inf}(n',\theta')$. Thus, (3.2) contradicts Fact 3.3 and Proposition 3.4. □

3.2. **Proof of Proposition 3.4.** Given a small $\varepsilon > 0$ we choose $\theta_0 = \theta_0(\varepsilon) \in (0,1)$ sufficiently close to one and fix $\theta \in (\theta_0,1)$. Additionally, pick $\xi = \xi(\varepsilon,\theta) \in (0,1)$ such that

$$2(1-\theta) < \xi < \theta\varepsilon. \tag{3.3}$$

We fix ε,θ,ξ throughout this section.

To avoid the (mild) stochastic dependencies that result from the total number of infected individuals being fixed, instead of $\boldsymbol{\sigma}$ we will consider a vector $\boldsymbol{\chi} \in \{0,1\}^{V_n}$ whose entries are stochastically independent. Specifically, every entry of $\boldsymbol{\chi}$ equals one with probability

$$p = \frac{k - \sqrt{k}\ln n}{n}$$

independently. Let $\hat{\boldsymbol{\chi}}_G \in \{0,1\}^{F_m}$ be the corresponding vector of test results. The following lemma shows that it suffices to estimate $|V_{0+}(G,\boldsymbol{\chi})|, |V_{1+}(G,\boldsymbol{\chi})|$. Let G denote an arbitrary test design with individuals $V_n = \{x_1,\ldots,x_n\}$ and tests $F_m = \{a_1,\ldots,a_m\}$.

Lemma 3.6. *There is* $n_0 = n_0(\theta,\varepsilon)$ *such that for all* $n > n_0$ *and for all* G *with* $m \le m_{\inf}$ *the following is true:*

$$\textit{if } \mathbb{P}\left[|V_{0+}(G,\boldsymbol{\chi})| \wedge |V_{1+}(G,\boldsymbol{\chi})| \ge 2\ln n\right] > 1-\varepsilon/4, \textit{ then } \mathbb{P}\left[|V_{0+}(G,\boldsymbol{\sigma})| \wedge |V_{1+}(G,\boldsymbol{\sigma})| \ge \ln n\right] > 1-\varepsilon.$$

Proof. Let $\mathscr{X} = \{k - 2\sqrt{k}\ln n \le \sum_{x \in V_n} \boldsymbol{\chi}_x \le k\}$. The Chernoff bound shows for large enough n,

$$\mathbb{P}[\mathscr{X}] > 1 - \eta/4. \tag{3.4}$$

Further, given $\mathscr{X}$ we can couple $\boldsymbol{\chi},\boldsymbol{\sigma}$ such that the latter is obtained by turning $k - \sum_{x \in V_n} \boldsymbol{\chi}_x$ random zero entries of the former into ones. Since turning zero entries into ones can only increase the number of disguised individuals, on $\mathscr{X}$ we have

$$V_{1+}(G,\boldsymbol{\sigma}) \ge V_{1+}(G,\boldsymbol{\chi}). \tag{3.5}$$

Of course, it is possible that $|V_{0+}(G,\boldsymbol{\sigma})| < |V_{0+}(G,\boldsymbol{\chi})|$. But since on $\mathscr{X}$ the two vectors $\boldsymbol{\sigma},\boldsymbol{\chi}$ differ in no more than $2\sqrt{k}\ln n$ entries, we obtain the bound

$$\mathbb{E}\left[|V_{0+}(G,\boldsymbol{\chi})| - |V_{0+}(G,\boldsymbol{\sigma})| \,|\, \mathscr{X}\right] \le \frac{2\sqrt{k}\ln n}{n-k} |V_{0+}(G,\boldsymbol{\chi})| < n^{-1/3} |V_{0+}(G,\boldsymbol{\chi})|,$$

provided n is sufficiently large. Hence, Markov's inequality shows that for large enough n,

$$\mathbb{P}\left[|V_{0+}(G,\boldsymbol{\chi})| - |V_{0+}(G,\boldsymbol{\sigma})| > |V_{0+}(G,\boldsymbol{\chi})|/2 \mid \mathscr{X}\right] < \varepsilon/4. \tag{3.6}$$

Combining (3.4), (3.5) and (3.6) completes the proof. □

As a next step we show that there is no point in having very big tests a that contain more than, say, $\Gamma = \Gamma(n,\theta) = n^{1-\theta}\ln n$ individuals. This is because anyway all such tests are positive w.h.p., so there is little point in actually conducting them. Indeed, the following lemma shows that w.h.p. all tests of very high degree contain at least two infected individuals.

Lemma 3.7. *There exists $n_0 = n_0(\theta,\varepsilon) > 0$ such that for all $n > n_0$ and all test designs G with $m \le m_{\inf}$ tests,*

$$\mathbb{P}\left[\exists a \in F_m : |\partial_G a| > \Gamma \wedge |\partial_G a \cap V_1(G,\boldsymbol{\chi})| \le 1\right] < \varepsilon/8.$$

Proof. Consider a test a of degree $\gamma = |\partial_G a| \ge \Gamma$. Because in $\boldsymbol{\chi}$ each of the γ individuals that take part in a is infected with probability p independently, we have

$$\mathbb{P}\left[|\partial_G a \cap V_1(G,\boldsymbol{\sigma})| \le 1\right] = \mathbb{P}\left[\mathrm{Bin}(\gamma,p) \le 1\right] = (1-p)^\gamma + \gamma p(1-p)^{\gamma-1} \le (1+\gamma p/(1-p))\exp(-\gamma p) = n^{o(1)-1}. \tag{3.7}$$

Since $m \le m_{\inf} = O(n^\theta)$ for a fixed $\theta < 1$, the assertion follows from (3.7) and the union bound. □

Let G^* be test design obtained from $G = G_{n,m}$ by deleting all tests of degree larger than Γ. If indeed every test of degree at least Γ contains at least two infected individuals, then $V_{0+}(G^*,\boldsymbol{\chi}) = V_{0+}(G,\boldsymbol{\chi})$ and $V_{1+}(G^*,\boldsymbol{\chi}) = V_{1+}(G,\boldsymbol{\chi})$. Hence, Lemma 3.7 shows that it suffices to bound $|V_{0+}(G^*,\boldsymbol{\chi})|, |V_{1+}(G^*,\boldsymbol{\chi})|$. To this end we observe that G^* contains few individuals of very high degree.

Lemma 3.8. *There is $n_0 = n_0(\theta,\varepsilon) > 0$ such that for all $n > n_0$ and all test designs G with $m \le m_{\inf}$ we have*

$$\left|\left\{x \in V_n : |\partial_{G^*} x| > \ln^3 n\right\}\right| \le \frac{n\ln\ln n}{\ln n}.$$

Proof. Since $\max_{a\in F_m} |\partial_{G^*} a| \le \Gamma = n^{1-\theta}\ln n$, double counting yields

$$\sum_{x\in V_n} |\partial_{G^*} x| = \sum_{a\in F_m} |\partial_{G^*} a| \le m_{\inf}\Gamma = O(n\ln^2 n).$$

Consequently, there are no more than $O(n/\ln n)$ individuals $x \in V_n$ with $|\partial_{G^*} x| > \ln^3 n$. □

Further, obtain $G^{(0)}$ from G^* by deleting all individuals of degree greater than $\ln^3 n$ (but keeping all tests). Then the degrees of $G^{(0)}$ satisfy

$$\max_{a\in F(G^{(0)})} |\partial_{G^{(0)}} a| \le \Gamma, \qquad \max_{x\in V(G^{(0)})} |\partial_{G^{(0)}} x| \le \ln^3 n. \tag{3.8}$$

Let $\boldsymbol{\chi}^{(0)} = (\boldsymbol{\chi}_x)_{x\in V(G^{(0)})}$ signify the restriction of $\boldsymbol{\chi}$ to the individuals that remain in $G^{(0)}$.

With these preparations in place we are ready to commence the main step of the proof of Proposition 3.4. Given a test design G with $m \le (1-\varepsilon)m_{\inf}$ we are going to construct a sequence $y_1, y_2, \ldots, y_N$, $N = \lceil n^{1-\xi}\rceil$, of individuals of $G^{(0)}$ such that each y_i individually has a moderately high probability of being disguised. Of course, to conclude that in the end a large number of disguised y_i actually materialise, we need to cope with stochastic dependencies. To this end we will pick individuals y_i that have pairwise distance at least five in $G^{(0)}$. The degree bounds (3.8) guarantee a sufficient supply of such far apart individuals.

To be precise, starting from $G^{(0)}$ we construct a sequence of test designs $G^{(1)}, G^{(2)}, \ldots, G^{(N)}$ inductively as follows. For each $i \ge 1$ select a variable $y_{i-1} \in V(G^{(i-1)})$ whose probability of being disguised is maximum; ties are broken arbitrarily. In formulas,

$$\mathbb{P}\left[y_{i-1} \in V_+(G^{(i-1)},\boldsymbol{\chi}^{(i-1)})\right] = \max_{y\in V(G^{(i-1)})} \mathbb{P}\left[y \in V_+(G^{(i-1)},\boldsymbol{\chi}^{(i-1)})\right],$$

where, of course, $\boldsymbol{\chi}^{(i-1)}$ is the only random object. Then obtain $G^{(i)}$ from $G^{(i-1)}$ by removing y_{i-1} along with all vertices (i.e., tests or individuals) at distance at most four from y_{i-1}. Moreover, let $\boldsymbol{\chi}^{(i)}$ denote the restriction $(\boldsymbol{\chi}_x)_{x\in V(G^{(i)})}$ of $\boldsymbol{\chi}$ to $G^{(i)}$. The following lemma estimates the probability of y_i being disguised. Let $m^* = |F(G^*)|$ be the total number of tests of G of degree at most Γ.

Lemma 3.9. *There exists $n_0 = n_0(\varepsilon,\theta,\xi)$ such that for all $n > n_0$ and all G with $m \le (1-\varepsilon)m_{\inf}$ we have*

$$\min_{1\le i\le N} \mathbb{P}\left[y_i \in V_+(G^{(i)})\right] \ge \exp\left(-\frac{m\ln^2 2}{n^\theta} - 1\right).$$

The proof of Lemma 3.9 requires three intermediate steps. First, we need a lower bound on number of individuals in $G^{(i)}$. Recall that $N = \lceil n^{1-\xi}\rceil$.

Claim 3.10. *We have* $\min_{0\le i\le N} |V(G^{(i)})| \ge n - N\Gamma^2\ln^6 n$.

Proof. Since throughout the construction of the $G^{(i)}$ we only delete vertices, the degree bound (3.8) implies

$$\max_{a\in F(G^{(i)})} |\partial_{G^{(i)}} a| \le \Gamma = n^{1-\theta}\ln n, \qquad \max_{x\in V(G^{(i)})} |\partial_{G^{(i)}} x| \le \ln^3 n \qquad \text{for all } i \le N. \tag{3.9}$$

We now proceed by induction on i. For $i = 0$ there is nothing to show. Going from i to $i+1 \le N$, we notice that because all individuals $x \in V(G^{(i)}) \setminus V(G^{(i+1)})$ have distance at most four from y_{i+1}, (3.9) ensures that

$$|V(G^{(i)}) \setminus V(G^{(i+1)})| \le \Gamma^2\ln^6 n. \tag{3.10}$$

Iterating (3.10), we obtain $|V(G^{(0)}) \setminus V(G^{(i+1)})| \le (i+1)\Gamma^2\ln^6 n$, whence $|V(G^{(i+1)})| \ge n - (i+1)\Gamma^2\ln^6 n$. □

The following claim resembles the proof of [5, Theorem 1] (where the case $k = \Omega(n)$ is considered).

Claim 3.11. *Let $\mathscr{D}^{(i)}(x) = \{x \in V_+(G^{(i)})\}$ and let*

$$L^{(i)} = \frac{1}{|V(G^{(i)})|}\sum_{x\in V(G^{(i)})} \ln\mathbb{P}\left[\mathscr{D}^{(i)}(x)\right]. \tag{3.11}$$

Then

$$L^{(i)} \ge \frac{|F(G^{(i)})|}{|V(G^{(i)})|}\min_{a\in F(G^{(i)})} \left|\partial_{G^{(i)}} a\right| \ln\left(1-(1-p)^{|\partial_{G^{(i)}} a|-1}\right). \tag{3.12}$$

Proof. For an individual $x \in V(G^{(i)})$ and a test $a \in \partial_{G^{(i)}} x$ let $\mathscr{D}^{(i)}(x,a)$ be the event that there is another individual $z \in \partial_{G^{(i)}} a \setminus \{x\}$ such that $\boldsymbol{\chi}_z = 1$. Then for every $x \in V(G^{(i)})$ we have

$$\mathbb{P}\left[\mathscr{D}^{(i)}(x)\right] = \mathbb{P}\left[\bigcap_{a\in\partial_{G^{(i)}} x} \mathscr{D}^{(i)}(x,a)\right]. \tag{3.13}$$

Furthermore, the events $\mathscr{D}^{(i)}(x,a)$ are increasing with respect to $\boldsymbol{\chi}$. Therefore, (3.13) and the FKG inequality imply

$$\mathbb{P}\left[\mathscr{D}^{(i)}(x)\right] \ge \prod_{a\in\partial_{G^{(i)}} x} \mathbb{P}\left[\mathscr{D}^{(i)}(x,a)\right]. \tag{3.14}$$

Moreover, because each entry of $\boldsymbol{\chi}$ is one with probability p independently, we obtain

$$\mathbb{P}\left[\mathscr{D}^{(i)}(x,a)\right] = 1-(1-p)^{|\partial_{G^{(i)}} a|-1} \tag{3.15}$$

Finally, combining (3.13)–(3.15), we obtain

$$\begin{aligned} |V(G^{(i)})|L^{(i)} &\ge \sum_{x\in V(G^{(i)})}\sum_{a\in F(G^{(i)})} \mathbf{1}\{a\in\partial_{G^{(i)}} x\}\ln\left(1-(1-p)^{|\partial_{G^{(i)}} a|-1}\right) \\ &= \sum_{a\in F(G^{(i)})} \left|\partial_{G^{(i)}} a\right|\ln\left(1-(1-p)^{|\partial_{G^{(i)}} a|-1}\right) \ge |F(G^{(i)})|\min_{a\in F(G^{(i)})}\left|\partial_{G^{(i)}} a\right|\ln\left(1-(1-p)^{|\partial_{G^{(i)}} a|-1}\right), \end{aligned}$$

as claimed. □

As a final preparation for the proof of Lemma 3.9 we need the following estimate.

Claim 3.12. *The function $z \in (0,\infty) \mapsto z\ln(1-(1-p)^{z-1})$ attains its minimum at $z = \left(1+O(n^{-\Omega(1)})\right)\ln(2)/p$.*

Proof. We consider three separate cases.

Case 1: $z = o(1/p)$**:** we obtain

$$\begin{aligned} z\ln\left(1-(1-p)^{z-1}\right) &= z\ln\left(1-\exp\left(-pz+O(p^2z)\right)\right) = z\ln\left(1-\left(1-pz+O(p^2z^2)\right)\right) \\ &= \frac{z}{\ln}(zp+O(zp)^2) = o(1/p). \end{aligned} \tag{3.16}$$

Case 2: $z = \omega(1/p)$**:** we find

$$z\ln\bigl(1-(1-p)^{z-1}\bigr) = z\ln\bigl(1-\exp\bigl(-pz+O(p^2z)\bigr)\bigr) = -z\bigl(\exp(-pz)+O\bigl(\exp(-2pz)\bigr)\bigr)$$
$$= -\frac{1}{p}pz\bigl(\exp\bigl(-pz\bigr)+\exp\bigl(-2pz\bigr)\bigr) = o(1/p). \tag{3.17}$$

Case 3: $z = \Theta(1/p)$**:** letting $d = zp$, we obtain

$$z\ln\bigl(1-(1-p)^{z-1}\bigr) = \frac{d}{p}\ln\bigl(1-\exp\bigl(-d+O(p)\bigr)\bigr) = \frac{d}{p}\ln\bigl(1-\exp(-d)\bigr)+O(1). \tag{3.18}$$

Since the strictly convex function $d \in (0,\infty) \mapsto d\ln(1-\exp(-d))$ attains its minimum at $d = \ln 2$, (3.18) dominates (3.16) and (3.17). Thus, the minimiser reads $z = \ln(2)/p + O(p^{-1/2})$. □

Proof of Lemma 3.9. Combining Claims 3.11 and 3.12, we see that for all test designs G with $m \le (1-\varepsilon)m_{\inf}$ and for all $i \le N$,

$$L^{(i)} \ge -\bigl(1+O(n^{-\Omega(1)})\bigr)\frac{|F(G^{(i)})|\ln^2 2}{|V(G^{(i)})|p} \ge -\bigl(1+O(n^{-\Omega(1)})\bigr)\frac{m\ln^2 2}{|V(G^{(i)})|p}.$$

Hence, Claim 3.10, (3.3) and the choice $p = (k+\sqrt{k}\ln n)/n$ imply that for all $i \le N$,

$$L^{(i)} \ge -\bigl(1+O(n^{-\Omega(1)})\bigr)\frac{m\ln^2 2}{(n-N\Delta^2\ln^6 n)p} \ge -\bigl(1+O(n^{-\Omega(1)})\bigr)\frac{m\ln^2 2}{n^\theta}. \tag{3.19}$$

Further, combining the definition (3.11) of $L^{(i)}$ with (3.19), we conclude that for every $i \le N$ there exists an individual $y_i \in V(G^{(i)})$ such that

$$\mathbb{P}\Bigl[y_i \in V_+(G^{(i)})\Bigr] = \mathbb{P}\Bigl[\mathscr{D}^{(i)}(y_i)\Bigr] \ge \exp\Bigl(L^{(i)}\Bigr) \ge \exp\biggl(-\bigl(1+O(n^{-\Omega(1)})\bigr)\frac{m\ln^2(2)}{n^\theta}\biggr),$$

which implies the assertion. □

Lemma 3.9 implies the following bound on $|V_{0+}(G^*,\boldsymbol{\chi})|$, $|V_{1+}(G^*,\boldsymbol{\chi})|$.

Corollary 3.13. *There exists $n_0 = n_0(\varepsilon,\theta,\xi)$ such that for all $n > n_0$ and all $G = G_{n,m}$ with $m \le (1-\varepsilon)m_{\inf}$ we have*

$$\mathbb{P}\bigl[\bigl|V_{0+}(G^*,\boldsymbol{\chi})\bigr| \wedge \bigl|V_{1+}(G^*,\boldsymbol{\chi})\bigr| < \ln^4 n\bigr] < \varepsilon/8.$$

Proof. We observe that $V_+(G^{(i)},\boldsymbol{\chi}) \subset V_+(G^*,\boldsymbol{\chi})$ for all $i \le N$ because by construction for any individual $x \in V(G^{(i)})$ every test $a \in \partial_{G^*}x$ of G^* that x belongs to is still present in $G^{(i)}$. Consequently, we obtain the bound

$$\mathbb{P}\bigl[x \in V_+(G^*)\bigr] \ge \mathbb{P}\Bigl[x \in V(G^{(i)})\Bigr] \qquad \text{for all } i \in [N],\ x \in V(G^*). \tag{3.20}$$

Combining (3.20) with Lemma 3.9 we obtain

$$\mathbb{P}\Bigl[y^{(i)} \in V_+(G^*)\Bigr] \ge \exp\Bigl(-\ln^2(2)n^{-\theta}m - 1\Bigr) \ge \exp\Bigl(-(1-\varepsilon)\ln^2(2)n^{-\theta}m_{\inf} - 1\Bigr) \qquad \text{for all } i \in [N].$$

Hence, recalling the definition of $m_{\inf}$ from (1.3), we obtain

$$\mathbb{P}\Bigl[y^{(i)} \in V_+(G^*)\Bigr] \ge \exp\bigl(-(1-\varepsilon)\theta\ln(n) - 1\bigr) = n^{(\varepsilon-1)\theta}/\mathrm{e}. \qquad \text{for all } i \in [N]. \tag{3.21}$$

Since the entry $\boldsymbol{\chi}_{y^{(i)}}$ is independent of the event $\{y^{(i)} \in V_+(G^*)\}$, the definitions (3.1) of $V_{0+}(G^*,\boldsymbol{\chi})$ and $V_{1+}(G^*,\boldsymbol{\chi})$ and (3.21) yield

$$\mathbb{P}\Bigl[y^{(i)} \in V_{0+}(G^*,\boldsymbol{\chi})\Bigr] \ge (1-p)\cdot\frac{n^{(\varepsilon-1)\theta}}{\mathrm{e}} \ge \frac{n^{\varepsilon\theta-1}}{3}, \quad \mathbb{P}\Bigl[y^{(i)} \in V_{1+}(G^*,\boldsymbol{\chi})\Bigr] \ge p\cdot\frac{n^{(\varepsilon-1)\theta}}{\mathrm{e}} \ge \frac{n^{\varepsilon\theta-1}}{3} \quad \text{for all } i \in [N],$$

provided n is sufficiently large. Therefore, recalling $N = \lceil n^{1-\xi}\rceil$ we obtain for large enough n,

$$\mathbb{E}\bigl|\{y^{(1)},\ldots,y^{(N)}\} \cap V_{0+}(G^*,\boldsymbol{\chi})\bigr| \ge n^{\varepsilon\theta-\xi}/3, \qquad \mathbb{E}\bigl|\{y^{(1)},\ldots,y^{(N)}\} \cap V_{1+}(G^*,\boldsymbol{\chi})\bigr| \ge n^{\varepsilon\theta-\xi}/3. \tag{3.22}$$

Further, because the pairwise distances of $y^{(1)},\dots,y^{(N)}$ in G^* exceed four, the events $\{y^{(i)} \in V_{0+}(G^*,\boldsymbol{\chi})\}_{i\le N}$ are mutually independent. So are the events $\{y^{(i)} \in V_{1+}(G^*,\boldsymbol{\chi})\}_{i\le N}$. Finally, since (3.3) ensures that $\varepsilon\theta-\xi>0$, (3.22) and the Chernoff bound yield

$$\mathbb{P}\left[\left|\{y^{(1)},\dots,y^{(N)}\}\cap V_{0+}(G^*,\boldsymbol{\chi})\right|\le \ln^2 n\right]\le \mathbb{P}\left[\mathrm{Bin}(N,n^{\varepsilon\theta-1}/3)\le \ln^2 n\right]\le \exp(-n^{\Omega(1)}),$$
$$\mathbb{P}\left[\left|\{y^{(1)},\dots,y^{(N)}\}\cap V_{1+}(G^*,\boldsymbol{\chi})\right|\le \ln^2 n\right]\le \mathbb{P}\left[\mathrm{Bin}(N,n^{\varepsilon\theta-1}/3)\le \ln^2 n\right]\le \exp(-n^{\Omega(1)}),$$

whence the assertion is immediate. □

Proof of Proposition 3.4. Suppose that $n>n_0(\varepsilon,\theta,\xi)$ is large enough and let $G=G_{n,m}$ be a test design with $m\le(1-\varepsilon)m_{\inf}$ tests. If for every test $a\in F_m$ of degree $|\partial_G a|>\Gamma$ we have $|\partial_G a\cap V_1(G,\boldsymbol{\chi})|\ge 2$, then $V_{0+}(G,\boldsymbol{\chi})=V_{0+}(G^*,\boldsymbol{\chi})$ and $V_{1+}(G,\boldsymbol{\chi})=V_{1+}(G^*,\boldsymbol{\chi})$. Therefore, the assertion is an immediate consequence of Lemma 3.6, Lemma 3.7 and Corollary 3.13. □

3.3. **Proof of Proposition 3.5.** Given $\varepsilon>0$ and $\ln(2)/(1+\ln(2))\le\theta<\theta'<1$ we choose a large enough $n_0=n_0(\varepsilon,\theta,\theta')$ and assume that $n>n_0$. Furthermore, let G be a test design with $m\le(1-\varepsilon)m_{\inf}(n,\theta)$ for the purpose of identifying $k=\lceil n^\theta\rceil$ infected individuals. Starting from the test design G infection for density θ we are going to construct a random test design $\boldsymbol{G}'$ for infection density θ' with the same number m of tests as G. The following lemma fixes the order of $\boldsymbol{G}'$.

Lemma 3.14. *There exists an integer $n^{\theta/\theta'}/2\le n'\le 2n^{\theta/\theta'}\wedge n$ such that $k'=\lceil n'^{\theta'}\rceil=k$.*

Proof. Let $n''=\lceil n^{\theta/\theta'}/2\rceil$. Then $(4n'')^{\theta'}>k$ but $n''^{\theta'}<k$ because the function $z\in(1,\infty)\mapsto z^{\theta'}$ has derivative less than one. For the same reason for any integer $n''<N<4n''$ we have $(N+1)^{\theta'}-N^{\theta'}\le 1$ and thus

$$\lceil(N+1)^{\theta'}\rceil-\lceil N^{\theta'}\rceil\le 1.$$

Consequently, there exists an integer $n'\in(n'',4n'')$ such that $\lceil n'^{\theta'}\rceil=k$. □

Given the test design G with individuals $V_n=\{x_1,\dots,x_n\}$ and tests $F_m=\{a_1,\dots,a_m\}$ we now construct the test design $\boldsymbol{G}'$ as follows. Choose a subset $V(\boldsymbol{G}')\subset V_n$ of n' individuals uniformly at random. Then $\boldsymbol{G}'$ is the subgraph that G induces on $V(\boldsymbol{G}')\cup F_m$. Thus, $\boldsymbol{G}'$ has the same tests as G but we simply leave out from every test the individuals that do not belong to the random subset $V(\boldsymbol{G}')$. Let $\boldsymbol{\tau}\in\{0,1\}^{V(\boldsymbol{G}')}$ be a random vector of Hamming weight k and let $\hat{\boldsymbol{\tau}}\in\{0,1\}^{F_m}$ be the induced vector of tests results

$$\hat{\boldsymbol{\tau}}_a=\max_{x\in\partial_{\boldsymbol{G}'}a}\boldsymbol{\tau}_x\qquad\qquad(a\in F_m).$$

Lemma 3.15. *For any integer $t>0$ we have $\mathbb{P}[Z_k(G,\hat{\boldsymbol{\sigma}})\ge t]\ge\mathbb{P}\left[Z_k(\boldsymbol{G}',\hat{\boldsymbol{\tau}})\ge t\right]$.*

Proof. The choice of n' ensures that $k'=\lceil n'^{\theta'}\rceil=k$. Therefore, the random sets $\{x\in V:\boldsymbol{\sigma}_x=1\}$ and $\{x\in V(\boldsymbol{G}'):\boldsymbol{\tau}_x=1\}$ are identically distributed. Indeed, we obtain the latter by first choosing the random subset $V(\boldsymbol{G}')$ of V_n and then choosing a random subset of $V(\boldsymbol{G}')$ size k. Clearly, this two-step procedure is equivalent to just choosing a random subset of size k out of V_n. Hence, we can couple $\boldsymbol{\sigma},\boldsymbol{\tau}$ such that the sets $\{x\in V:\boldsymbol{\sigma}_x=1\}$, $\{x\in V:\boldsymbol{\tau}_x=1\}$ are identical. Then the construction of $\boldsymbol{G}'$ ensures that the vectors $\hat{\boldsymbol{\sigma}},\hat{\boldsymbol{\tau}}$ coincide as well.

Now consider a vector $\sigma'\in\mathscr{S}_k(\boldsymbol{G}',\hat{\boldsymbol{\tau}})$ that explains the test results. Extend σ' to a vector $\sigma\in\{0,1\}^{V_n}$ by setting $\sigma_x=0$ for all $x\in V_n\setminus V(\boldsymbol{G}')$. Then $\sigma\in\mathscr{S}_k(\boldsymbol{G},\hat{\boldsymbol{\sigma}})$. Hence, $Z_k(G,\hat{\boldsymbol{\sigma}})\ge Z_k(\boldsymbol{G}',\hat{\boldsymbol{\tau}})$. □

Proof of Proposition 3.5. Lemma 3.15 shows that for any $t>0$,

$$\mathbb{P}[Z_k(G,\hat{\boldsymbol{\sigma}})\ge t]\ge\mathbb{P}\left[Z_k(\boldsymbol{G}',\hat{\boldsymbol{\tau}})\ge t\right]=\mathbb{E}\left[\mathbb{P}\left[Z_k(\boldsymbol{G}',\hat{\boldsymbol{\tau}})\ge t\mid\boldsymbol{G}'\right]\right].$$

Consequently, there exists an outcome G' of $\boldsymbol{G}'$ such that $\mathbb{P}[Z_k(G,\hat{\boldsymbol{\sigma}})\ge t]\ge\mathbb{P}\left[Z_k(G',\hat{\boldsymbol{\tau}})\ge t\right]$. □

4. The non-adaptive group testing algorithm SPIV

In this section we describe the new test design and the associated inference algorithm SPIV for Theorem 1.2. Throughout we fix $\theta\in(0,1)$ and $\varepsilon>0$ and we tacitly assume that $n>n_0(\varepsilon,\theta)$ is large enough for the various estimates to hold.

4.1. **The random bipartite graph and the DD algorithm.** To motivate the new test design we begin with a brief discussion of the plain random design used in prior work and the best previously known inference algorithm DD [11, 22]. At first glance a promising candidate test design appears to be a random bipartite graph with one vertex class $V_n = \{x_1, \dots, x_n\}$ representing individuals and the other class $F_m = \{a_1, \dots, a_m\}$ representing tests. Indeed, two slightly different random graph models have been proposed [6]. First, in the *Bernoulli model* each V_n–F_m-edge is present with a certain probability (the same for every pair) independently of all others. However, due to the relatively heavy lower tail of the degrees of the individuals, this test design turns out to be inferior to a second model where the degrees of the individuals are fixed. Specifically, in the Δ-*out model* every individual independently joins an equal number of Δ tests drawn uniformly at random without replacement [29].

Clearly, in order to extract the maximum amount of information Δ should be chosen so as to maximise the entropy of the vector of test results. Specifically, since the average test degree equals $\Delta n/m$ and a total of k individuals are infected, the average number of infected individuals per test comes to $\Delta k/m$. Indeed, since $k \sim n^\theta$ for a fixed $\theta < 1$, the number of infected individuals in test a_i can be well approximated by a Poisson variable. Therefore, setting

$$\Delta \sim \frac{m}{k} \ln 2 \tag{4.1}$$

ensures that about half the tests are positive w.h.p.

With respect to the performance of the Δ-out model, [11, Theorem 1.1] implies together with Theorem 1.1 that this simple construction is information-theoretically optimal. Indeed, $m = (1+\varepsilon+o(1))m_{\mathrm{inf}}$ test suffice so that an exponential time algorithm correctly infers the set of infected individuals. Specifically, the algorithm solves a minimum hypergraph vertex cover problem with the individuals as the vertex set and the positive test groups as the hyperedges. For $m = (1+\varepsilon+o(1))m_{\mathrm{inf}}$ the unique optimal solution is precisely the correct set of infected individuals w.h.p. While the worst case NP-hardness of hypergraph vertex cover does not, of course, preclude the existence of an algorithm that is efficient on random hypergraphs, despite considerable efforts no such algorithm has been found. In fact, as we saw in Section 1.4 for a good number of broadly similar inference and optimisation problems on random graphs no efficient information-theoretically optimal algorithms are known.

But for m exceeding the threshold $m_{\mathtt{DD}}$ from (1.2) an efficient greedy algorithm DD correctly recovers $\boldsymbol{\sigma}$ w.h.p. The algorithm proceeds in three steps.

- **DD1:** declare every individual that appears in a negative test uninfected and subsequently remove all negative tests and all individuals that they contain.
- **DD2:** for every remaining (positive) test of degree one declare the individual that appears in the test infected.
- **DD3:** declare all other individuals as uninfected.

The decisions made by the first two steps **DD1**–**DD2** are clearly correct but **DD3** might produce false negatives. Prior to the present work DD was the best known polynomial time group testing algorithm. While DD correctly identifies the set of infected individuals w.h.p. if $m > (1+\varepsilon)m_{\mathtt{DD}}$ [22], the algorithm fails if $m < (1-\varepsilon)m_{\mathtt{DD}}$ w.h.p. [11].

4.2. **Spatial coupling.** The new efficient algorithm SPIV for Theorem 1.2 that gets by with the optimal number $(1+\varepsilon+o(1))m_{\mathrm{inf}}$ of tests comes with a tailor-made test design that, inspired by spatially coupled codes [18, 26, 27], combines randomisation with a superimposed geometric structure. Specifically, we divide both the individuals and the tests into

$$\ell = \lceil \ln^{1/2} n \rceil \tag{4.2}$$

compartments of equal size. The compartments are arranged along a ring and each individual joins an equal number of random tests in the

$$s = \lceil \ln \ln n \rceil = o(\ell) \tag{4.3}$$

topologically subsequent compartments. Additionally, to get the algorithm started we equip the first s compartments with extra tests so that they can be easily diagnosed via the DD algorithm. Then, having diagnosed the initial compartments correctly, SPIV will work its way along the ring, diagnosing one compartment after the other.

To implement this idea precisely we partition the set $V = V_n = \{x_1, \dots, x_n\}$ of individuals into pairwise disjoint subsets $V[1], \dots, V[\ell]$ of sizes $|V[j]| \in \{\lfloor n/\ell \rfloor, \lceil n/\ell \rceil\}$. With each compartment $V[i]$ of individuals we associate a compartment $F[i]$ of tests of size $|F[i]| = m/\ell$ for an integer m that is divisible by ℓ. Additionally, we introduce

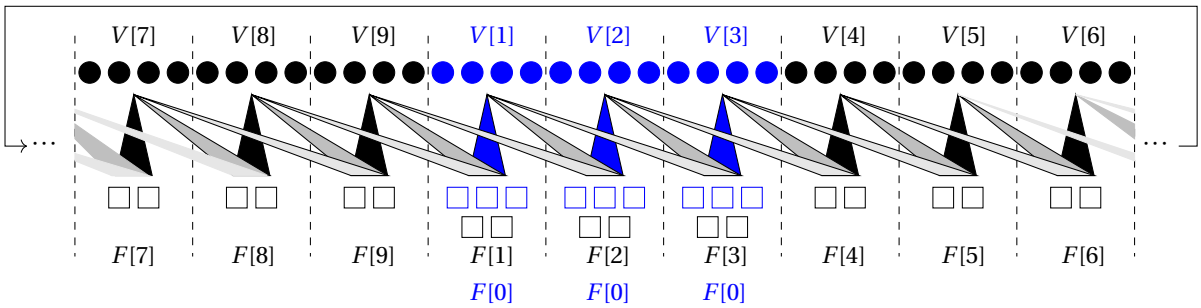

FIGURE 2. The spatially coupled test design with $n = 36, \ell = 9, s = 3$. The individuals in the seed groups $V[1] \cup \cdots \cup V[s]$ (blue) are equipped with additional test $F[0]$ (blue rectangles). The black rectangles represent the tests $F[1] \cup \cdots \cup F[\ell]$.

a set $F[0]$ of $10\lceil (ks/\ell)\ln n\rceil$ extra tests to facilitate the greedy algorithm for diagnosing the first s compartments. Thus, the total number of tests comes to

$$|F[0]| + \sum_{i=1}^{\ell} |F[i]| = (1 + O(s/\ell))m = (1 + o(1))m. \tag{4.4}$$

Finally, for notational convenience we define $V[\ell + i] = V[i]$ and $F[\ell + i] = F[i]$ for $i = 1, \ldots, s$.

The test groups are composed as follows: let

$$k = \lceil n^{\theta} \rceil \qquad \text{and let} \qquad \Delta = \frac{m \ln 2}{k} + O(s) \tag{4.5}$$

be an integer divisible by s; cf. (4.1). Then we construct a random bipartite graph as follows.

SC1: for $i = 1, \ldots, \ell$ and $j = 1, \ldots, s$ every individual $x \in V[i]$ joins Δ/s tests from $F[i + j - 1]$ chosen uniformly at random without replacement. The choices are mutually independent for all individuals x and all j.

SC2: additionally, each individual from $V[1] \cup \cdots \cup V[s]$ independently joins $\lceil 10 \ln(2) \ln n \rceil$ random tests from $F[0]$, drawn uniformly without replacement.

Thus, **SC1** provides that the individuals in compartment $V[i]$ take part in the next s compartments $F[i], \ldots, F[i + s - 1]$ of tests along the ring. Furthermore, **SC2** supplies the tests required by the DD algorithm to diagnose the first s compartments. Figure 2 provides an illustration of the resulting random test design,

From here on the test design produced by **SC1–SC2** is denoted by $\boldsymbol{G}$. Furthermore $\boldsymbol{\sigma} \in \{0,1\}^V$ denotes a uniformly random vector of Hamming weight k, drawn independently of $\boldsymbol{G}$, and $\hat{\boldsymbol{\sigma}} = (\hat{\boldsymbol{\sigma}}_a)_{a \in F[0] \cup \cdots \cup F[\ell]}$ signifies the vector of test results

$$\hat{\boldsymbol{\sigma}}_a = \max_{x \in \partial a} \boldsymbol{\sigma}_x.$$

In addition, let $V_1 = \{x \in V : \boldsymbol{\sigma}_x = 1\}$ be the set of infected individuals and let $V_0 = V \setminus V_1$ be the set of healthy individuals. Moreover, let $F = F[0] \cup F[1] \cup \cdots \cup F[\ell]$ be the set of all tests, let $F_1 = \{a \in F : \hat{\boldsymbol{\sigma}}_a = 1\}$ be the set of all positive tests and let $F_0 = F \setminus F_1$ be the set of all negative tests. Finally, let

$$V_0[i] = V[i] \cap V_0, \qquad V_1[i] = V[i] \cap V_1, \qquad F_0[i] = F[i] \cap F_0, \qquad F_1[i] = F[i] \cap F_1.$$

The following proposition summarises a few basic properties of the test design $\boldsymbol{G}$.

Proposition 4.1. *If $m = \Theta(n^{\theta} \ln n)$ then $\boldsymbol{G}$ enjoys the following properties with probability $1 - o(n^{-2})$.*

(i) The infected individual counts in the various compartments satisfy

$$\frac{k}{\ell} - \sqrt{\frac{k}{\ell}} \ln n \le \min_{i \in [\ell]} |V_1[i]| \le \max_{i \in [\ell]} |V_1[i]| \le \frac{k}{\ell} + \sqrt{\frac{k}{\ell}} \ln n.$$

(ii) For all $i \in [\ell]$ and all $j \in [s]$ the test degrees satisfy

$$\frac{\Delta n}{ms} - \sqrt{\frac{\Delta n}{ms}} \ln n \le \min_{a \in F[i+j-1]} |V[i] \cap \partial a| \le \max_{a \in F[i+j-1]} |V[i] \cap \partial a| \le \frac{\Delta n}{ms} + \sqrt{\frac{\Delta n}{ms}} \ln n.$$

(iii) For all $i \in [\ell]$ the number of negative tests in compartment $F[i]$ satisfies

$$\frac{m}{2\ell} - \sqrt{m}\ln^3 n \le |F_0[i]| \le \frac{m}{2\ell} + \sqrt{m}\ln^3 n.$$

We prove Proposition 4.1 in Section 4.4. Finally, as a preparation for things to come we point out that for any specific individual $x \in V[i]$ and any particular test $a \in F[i+j]$, $j = 0,\dots,s-1$, we have

$$\mathbb{P}[x \in \partial a] = 1 - \mathbb{P}[x \notin \partial a] = 1 - \binom{|F[i+j]|-1}{\Delta/s}\binom{|F[i+j]|}{\Delta/s}^{-1} = \frac{\Delta\ell}{ms} + O\left(\left(\frac{\Delta\ell}{ms}\right)^2\right). \tag{4.6}$$

4.3. **The Spatial Inference Vertex Cover ('SPIV') algorithm.** The SPIV algorithm for Theorem 1.2 proceeds in three phases. The plan of attack is for the algorithm to work its way along the ring, diagnosing one compartment after the other aided by what has been learned about the preceding compartments. Of course, we need to start somewhere. Hence, in its first phase SPIV diagnoses the seed compartments $V[1],\dots,V[s]$.

4.3.1. *Phase 1: the seed.* Specifically, the first phase of SPIV applies the DD greedy algorithm from Section 4.1 to the subgraph of $\boldsymbol{G}$ induced on the individuals $V[1] \cup \dots \cup V[s]$ and the tests $F[0]$. Throughout the vector $\tau \in \{0,1\}^V$ signifies the algorithm's current estimate of the ground truth $\boldsymbol{\sigma}$.

Input: $\boldsymbol{G}$, $\hat{\boldsymbol{\sigma}}$
Output: an estimate of $\boldsymbol{\sigma}$
1 Let $(\tau_x)_{x \in V[1]\cup\dots\cup V[s]} \in \{0,1\}^{V[1]\cup\dots\cup V[s]}$ be the result of applying DD to the tests $F[0]$;
2 Set $\tau_x = 0$ for all individuals $x \in V \setminus (V[1] \cup \dots \cup V[s])$;

Algorithm 1: SPIV, phase 1

The following proposition, whose proof can be found in Section 4.5, summarises the analysis of phase 1.

Proposition 4.2. *W.h.p. the output of* DD *satisfies $\tau_x = \boldsymbol{\sigma}_x$ for all $x \in V[1] \cup \dots \cup V[s]$.*

4.3.2. *Phase 2: enter the ring.* This is the main phase of the algorithm. Thanks to Proposition 4.2 we may assume that the seed has been diagnosed correctly. Now, the programme is to diagnose one compartment after the other, based on what the algorithm learned previously. Hence, assume that we managed to diagnose compartments $V[1],\dots,V[i]$ correctly. How do we proceed to compartment $V[i+1]$?

For a start, we can safely mark as uninfected all individuals in $V[i+1]$ that appear in a negative test. But a simple calculation reveals that this will still leave us with many more than k undiagnosed individuals w.h.p. To be precise, consider the set of uninfected disguised individuals

$$V_{0+}[i+1] = \{x \in V_0[i+1] : \hat{\boldsymbol{\sigma}}_a = 1 \text{ for all } a \in \partial x\},$$

i.e., uninfected individuals that fail to appear in a negative test. In Section 4.6 we prove the following.

Lemma 4.3. *Suppose that $(1+\varepsilon)m_{\mathrm{ad}} \le m = O(n^\theta \ln n)$. Then w.h.p. for all $s \le i < \ell$ we have*

$$|V_{0+}[i+1]| = \left(1 + O\left(n^{-\Omega(1)}\right)\right)\frac{n}{\ell 2^\Delta}.$$

Hence, by the definition (4.5) of Δ for m close to m_{inf} the set $V_{0+}[i+1]$ has size $k^{1+\Omega(1)} \gg k$ w.h.p.

Thus, the challenge is to discriminate between $V_{0+}[i+1]$ and the set $V_1[i+1]$ of actual infected individuals in compartment $i+1$. The key observation is that we can tell these sets apart by counting currently 'unexplained' positive tests. To be precise, for an individual $x \in V[i+1]$ and $1 \le j \le s$ let $\boldsymbol{W}_{x,j}$ be the number of tests in compartment $F[i+j]$ that contain x but that do not contain an infected individual from the preceding compartments $V[1] \cup \dots \cup V[i]$. In formulas,

$$\boldsymbol{W}_{x,j} = \left|\left\{a \in \partial x \cap F[i+j] : \partial a \cap (V_1[1] \cup \dots \cup V_1[i]) = \emptyset\right\}\right|. \tag{4.7}$$

Crucially, the following back-of-the-envelope calculation shows that the mean of this random variable depends on whether x is infected or healthy but disguised.

Infected individuals ($x \in V_1[i+1]$): consider a test $a \in \partial x \cap F[i+j]$, $j = 1,\dots,s$. Because the individuals join tests independently, conditioning on x being infected does not skew the distribution of the individuals from the $s-j$ prior compartments $V[i+j-s+1],\dots,V[i]$ that appear in a. Furthermore, we chose Δ so that for each of these compartments $V[h]$ the expected number of infected individuals that join a has mean $(\ln 2)/s$. Indeed, due to independence it is not difficult to see that $|V_1[h] \cap \partial a|$ is approximately a Poisson variable. Consequently,

$$\mathbb{P}\big[(V_1[i+j-s+1] \cup \dots \cup V_1[i]) \cap \partial a = \emptyset\big] \sim 2^{-(s-j)/s}. \tag{4.8}$$

Hence, because x appears in Δ/s tests $a \in F[i+j]$, the linearity of expectation yields

$$\mathbb{E}\big[\boldsymbol{W}_{x,j} \mid x \in V_1[i+1]\big] \sim 2^{j/s-1}\frac{\Delta}{s}. \tag{4.9}$$

Disguised healthy individuals ($x \in V_{0+}[i+1]$): similarly as above, for any individual $x \in V[i+1]$ and any $a \in \partial x \cap F[i+j]$ the *unconditional* number of infected individuals in a is asymptotically $\mathrm{Po}(\ln 2)$. But given $x \in V_{0+}[i+1]$ we know that a is positive. Thus, $\partial a \setminus \{x\}$ contains at least one infected individual. In effect, the number of positives in a approximately turns into a conditional Poisson $\mathrm{Po}_{\geq 1}(\ln 2)$. Consequently, for test a not to include any infected individual from one of the known compartments $V[h]$, $h = i+j-s+1,\dots,i$, every infected individual in test a must stem from the j yet undiagnosed compartments. Summing up the conditional Poisson and recalling that x appears in Δ/s tests $a \in F[j]$, we thus obtain

$$\mathbb{E}\big[\boldsymbol{W}_{x,j} \mid x \in V_{0+}[i+1]\big] \sim \frac{\Delta}{s}\sum_{t\geq 1}\mathbb{P}\left[\mathrm{Po}_{\geq 1}(\ln 2) = t\right](j/s)^t = (2^{j/s}-1)\frac{\Delta}{s}. \tag{4.10}$$

A first idea to tell $V_{0+}[i+1]$ and $V_1[i+1]$ apart might thus be to simply calculate

$$\boldsymbol{W}_x = \sum_{j=1}^{s-1}\boldsymbol{W}_{x,j} \qquad\qquad (x \in V[i+1]). \tag{4.11}$$

Indeed, (4.9) and (4.10) yield

$$\mathbb{E}[\boldsymbol{W}_x \mid x \in V_1[i+1]] \sim \frac{\Delta}{2\ln 2} = 0.721\dots\,\Delta \qquad \text{whereas} \qquad \mathbb{E}[\boldsymbol{W}_x \mid x \in V_{0+}[i+1]] \sim \frac{\Delta(1-\ln 2)}{\ln 2} = 0.442\dots\,\Delta.$$

But unfortunately a careful large deviations analysis reveals that $\boldsymbol{W}_x$ is not sufficiently concentrated. More precisely, even for $m = (1+\varepsilon+o(1))\,m_{\mathrm{inf}}$ there are as many as $k^{1+\Omega(1)}$ 'outliers' $x \in V_{0+}[i+1]$ whose $\boldsymbol{W}_x$ grows as large as the mean $\Delta/(2\ln 2)$ of actual infected individuals w.h.p.

At second thought the plain sum (4.11) does seem to leave something on the table. While $\boldsymbol{W}_x$ counts all as yet unexplained positive tests equally, not all of these tests reveal the same amount of information. In fact, we should really be paying more attention to 'early' unexplained tests $a \in F[i+1]$ than to 'late' ones $b \in F[i+s]$. For we already diagnosed $s-1$ out of the s compartments of individuals that a draws on, whereas only one of the s compartments that contribute to b has already been diagnosed. Thus, the unexplained test a is a much stronger indication that x might be infected. Consequently, it seems promising to replace $\boldsymbol{W}_x$ by a weighted sum

$$\boldsymbol{W}^\star_x = \sum_{j=1}^{s-1} w_j \boldsymbol{W}_{x,j} \tag{4.12}$$

with $w_1,\dots,w_{s-1} \geq 0$ chosen so as to gauge the amount of information carried by the different compartments.

To find the optimal weights $w_1,\dots,w_{s-1}$ we need to investigate the rate function of $\boldsymbol{W}^\star_x$ given $x \in V_{0+}[i+1]$. More specifically, we should minimise the probability that $\boldsymbol{W}^\star_x$ given $x \in V_{0+}[i+1]$ grows as large as the mean of $\boldsymbol{W}^\star_x$ given $x \in V_1[i+1]$, which we read off (4.9) easily:

$$\mathbb{E}\big[\boldsymbol{W}^\star_x \mid x \in V_1[i+1]\big] \sim \frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j. \tag{4.13}$$

A careful large deviations analysis followed by a Lagrangian optimisation leads to the optimal choice

$$w_j = \ln\frac{(1-2\zeta)2^{j/s-1}(2-2^{j/s})}{(1-(1-2\zeta)2^{j/s-1})(2^{j/s}-1)} \qquad \text{where} \qquad \zeta = 1/s^2. \tag{4.14}$$

The following two lemmas show that with these weights the scores $\boldsymbol{W}_x^\star$ discriminate well between the potential false positives and the infected individuals. More precisely, thresholding $\boldsymbol{W}_x^\star$ we end up misclassifying no more than $o(k)$ individuals x w.h.p.

Lemma 4.4. *Suppose that* $(1+\varepsilon)m_{\mathrm{ad}} \le m = O(n^\theta \ln n)$. *W.h.p. we have*

$$\sum_{s\le i<\ell}\sum_{x\in V_1[i]} \mathbf{1}\left\{\boldsymbol{W}_x^\star < (1-\zeta/2)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1}w_j\right\} \le k\exp\left(-\frac{\Omega(\ln n)}{(\ln\ln n)^4}\right). \tag{4.15}$$

Lemma 4.5. *Suppose that* $(1+\varepsilon)m_{\mathrm{ad}} \le m = O(n^\theta \ln n)$. *W.h.p. we have*

$$\sum_{s\le i<\ell}\sum_{x\in V_{0+}[i]} \mathbf{1}\left\{\boldsymbol{W}_x^\star > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1}w_j\right\} \le k^{1-\Omega(1)}. \tag{4.16}$$

We prove these two lemmas in Sections 4.7 and 4.8.

Lemmas 4.4–4.5 leave us with only one loose end. Namely, calculating the scores $\boldsymbol{W}_x^\star$ requires knowledge of the correct infection status $\boldsymbol{\sigma}_x$ of *all* the individuals $x \in V[1]\cup\cdots\cup V[i]$ from the previous compartments. But since the r.h.s. expressions in (4.15) and (4.16) are non-zero, it is unrealistic to assume that the algorithm's estimates τ_x will consistently match the ground truth $\boldsymbol{\sigma}_x$ beyond the seed compartments. Hoping that the algorithm's estimate will not stray too far, we thus have to make do with the approximate scores

$$W_x^\star(\tau) = \sum_{j=1}^{s-1} w_j W_{x,j}(\tau), \qquad \text{where} \qquad W_{x,j}(\tau) = \left|\left\{a\in\partial x\cap F[i+j-1] : \max_{y\in\partial a\cap(V[1]\cup\cdots V[i])}\tau_y = 0\right\}\right|. \tag{4.17}$$

Hence, phase 2 of SPIV reads as follows.

```
for i = s, ..., ℓ − 1 do
    for x ∈ V[i + 1] do
        if ∃a ∈ ∂x : σ̂_a = 0 then
            τ_x = 0 // classify as uninfected
        else if W*_x(τ) < (1 − ζ)(Δ/s) Σ_{j=1}^{s−1} 2^{j/s−1} w_j then
            τ_x = 0 // tentatively classify as uninfected
        else
            τ_x = 1 // tentatively classify as infected
```

Algorithm 2: SPIV, phase 2.

Since phase 2 of SPIV uses the approximations from (4.17), there seems to be a risk of errors amplifying as we move along. Fortunately, it turns out that errors proliferate only moderately and the second phase of SPIV will misclassify only $o(k)$ individuals. The following proposition summarises the analysis of phase 2.

Proposition 4.6. *Suppose that* $(1+\varepsilon)m_{\mathrm{ad}} \le m = O(k\ln n)$. *W.h.p. the assignment* τ *obtained after steps 1–10 satisfies*

$$\sum_{x\in V}\mathbf{1}\{\tau_x \ne \boldsymbol{\sigma}_x\} \le k\exp\left(-\frac{\ln n}{(\ln\ln n)^6}\right).$$

The proof of Proposition 4.6 can be found in Section 4.9.

4.3.3. *Phase 3: cleaning up.* The final phase of the algorithm rectifies the errors incurred during phase 2. The combinatorial insight that makes this possible is that for $m \ge (1+\varepsilon)m_{\mathrm{inf}}$ every infected individual has at least $\Omega(\Delta)$ positive tests to itself w.h.p. Thus, these tests do not feature a second infected individual. Phase 3 of the algorithm exploits this observation by simply thresholding the number S_x of tests where there is no other infected individual besides potentially x. Thanks to the expansion properties of the graph $\boldsymbol{G}$, each iteration of the thresholding procedure reduces the number of misclassified individuals by at least a factor of three. In effect, after $\ln n$ iterations all individuals will be classified correctly w.h.p. Of course, due to Proposition 4.2 we do not need to reconsider the seed $V[1]\cup\cdots\cup V[s]$.

Let $\tau^{(1)} = \tau$;
for $i = 1, \dots, \lceil \ln n \rceil$ **do**
  For all $x \in V[s+1] \cup \dots \cup V[\ell]$ calculate
    $S_x(\tau^{(i)}) = \sum_{a \in \partial x : \hat{\boldsymbol{\sigma}}_a = 1} \mathbf{1}\left\{\forall y \in \partial a \setminus \{x\} : \tau_y^{(i)} = 0\right\}$;
  Let $\tau_x^{(i+1)} = \begin{cases} \tau_x^{(i)} & \text{if } x \in V[1] \cup \dots \cup V[s], \\ \mathbf{1}\left\{S_x\left(\tau^{(i)}\right) > \ln^{1/4} n\right\} & \text{otherwise} \end{cases}$;
return $\tau^{(\lceil \ln n \rceil)}$

Algorithm 3: `SPIV`, phase 3.

Proposition 4.7. *Suppose that* $(1+\varepsilon)m_{\inf} \le m = O(n^\theta \ln n)$. *W.h.p. for all* $1 \le i \le \lceil \ln n \rceil$ *we have*

$$\sum_{x \in V} \mathbf{1}\{\tau_x^{(i+1)} \neq \boldsymbol{\sigma}_x\} \le \frac{1}{3} \sum_{x \in V} \mathbf{1}\{\tau_x^{(i)} \neq \boldsymbol{\sigma}_x\}.$$

We prove Proposition 4.7 in Section 4.10.

Proof of Theorem 1.2. The theorem is an immediate consequence of Propositions 4.2, 4.6 and 4.7. □

4.4. **Proof of Proposition 4.1.** The number $|V_1[i]|$ of infected individuals in compartment $V[i]$ has distribution $\mathrm{Hyp}(n, k, |V[i]|)$. Since $||V[i]| - n/\ell| \le 1$, (i) is an immediate consequence of the Chernoff bound from Lemma 2.2.

With respect to (ii), we recall from (4.6) that $\mathbb{P}[x \in \partial a] = \frac{\Delta\ell}{ms}(1 + O(\frac{\Delta\ell}{ms}))$. Hence, because the various individuals $x \in V[i]$ join tests independently, the number $|V[i] \cap \partial a|$ of test participants from $V[i]$ has distribution

$$|V[i] \cap \partial a| \sim \mathrm{Bin}(|V[i]|, \Delta\ell/(ms) + O((\Delta\ell/ms)^2)).$$

Since $|V[i]| = n/\ell + O(1)$, assertion (ii) follows from (4.5) and the Chernoff bound from Lemma 2.1.

Coming to (iii), due to part (i) we may condition on $\mathscr{E} = \{\forall i \in [\ell] : |V_1[i]| = k/\ell + O(\sqrt{k/\ell}\ln n)\}$. Hence, with h ranging over the s compartments whose individuals join tests in $F[i]$, (4.6) implies that for every test $a \in F[i]$ the number of infected individuals $|V_1 \cap \partial a|$ is distributed as a sum of independent binomial variables

$$|V_1 \cap \partial a| \sim \sum_h \boldsymbol{X}_h \qquad \text{with} \qquad \boldsymbol{X}_h \sim \mathrm{Bin}\left(V_1[h], \frac{\Delta\ell}{ms} + O\left(\left(\frac{\Delta\ell}{ms}\right)^2\right)\right).$$

Consequently, (4.5) ensures that the event $V_1 \cap \partial a = \emptyset$ has conditional probability

$$\begin{aligned}\mathbb{P}[V_1 \cap \partial a = \emptyset \mid \mathscr{E}] = \prod_h \mathbb{P}[\boldsymbol{X}_h = 0 \mid \mathscr{E}] &= \exp\left[s\left(\frac{k}{\ell} + O\left(\sqrt{\frac{k}{\ell}}\ln n\right)\right)\ln\left(1 - \frac{\Delta\ell}{ms} + O\left(\left(\frac{\Delta\ell}{ms}\right)^2\right)\right)\right] \\ &= \exp\left[-\frac{sk}{\ell}\cdot\frac{\Delta\ell}{ms} + O\left(\sqrt{\frac{k}{\ell}}\cdot\frac{\Delta\ell}{m}\right) + O\left(\frac{sk}{\ell}\cdot\left(\frac{\Delta\ell}{ms}\right)^2\right)\right] = \frac{1}{2} + O(\sqrt{\ell/k}).\end{aligned}$$

Therefore, we obtain the estimate

$$\mathbb{E}[|F_0[i]| \mid \mathscr{E}] = \frac{m}{2\ell} + O(\sqrt{m}\ln n). \tag{4.18}$$

Finally, changing the set of tests that a specific infected individual $x \in V_1[h]$ joins shifts $|F_0[i]|$ by at most Δ (while tinkering with uninfected ones does not change $|F_0[i]|$ at all). Therefore, the Azuma–Hoeffding inequality yields

$$\mathbb{P}[||F_0[i]| - \mathbb{E}[|F_0[i]| \mid \mathscr{E}]| \ge t \mid \mathscr{E}] \le 2\exp\left(-\frac{t^2}{2k\Delta^2}\right) \qquad \text{for any } t > 0. \tag{4.19}$$

Thus, (iii) follows from (4.5), (4.18) and (4.19) with $t = \sqrt{m}\ln^3 n$.

4.5. **Proof of Proposition 4.2.** Let $D = \lceil 10\ln(2)\ln n\rceil$ and recall that $|F[0]| = \lceil 10ks\ln(n)/\ell\rceil$. Since by **SC2** every individual from $\in V[1]\cup\cdots\cup V[s]$ joins D random tests from $F[0]$, in analogy to (4.6) for every $x\in V[1]\cup\cdots\cup V[s]$ and every test $a\in F[0]$ we obtain

$$\mathbb{P}[x\in\partial a] = 1-\mathbb{P}[x\notin\partial a] = 1-\binom{|F[0]|-1}{D}\binom{|F[0]|}{D}^{-1} = \frac{D}{|F[0]|}\left(1+O\left(\frac{D}{|F[0]|}\right)\right) = \frac{\ell\ln 2}{ks}\left(1+O(n^{-\Omega(1)})\right). \tag{4.20}$$

Let $F_1[0]$ be the set of tests $a\in F[0]$ with $\hat{\boldsymbol{\sigma}}_a = 1$.

Lemma 4.8. *W.h.p. the number of positive tests $a\in F[0]$ satisfies $|F_1[0]| = |F[0]|(\frac{1}{2}+O(n^{-\Omega(1)}))$.*

Proof. By Proposition 4.1 we may condition on the event $\mathscr{E}$ that $|V_1[1]\cup\cdots\cup V_1[s]| = \frac{ks}{\ell}(1+O(n^{-\Omega(1)}))$. Hence, (4.20) implies that given $\mathscr{E}$ the expected number of infected individuals in a test $a\in F[0]$ comes to

$$\mathbb{E}[|\partial a\cap V_1|\,|\,\mathscr{E}] = \ln 2+O(n^{-\Omega(1)}). \tag{4.21}$$

Moreover, since individuals join tests independently, $|\partial a\cap V_1|$ is a binomial random variable. Hence, (4.21) implies $\mathbb{P}[\partial a\cap V_1=\emptyset\,|\,\mathscr{E}] = \frac{1}{2}+O(n^{-\Omega(1)})$. Consequently, since $\mathbb{P}[\mathscr{E}] = 1-o(n^{-2})$ by Proposition 4.1,

$$\mathbb{E}|F_1\cap F[0]| = \mathbb{E}|F_1[0]| = \frac{|F[0]|}{2}(1+O(n^{-\Omega(1)})). \tag{4.22}$$

Finally, changing the set ∂x of neighbours of an infected individual can shift $|F_1[0]|$ by at most D. Therefore, the Azuma–Hoeffding inequality implies that

$$\mathbb{P}[||F_1[0]|-\mathbb{E}|F_1[0]||>t]\le 2\exp\left(-\frac{t^2}{2D^2k}\right)\qquad\text{for any } t>0. \tag{4.23}$$

Since $D=O(\ln n)$, combining (4.22) and (4.23) and setting, say, $t=k^{2/3}$ completes the proof. □

As an application of Lemma 4.8 we show that w.h.p. every seed individual x appears in a test $a\in F[0]$ whose other individuals are all healthy.

Corollary 4.9. *W.h.p. every individual $x\in V[1]\cup\cdots\cup V[s]$ appears in a test $a\in F[0]\cap\partial x$ such that $\partial a\setminus\{x\}\subset V_0$.*

Proof. We expose the random bipartite graph induced on $V[1]\cup\cdots\cup V[s]$ and $F[0]$ in two rounds. In the first round we expose $\boldsymbol{\sigma}$ and all neighbourhoods $(\partial y)_{y\in(V[1]\cup\cdots\cup V[s])\setminus\{x\}}$. In the second round we expose ∂x. Let $\boldsymbol{X}$ be the number of negative tests $a\in F[0]$ after the first round. Since x has degree $D=O(\ln n)$, Lemma 4.8 implies that $\boldsymbol{X} = |F[0]|(\frac{1}{2}+O(n^{-\Omega(1)}))$ w.h.p. Furthermore, given $\boldsymbol{X}$ the number of tests $a\in\partial x$ all of whose other individuals are uninfected has distribution $\mathrm{Hyp}(|F[0]|,\boldsymbol{X},D)$. Hence,

$$\mathbb{P}[\forall a\in\partial x: V_1\cap\partial a\setminus\{x\}\neq\emptyset\,|\,\boldsymbol{X}] = \binom{|F[0]|-\boldsymbol{X}}{D}\binom{|F[0]|}{D}^{-1}\le\exp(-D\boldsymbol{X}/|F[0]|). \tag{4.24}$$

Assuming $\boldsymbol{X}/|F[0]| = \frac{1}{2}+O(n^{-\Omega(1)})$ and recalling that $D=\lceil 10\ln(2)\ln n\rceil$, we obtain $\exp(-D\boldsymbol{X}/|F[0]|) = o(1/n)$. Thus, the assertion follows from (4.24) and the union bound. □

Proof of Proposition 4.2. Due to Corollary 4.9 we may assume that for every $x\in V[1]\cup\cdots\cup V[s]$ there is a test $a_x\in F[0]$ such that $\partial a_x\setminus\{x\}\subset V_0$. Hence, recalling the DD algorithm from Section 4.1, we see that the first step **DD1** will correctly identify all healthy individuals $x\in V_0[1]\cup\cdots\cup V_0[s]$. Moreover, the second step **DD2** will correctly classify all remaining individuals $V_1[1]\cup\cdots\cup V_1[s]$ as infected, and the last step **DD3** will be void. □

4.6. **Proof of Lemma 4.3.** Let $\mathscr{E}$ be the event that properties (i) and (iii) from Proposition 4.1 hold; then $\mathbb{P}[\mathscr{E}] = 1-o(n^{-2})$. Moreover, let $\mathfrak{E}$ be the σ-algebra generated by $\boldsymbol{\sigma}$ and the neighbourhoods $(\partial x)_{x\in V_1}$. Then the event $\mathscr{E}$ is $\mathfrak{E}$-measurable while the neighbourhoods $(\partial x)_{x\in V_0}$ of the healthy individuals are independent of $\mathfrak{E}$. Recalling from **SC1** that the individuals $x\in V_0[i]$ choose Δ/s random tests in each of the compartments $F[i+j]$, $0\le j\le s-1$ independently and remembering that $x\in V_{0+}[i]$ iff none of these tests is negative, on $\mathscr{E}$ we obtain

$$\begin{aligned}\mathbb{P}[x\in V_{0+}[i]\,|\,\mathfrak{E}] &= \binom{m/(2\ell)+O(\sqrt{m}\ln^3 n)}{\Delta/s}^s\binom{m/\ell}{\Delta/s}^{-s} = \left(\frac{1+O(m^{-1/2}\ell\ln^3 n)}{2}\right)^{\Delta}\\ &= 2^{-\Delta}+O(m^{-1/2}\Delta\ell\ln^3 n) = 2^{-\Delta}(1+O(n^{-\theta/2}\ln^4 n))\qquad\text{[due to (4.2) and (4.5)]}.\end{aligned} \tag{4.25}$$

Because all $x \in V_0[i]$ choose their neighbourhoods independently, (4.25) implies that the conditional random variable $|V_{0+}[i]|$ given $\mathfrak{E}$ has distribution $\mathrm{Bin}(|V_0[i]|, 2^{-\Delta}(1+O(n^{-\Omega(1)})))$. Therefore, since on $\mathscr{E}$ we have $|V_0[i]| = |V[i]| + O(n^\theta) = n/\ell + O(n^\theta)$, the assertion follows from the Chernoff bound from Lemma 2.1.

4.7. **Proof of Lemma 4.4.** The aim is to estimate the weighted sum $\boldsymbol{W}^\star_x$ for infected individuals $x \in V[i+1]$ with $s \le i < \ell$. These individuals join tests in the s compartments $F[i+j]$, $j \in [s]$. Conversely, for each such j the tests $a \in F[i+j]$ recruit their individuals from the compartments $V[i+j-s+1], \ldots, V[i+j]$. Thus, the compartments preceding $V[i+1]$ that the tests in $F[i+j]$ draw upon are $V[h]$ with $i+j-s < h \le i$. We begin by investigating the set $\mathscr{W}_{i,j}$ of tests $a \in F[i+j]$ without an infected individual from these compartments, i.e.,

$$\mathscr{W}_{i,j} = \big\{a \in F[i+j] : (V_1[1] \cup \cdots \cup V_1[i]) \cap \partial a = \emptyset\big\} = \left\{a \in F[i+j] : \bigcup_{i+j-s+1<h\le i} V_1[h] \cap \partial a = \emptyset\right\}.$$

Claim 4.10. *With probability* $1-o(n^{-2})$ *for all* $s \le i < \ell$, $j \in [s]$ *we have* $|\mathscr{W}_{i,j}| = 2^{-(s-j)/s}\frac{m}{\ell}(1+O(n^{-\Omega(1)}))$.

Proof. We may condition on the event $\mathscr{E}$ that (i) from Proposition 4.1 occurs. To compute the mean of $|\mathscr{W}_{i,j}|$ fix a test $a \in F[i+j]$ and an index $i+j-s < h \le i$. Then (4.6) shows that the probability that a fixed individual $x \in V[h]$ joins a equals $\mathbb{P}[x \in \partial a] = \frac{\Delta\ell}{ms}(1+O(\frac{\Delta\ell}{ms}))$. Hence, the choices (4.2) and (4.5) of Δ and ℓ and the assumption $m = \Theta(k\ln n)$ ensure that

$$\begin{aligned}\mathbb{E}\big[\big|(V_1[i+j-s+1] \cup \cdots \cup V_1[i]) \cap \partial a\big| \mid \mathscr{E}\big] &= (s-j)\left(\frac{\Delta\ell}{ms}\cdot\frac{k}{\ell} + O\left(\frac{\Delta^2 k}{m^2 s^2}\right) + O\left(\frac{\Delta\ell\sqrt{k}\ln n}{ms}\right)\right)\\ &= \frac{s-j}{s}\ln 2 + O(n^{-\Omega(1)}).\end{aligned} \tag{4.26}$$

Since by **SC1** the events $\{x \in \partial a\}_x$ are independent, $|V_1[h] \cap \partial a|$ is a binomial random variable for every h and all these random variables $(|V_1[h] \cap \partial a|)_h$ are mutually independent. Therefore, (4.26) implies that

$$\mathbb{P}\big[(V_1[i+j-s+1] \cup \cdots V_1[i]) \cap \partial a = \emptyset \mid \mathscr{E}\big] = 2^{-(s-j)/s} + O(n^{-\Omega(1)}). \tag{4.27}$$

Hence,

$$\mathbb{E}\big[|\mathscr{W}_{i,j}| \mid \mathscr{E}\big] = \sum_{a\in F[i+j]} \mathbb{P}\big[(V_1[i+j-s+1] \cup \cdots \cup V_1[i]) \cap \partial a = \emptyset \mid \mathscr{E}\big] = \frac{m}{\ell}2^{-(s-j)/s}(1+O(n^{-\Omega(1)})). \tag{4.28}$$

Finally, changing the neighbourhood ∂x of one infected individual $x \in V_1$ can alter $|\mathscr{W}_{i,j}|$ by at most Δ. Therefore, the Azuma–Hoeffing inequality shows that for any $t > 0$,

$$\mathbb{P}\big[\big||\mathscr{W}_{i,j}| - \mathbb{E}[|\mathscr{W}_{i,j}| \mid \mathscr{E}]\big| > t \mid \mathscr{E}\big] \le 2\exp\left(-\frac{t^2}{2k\Delta^2}\right). \tag{4.29}$$

Combining (4.28) and (4.29), applied with $t = \sqrt{m}\ln^2 n$, and taking a union bound on i, j completes the proof. □

As a next step we use Claim 4.10 to estimate the as yet unexplained tests counts $\boldsymbol{W}_{x,j}$ from (4.7).

Claim 4.11. *For all* $s \le i < \ell$, $x \in V_1[i+1]$ *and* $j \in [s]$ *we have*

$$\mathbb{P}\left[\boldsymbol{W}_{x,j} < (1-\varepsilon/2)2^{j/s-1}\Delta/s\right] \le \exp\left(-\frac{\Omega(\ln n)}{(\ln\ln n)^4}\right).$$

Proof. Fix a pair of indices i, j and an individual $x \in V_1[i+1]$. We also condition on the event $\mathscr{E}$ that (i) from Proposition 4.1 occurs. Additionally, thanks to Claim 4.10 we may condition on the event

$$\mathscr{E}' = \left\{|\mathscr{W}_{i,j}| = 2^{-(s-j)/s}\frac{m}{\ell}(1+O(n^{-\Omega(1)}))\right\}.$$

Further, let $\mathfrak{E}$ be the σ-algebra generated by $\boldsymbol{\sigma}$ and by the neighbourhoods $(\partial y)_{y\in V[1]\cup\cdots\cup V[i]}$. Recall from **SC1** that x simply joins Δ/s random tests in compartment $F[i+j]$, independently of all other individuals, and remember from (4.7) that $\boldsymbol{W}_{x,j}$ counts tests $a \in \mathscr{W}_{i+j} \cap \partial x$. Therefore, since the events $\mathscr{E}, \mathscr{E}'$ and the random variable $|\mathscr{W}_{i,j}|$ are $\mathfrak{E}$-measurable while ∂x is independent of $\mathfrak{E}$, given $\mathfrak{E}$ the random variable $\boldsymbol{W}_{x,j}$ has a hypergeometric distribution $\mathrm{Hyp}(m/\ell, |\mathscr{W}_{i,j}|, \Delta/s)$. Thus, the assertion follows from the hypergeometric Chernoff bound from Lemma 2.2 and the choice (4.14) of ζ. □

Proof of Lemma 4.4. Since $\boldsymbol{W}^\star_x = \sum_{j=1}^s w_j \boldsymbol{W}_{x,j}$, the lemma is an immediate consequence of Markov's inequality and Claim 4.11. □

4.8. **Proof of Lemma 4.5.** We need to derive the rate functions of the random variable $\boldsymbol{W}_{x,j}$ that count as yet unexplained tests for $x \in V_{0+}[i+1]$. To this end we first investigate the set of positive tests in compartment $i+j$ that do not contain any infected individuals from the first i compartments. In symbols,

$$\mathscr{P}_{i+1,j} = \bigl\{a \in F_1[i+j] : \partial a \cap (V_1[1] \cup \cdots \cup V_1[i]) = \emptyset\bigr\} \qquad (s \le i < \ell,\ j \in [s]).$$

Claim 4.12. *W.h.p. for all $s \le i < \ell$, $j \in [s]$ we have $|\mathscr{P}_{i+1,j}| = \bigl(1+O(n^{-\Omega(1)})\bigr)\bigl(2^{j/s}-1\bigr)\frac{m}{2\ell}$.*

Proof. We may condition on the event $\mathscr{E}$ that (i) from Proposition 4.1 occurs. As a first step we calculate the probability that $(V_1[i+1] \cup \cdots \cup V_1[i+j]) \cap \partial a \neq \emptyset$ for a specific test $a \in F[i+j]$. To this end we follow the steps of the proof of Claim 4.10. Since by (4.6) a specific individual $x \in V[h]$, $i < h \le i+j$, joins a with probability $\mathbb{P}[x \in \partial a] = (\Delta\ell/(ms))(1+O(\Delta\ell/(ms)))$ and since given $\mathscr{E}$ each compartment $V[h]$ contains $k/\ell + O(\sqrt{k/\ell}\ln n)$ infected individuals, we obtain, in perfect analogy to (4.26),

$$\mathbb{E}\bigl[\bigl|(V_1[i+1] \cup \cdots \cup V_1[i+j]) \cap \partial a\bigr| \mid \mathscr{E}\bigr] = \frac{j}{s}\ln 2 + O(n^{-\Omega(1)}). \tag{4.30}$$

Since the individuals $x \in V[i+1] \cup \cdots \cup V[i+j]$ join tests independently, (4.30) implies that

$$\mathbb{P}\bigl[(V_1[i+1] \cup \cdots \cup V_1[i+j]) \cap \partial a \neq \emptyset \mid \mathscr{E}\bigr] = 1 - 2^{-j/s} + O(n^{-\Omega(1)}). \tag{4.31}$$

Furthermore, we already verified in (4.27) that

$$\mathbb{P}\bigl[(V_1[i+j-s+1] \cup \cdots V_1[i]) \cap \partial a = \emptyset \mid \mathscr{E}\bigr] = 2^{-(s-j)/s} + O(n^{-\Omega(1)}). \tag{4.32}$$

Because the choices for the compartments $V[i+j-s+1] \cup \cdots \cup V[i+j]$ from which a draws its individuals are mutually independent, we can combine (4.31) with (4.32) to obtain

$$\mathbb{P}\left[\bigcup_{i+j-s<h\le i} V_1[h] \cap \partial a = \emptyset \neq \bigcup_{i<h\le i+j} V_1[h] \cap \partial a \mid \mathscr{E}\right] = \frac{2^{j/s}-1}{2} + O(n^{-\Omega(1)}). \tag{4.33}$$

Further, (4.33) implies

$$\mathbb{E}\bigl[|\mathscr{P}_{i+1,j}| \mid \mathscr{E}\bigr] = \mathbb{E}\left[\left|\left\{a \in F_1[i+j] : \bigcup_{h\le i} V_1[h] \cap \partial a = \emptyset \neq \bigcup_{i<h} V_1[h] \cap \partial a\right\}\right| \mid \mathscr{E}\right] = (2^{j/s}-1)\frac{m}{2\ell}\bigl(1+O(n^{-\Omega(1)})\bigr). \tag{4.34}$$

Finally, altering the neighbourhood ∂x of any infected individual can shift $|\mathscr{P}_{i+1,j}|$ by at most Δ. Therefore, the Azuma–Hoeffding inequality implies that

$$\mathbb{P}\bigl[\bigl||\mathscr{P}_{i+1,j}| - \mathbb{E}[|\mathscr{P}_{i+1,j}| \mid \mathscr{E}]\bigr| > t \mid \mathscr{E}\bigr] \le 2\exp\left(-\frac{t^2}{2k\Delta^2}\right). \tag{4.35}$$

Thus, the assertion follows from (4.5), (4.34) and (4.35) by setting $t = \sqrt{m}\ln^2 n$. □

Thanks to Proposition 4.1 (iii) and Lemma 4.12 in the following we may condition on the event

$$\mathscr{U} = \left\{\forall s < i \le \ell, j \in [s] : |F_1[i+j]| = \bigl(1+O(n^{-\Omega(1)})\bigr)\frac{m}{2\ell} \wedge |\mathscr{P}_{i+1,j}| = \bigl(1+O(n^{-\Omega(1)})\bigr)\Bigl(2^{j/s}-1\Bigr)\frac{m}{2\ell}\right\}. \tag{4.36}$$

As a next step we will determine the conditional distribution of $\boldsymbol{W}_{x,j}$ for $x \in V_{0+}[i+1]$ given $\mathscr{U}$.

Claim 4.13. *Let $s < i \le \ell$ and $j \in [s]$. Given $\mathscr{U}$ for every $x \in V_{0+}[i+1]$ we have*

$$\boldsymbol{W}_{x,j} \sim \mathrm{Hyp}\left(\bigl(1+O(n^{-\Omega(1)})\bigr)\frac{m}{2\ell}, \bigl(1+O(n^{-\Omega(1)})\bigr)\Bigl(2^{j/s}-1\Bigr)\frac{m}{2\ell}, \frac{\Delta}{s}\right). \tag{4.37}$$

Proof. By **SC1** each individual $x \in V_{0+}[i+1]$ joins Δ/s positive test from $F[i+j]$, drawn uniformly without replacement. Moreover, by (4.7) given $x \in V_{0+}[i+1]$ the random variable $\boldsymbol{W}_{x,j}$ counts the number of tests $a \in \mathscr{P}_{i+1,j} \cap \partial x$. Therefore, $\boldsymbol{W}_{x,j} \sim \mathrm{Hyp}(|F_1[i+j]|, |\mathscr{P}_{i+1,j}|, \Delta/s)$. Hence, given $\mathscr{U}$ we obtain (4.37). □

The estimate (4.37) enables us to bound the probability that $\boldsymbol{W}^\star_x$ gets 'too large'.

Claim 4.14. *Let*

$$\mathscr{M} = \min \frac{1}{s}\sum_{j=1}^{s-1} \mathbf{1}\left\{z_j \geq 2^{j/s}-1\right\} D_{\mathrm{KL}}\left(z_j \| 2^{j/s}-1\right)$$
$$s.t. \quad \sum_{j=1}^{s-1}\left(z_j - (1-2\zeta)2^{j/s-1}\right) w_j = 0, \qquad z_1,\ldots,z_{s-1} \in [0,1].$$

Then for all $s \leq i < \ell$ and all $x \in V[i+1]$ we have

$$\mathbb{P}\left[\boldsymbol{W}^\star_x > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j \mid \mathscr{U}, x \in V_{0+}[i+1]\right] \leq \exp(-(1+o(1))\mathscr{M}\Delta).$$

Proof. Let $s \leq i < \ell$ and $x \in V_{0+}[i+1]$. Step **SC1** of the construction of $\boldsymbol{G}$ ensures that the random variables $(\boldsymbol{W}_{x,j})_{j\in[s]}$ are independent because the tests in the various compartments $F[i+j]$, $j \in [s]$, that x joins are drawn independently. Therefore, the definition (4.12) of $\boldsymbol{W}^\star_x$ and Lemma 4.13 yield

$$\begin{aligned}\mathbb{P}\left[\boldsymbol{W}^\star_x > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j \mid \mathscr{U}, x \in V_{0+}[i+1]\right] &= \mathbb{P}\left[\sum_{j=1}^{s-1} w_j \boldsymbol{W}_{x,j} \geq \frac{1-2\zeta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j \mid \mathscr{U}, x \in V_{0+}[i+1]\right] \\ &\leq \sum_{y_1,\ldots,y_s=0}^{\Delta} \mathbf{1}\left\{\sum_{j=1}^{s-1} w_j y_j \geq \frac{1-2\zeta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j\right\} \prod_{j=1}^{s-1} \mathbb{P}[\boldsymbol{W}_{x,j} \geq y_j \mid \mathscr{U}, x \in V_{0+}[i+1]]. \end{aligned} \tag{4.38}$$

Further, let

$$\mathscr{Z} = \left\{(z_1,\ldots,z_{s-1}) \in [0,1]^{s-1} : \sum_{j=1}^{s-1}\left(z_j - (1-2\zeta)2^{j/s-1}\right) w_j = 0\right\}.$$

Substituting $y_j = \Delta z_j/s$ in (4.38) and bounding the total number of summands by $(\Delta+1)^s$, we obtain

$$\mathbb{P}\left[\boldsymbol{W}^\star_x > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j \mid \mathscr{U}, x \in V_{0+}[i+1]\right] \leq (\Delta+1)^s \max_{(z_1,\ldots,z_s)\in\mathscr{Z}} \prod_{j=1}^{s-1} \mathbb{P}[\boldsymbol{W}_{x,j} \geq \Delta z_j/s \mid \mathscr{U}, x \in V_{0+}[i+1]]. \tag{4.39}$$

Moreover, Claim 4.13 and the Chernoff bound from Lemma 2.2 yield

$$\mathbb{P}[\boldsymbol{W}_{x,j} \geq \Delta z_j/s \mid \mathscr{U}, x \in V_{0+}[i+1]] \leq \exp\left(-\mathbf{1}\left\{z_j \geq p_j\right\}\frac{\Delta}{s} D_{\mathrm{KL}}\left(z_j \| p_j\right)\right) \qquad \text{where } p_j = 2^{j/s}-1+O(n^{-\Omega(1)}).$$

Consequently, since (4.5) and the assumption $m = \Theta(k\ln n)$ ensure that $\Delta = \Theta(\ln n)$, we obtain

$$\mathbb{P}[\boldsymbol{W}_{x,j} \geq \Delta z_j/s \mid \mathscr{U}, x \in V_{0+}[i+1]] \leq \exp\left(-\mathbf{1}\left\{z_j \geq 2^{j/s}-1\right\}\frac{\Delta}{s} D_{\mathrm{KL}}\left(z_j \| 2^{j/s}-1\right) + O(n^{-\Omega(1)})\right). \tag{4.40}$$

Finally, the assertion follows from (4.39) and (4.40). □

As a next step we solve the optimisation problem $\mathscr{M}$ from Claim 4.14.

Claim 4.15. *We have $\mathscr{M} = 1-\ln 2 + O(\ln(s)/s)$.*

Proof. Fixing an auxiliary parameter $\delta \geq 0$ we set up the Lagrangian

$$\mathscr{L}_\delta(z_1,\ldots,z_s,\lambda) = \sum_{j=1}^{s-1}\left(\mathbf{1}\left\{z_j \geq 2^{j/s}-1\right\} + \delta\mathbf{1}\left\{z_j < 2^{j/s}-1\right\}\right) D_{\mathrm{KL}}\left(z_j \| 2^{j/s}-1\right) + \frac{\lambda}{s}\sum_{j=1}^{s-1} w_j\left(z_j - (1-2\zeta)2^{j/s-1}\right).$$

The partial derivatives come out as

$$\frac{\partial\mathscr{L}_\delta}{\partial\lambda} = -\frac{1}{s}\sum_{j=1}^{s-1}((1-2\zeta)2^{j/s-1} - z_j)\, w_j, \quad \frac{\partial\mathscr{L}_\delta}{\partial z_j} = -\lambda w_j + \left(\mathbf{1}\left\{z_j \geq 2^{j/s}-1\right\} + \delta\mathbf{1}\left\{z_j < 2^{j/s}-1\right\}\right)\ln\frac{z_j(2-2^{j/s})}{(1-z_j)(2^{j/s}-1)}.$$

Set $z^*_j = (1-2\zeta)2^{j/s-1}$ and $\lambda^* = 1$. Then clearly

$$\frac{\partial\mathscr{L}_\delta}{\partial\lambda}\Big|_{\lambda^*, z^*_1,\ldots,z^*_{s-1}} = 0. \tag{4.41}$$

Moreover, the choice (4.14) of ζ guarantees that $z_j^* \geq 2^{j/s}-1$. Hence, by the choice (4.14) of the weights w_j,

$$\frac{\partial \mathscr{L}_\delta}{\partial z_j}\Big|_{\lambda^*,z_1^*,\dots,z_{s-1}^*} = 0. \tag{4.42}$$

Since $\mathscr{L}_\delta(y_1,\dots,y_s,\lambda)$ is strictly convex in $z_1,\dots,z_s$ for every $\delta>0$, (4.41)–(4.42) imply that $\lambda^*, z_1^*,\dots,z_{s-1}^*$ is a global minimiser. Furthermore, since this is true for any $\delta>0$ and since $z_j^* \geq 2^{j/s}-1$, we conclude that $(z_1^*,\dots,z_{s-1}^*)$ is an optimal solution to the minimisation problem $\mathscr{M}$. Hence,

$$\mathscr{M} = \frac{1}{s}\sum_{j=1}^{s-1} D_{\mathrm{KL}}\left(z_j^* \| 2^{j/s}-1\right) = \frac{1}{s}\sum_{j=1}^{s-1} D_{\mathrm{KL}}\left((1-2\zeta)2^{j/s-1}\|2^{j/s}-1\right). \tag{4.43}$$

Since

$$\frac{\partial}{\partial\alpha} D_{\mathrm{KL}}\left((1-2\alpha)2^{z-1}\|2^z-1\right) = 2^z\left[-z\ln(2)+\ln(1-2^{z-1}+\alpha 2^z)-\ln(1-2^{z-1})-\ln(1-2\alpha)+\ln(2^z-1)\right],$$

we obtain $\frac{\partial}{\partial\alpha} D_{\mathrm{KL}}\left((1-2\alpha)2^{z-1}\|2^z-1\right) = O(\ln s)$ for all $z = 1/s,\dots,(s-1)/s$ and $\alpha\in[0,2\zeta]$. Combining this bound with (4.43), we arrive at the estimate

$$\mathscr{M} = O(\zeta\ln s) + \frac{1}{s}\sum_{j=1}^{s-1} D_{\mathrm{KL}}\left(2^{j/s-1}\|2^{j/s}-1\right). \tag{4.44}$$

Additionally, the function $f: z\in[0,1]\mapsto D_{\mathrm{KL}}\left(2^{z-1}\|2^z-1\right)$ is strictly decreasing and convex. Indeed,

$$f'(z) = \frac{2^{z-1}\ln 2}{2^z-1}\left((2^z-1)\ln\left(\frac{2^z}{2^z-1}\right)-1\right), \qquad f''(z) = \left(2^{z-1}\ln^2 2\right)\left(\ln\left(\frac{2^z}{2^z-1}\right)+\frac{2-2^z}{(2^z-1)^2}\right).$$

The first derivative is negative because $2^{z-1}/(2^z-1)>0$ while $(2^z-1)\ln(2^z/(2^z-1))<1$ for all $z\in(0,1)$. Moreover, since evidently $f''(z)>0$ for all $z\in(0,1)$, we obtain convexity. Further, l'Hôpital's rule yields

$$D_{\mathrm{KL}}\left(2^{1/s-1}\|2^{1/s}-1\right) = O(\ln s).$$

As a consequence, we can approximate the sum (4.44) by an integral and obtain

$$\begin{aligned}\mathscr{M} &= O(\ln(s)/s) + \int_0^1 D_{\mathrm{KL}}\left(2^{z-1}\|2^z-1\right)\mathrm{d}z\\ &= O(\ln(s)/s) + \frac{2(1-z)\ln^2(2)+2^z\ln 2^z+(1-2^z)\ln(2^z-1)}{2\ln 2}\Big|_{z=0}^{z=1} = 1-\ln(2)+O(\ln(s)/s),\end{aligned}$$

as claimed. □

Proof of Lemma 4.5. Fix $s\leq i<\ell$ and let $\boldsymbol{X}_i$ be the number of $x\in V_{0+}[i]$ such that $\boldsymbol{W}_x^\star > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1}2^{j/s-1}w_j$. Also recall that Proposition 4.1 (iii) and Claim 4.12 imply that $\mathbb{P}[\mathscr{U}] = 1-o(1)$. Combining Lemma 4.3 with Claims 4.14 and 4.15, we conclude that

$$\mathbb{E}[\boldsymbol{X}_i \mid \mathscr{U}] \leq \left(1+O\left(n^{-\Omega(1)}\right)\right)2^{-\Delta} n\exp(-(1-\ln(2)+o(1))\Delta) = \exp\left(\ln n-(1+o(1))\Delta\right). \tag{4.45}$$

Recalling the definition (4.5) of Δ and using the assumption that $m\geq(1+\varepsilon)m_{\mathrm{ad}}$ for a fixed $\varepsilon>0$, we obtain $\Delta\geq(1-\theta+\Omega(1))\ln n$. Combining this estimate with (4.45), we find

$$\mathbb{E}[\boldsymbol{X}_i \mid \mathscr{U}] \leq n^{\theta-\Omega(1)}. \tag{4.46}$$

Finally, the assertion follows from (4.46) and Markov's inequality. □

4.9. **Proof of Proposition 4.6.** The following lemma establishes an expansion property of $\boldsymbol{G}$. Specifically, if T is a small set of individuals, then there are few individuals x that share many tests with another individual from T.

Lemma 4.16. *Suppose that $m=\Theta(n^\theta\ln n)$. W.h.p. for any set $T\subset V$ of size at most $\exp(-\ln^{7/8}n)k$ we have*

$$\left|\left\{x\in V: \sum_{a\in\partial x\setminus F[0]}\mathbf{1}\{T\cap\partial a\setminus\{x\}\neq\emptyset\}\geq\ln^{1/4}n\right\}\right| \leq \frac{|T|}{3}.$$

Proof. Fix a set $T \subset V$ of size $t = |T| \le \exp(-\ln^{7/8} n)k$, a set $R \subset V$ of size $r = \lceil t/3 \rceil$ and let $\gamma = \lceil \ln^{1/4} n \rceil$. Furthermore, let $U \subset F[1] \cup \cdots \cup F[\ell]$ be a set of tests of size $\gamma r \le u \le \Delta t$. Additionally, let $\mathscr{E}(R,T,U)$ be the event that every test $a \in U$ contains two individuals from $R \cup T$. Then

$$\mathbb{P}\left[R \subset \left\{x \in V : \sum_{a \in \partial x \setminus F[0]} \mathbf{1}\{T \cap \partial a \setminus \{x\} \neq \emptyset\} \ge \gamma\right\}\right] \le \mathbb{P}\left[\mathscr{E}(R,T,U)\right]. \tag{4.47}$$

Hence, it suffices to estimate $\mathbb{P}[\mathscr{E}(R,T,U)]$.

Given a test $a \in U$ there are at most $\binom{r+t}{2}$ way to choose two individuals $x_a, x_a' \in R \cup T$. Moreover, (4.6) shows that the probability of the event $\{x_a, x_a' \in \partial a\}$ is bounded by $(1+o(1))(\Delta\ell/(ms))^2$. Therefore,

$$\mathbb{P}\left[\mathscr{E}(R,T,U)\right] \le \left[\binom{r+t}{2}\left(\frac{(1+o(1))\Delta\ell}{ms}\right)^2\right]^u.$$

Consequently, the event $\mathscr{E}(t,u)$ that there exist sets R,T,U of sizes $|R| = r = \lceil t/3 \rceil, |T| = t, |U| = u$ such that $\mathscr{E}(R,T,U)$ occurs has probability

$$\mathbb{P}\left[\mathscr{E}(t,u)\right] \le \binom{n}{r}\binom{n}{t}\binom{m}{u}\left[\binom{r+t}{2}\left(\frac{(1+o(1))\Delta\ell}{ms}\right)^2\right]^u.$$

Hence, the bounds $\gamma t/3 \le \gamma r \le u \le \Delta t$ yield

$$\begin{aligned}\mathbb{P}\left[\mathscr{E}(t,u)\right] &\le \binom{n}{t}^2\binom{m}{u}\left[\binom{2t}{2}\left(\frac{(1+o(1))\Delta\ell}{ms}\right)^2\right]^u \le \left(\frac{\mathrm{e}n}{t}\right)^{2t}\left(\frac{2\mathrm{e}\Delta^2\ell^2t^2}{ms^2u}\right)^u \\ &\le \left[\left(\frac{\mathrm{e}n}{t}\right)^{3/\gamma}\frac{6\mathrm{e}\Delta^2\ell^2t}{\gamma ms^2}\right]^u \le \left[\left(\frac{\mathrm{e}n}{t}\right)^{3/\gamma}\cdot\frac{t\ln^4 n}{m}\right]^u \qquad \text{[due to (4.2), (4.5)].}\end{aligned}$$

Further, since $\gamma = \Omega(\ln^{1/4} n)$ and $m = \Omega(k\ln n)$ while $t \le \exp(-\ln^{7/8} n)k$, we obtain $\mathbb{P}[\mathscr{E}(t,u)] \le \exp(-u\sqrt{\ln n})$. Thus,

$$\sum_{\substack{1 \le t \le k^{1-\alpha} \\ \gamma t/3 \le u \le \Delta t}} \mathbb{P}\left[\mathscr{E}(t,u)\right] \le \sum_{1 \le u \le \Delta t} u\exp(-u\sqrt{\ln n}) = o(1). \tag{4.48}$$

Finally, the assertion follows from (4.47) and (4.48). □

Proof of Proposition 4.6. With τ the result of steps 1–10 of SPIV let $\mathscr{M}[i] = \{x \in V[i] : \tau_x \neq \boldsymbol{\sigma}_x\}$ be the set of misclassified individuals in compartment $V[i]$. Proposition 4.2 shows that w.h.p. $\mathscr{M}[i] = \emptyset$ for all $i \le s$. Further, we claim that for every $s \le i < \ell$ and any individual $x \in \mathscr{M}[i+1]$ one of the following three statements is true.

M1: $x \in V_1[i+1]$ and $\boldsymbol{W}_x^\star < (1-\zeta/2)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1}w_j$,

M2: $x \in V_{0+}[i+1]$ and $\boldsymbol{W}_x^\star > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1}w_j$, or

M3: $x \in V[i+1]$ and $\sum_{a \in \partial x}\mathbf{1}\{\partial a \cap (\mathscr{M}[1] \cup \cdots \cup \mathscr{M}[i]) \neq \emptyset\} \ge \ln^{1/4} n$.

To see this, assume that $x \in \mathscr{M}[i+1]$ while **M3** does not hold. Then comparing (4.7) and (4.17) we obtain

$$\left|W_{x,j}(\tau) - \boldsymbol{W}_{x,j}\right| \le \ln^{1/4} n \qquad \text{for all } 1 \le j < s. \tag{4.49}$$

Moreover, the definition (4.14) of the weights, the choice (4.3) of s, and the choices (4.14) of ζ and the weights w_j ensure that $0 \le w_j \le O(s) = O(\ln\ln n)$. This bound implies together with the definition (4.12) of the scores $\boldsymbol{W}_x^\star$ and (4.49) that

$$|\boldsymbol{W}_x^\star - W_x^\star(\tau)| = o(\zeta\Delta). \tag{4.50}$$

Thus, combining (4.50) with the definition of τ_x in Steps 5–10 of SPIV, we conclude that either **M1** or **M2** occurs.

Finally, to bound $\mathscr{M}[i+1]$ let $\mathscr{M}_1[i+1]$, $\mathscr{M}_2[i+1]$, $\mathscr{M}_3[i+1]$ be the sets of individuals $x \in V[i+1]$ for which **M1**, **M2** or **M3** occurs, respectively. Then Lemmas 4.4 and 4.5 imply that w.h.p.

$$|\mathscr{M}_1[i+1]|, |\mathscr{M}_2[i+1]| \le k\exp\left(-\frac{\ln n}{(\ln\ln n)^5}\right).$$

Furthermore, Lemma 4.16 shows that $|\mathscr{M}_3[i+1]| \le \sum_{h=1}^{i} |\mathscr{M}[h]|$ w.h.p. Hence, we obtain the relation

$$|\mathscr{M}[i+1]| \le k\exp\left(-\frac{\ln n}{(\ln\ln n)^5}\right) + \sum_{h=1}^{i} |\mathscr{M}[h]|. \tag{4.51}$$

Because (4.2) ensures that the total number of compartments is $\ell = O(\ln^{1/2} n)$, the bound (4.51) implies that $|\mathscr{M}[i+1]| \le O(\ell^2 k\exp(-(\ln n)/(\ln\ln n)^5)$ for all $i \in [\ell]$ w.h.p. Summing on i completes the proof. □

4.10. **Proof of Proposition 4.7.** For an infected individual $x \in V$ let

$$\boldsymbol{S}_x[j] = \left|\left\{a \in F[j] \cap \partial x : V_1 \cap \partial a = \{x\}\right\}\right| \qquad \text{and} \qquad \boldsymbol{S}_x = \sum_{j=1}^{\ell} \boldsymbol{S}_x[j].$$

Thus, $\boldsymbol{S}_x[j]$ is the number of positive sets $a \in F[j]$ that x has to itself, i.e., tests that do not contain a second infected individual, and $\boldsymbol{S}_x$ is the total number of such tests.

Lemma 4.17. *Assume that $m \ge (1+\varepsilon)m_{\inf}$. W.h.p. we have $\min_{x\in V_1} \boldsymbol{S}_x \ge \sqrt{\Delta}$.*

Proof. Due to Proposition 4.1 we may condition on the event

$$\mathscr{N} = \left\{\forall i \in [\ell] : \frac{m}{2\ell} - \sqrt{m}\ln n \le |F_0[i]| \le \frac{m}{2\ell} + \sqrt{m}\ln n\right\}.$$

We claim that given $\mathscr{N}$ for each $x \in V_1[i]$, $i \in [\ell]$, the random variable $\boldsymbol{S}_x$ has distribution

$$\boldsymbol{S}_x[i+j-1] \sim \mathrm{Hyp}\left(\frac{m}{\ell}, \frac{m}{2\ell} + O(\sqrt{m}\ln n), \frac{\Delta}{s}\right). \tag{4.52}$$

To see this, consider the set $F_x[i+j-1] = \left\{a \in F[i+j-1] : \partial a \cap V_1 \setminus \{x\} = \emptyset\right\}$ of all tests in compartment $F[i+j-1]$ without an infected individual besides possibly x. Since x joins $\Delta/s = O(\ln n)$ tests in $F[i+j-1]$, given $\mathscr{N}$ we have

$$\left|F_{0,x}[i+j]\right| = |F_0[i+j]| + O(\ln n) = \frac{m}{2\ell} + O(\sqrt{m}\ln n). \tag{4.53}$$

Furthermore, consider the experiment of first constructing the test design $\boldsymbol{G}$ and then re-sampling the set ∂x of neighbours of x; i.e., independently of $\boldsymbol{G}$ we have x join Δ/s random tests in each compartment $F[i+j]$. Then the resulting test design $\boldsymbol{G}'$ has the same distribution as $\boldsymbol{G}$ and hence the random variable $\boldsymbol{S}'_x[i+j-1]$ that counts tests $a \in F[i+j-1] \cap \partial x$ that do not contain another infected individual has the same distribution as $\boldsymbol{S}_x[i+j-1]$. Moreover, the conditional distribution of $\boldsymbol{S}'_x[i+j-1]$ given $\boldsymbol{G}$ reads

$$\boldsymbol{S}'_x[i+j-1] \sim \mathrm{Hyp}\left(\frac{m}{\ell}, |F_{0,x}[i+j-1]|, \frac{\Delta}{s}\right). \tag{4.54}$$

Combining (4.53) and (4.54), we obtain (4.52).

To complete the proof we combine (4.52) with Lemma 2.2, which implies that

$$\mathbb{P}\left[\boldsymbol{S}_x[i+j-1] \le \sqrt{\Delta} \mid x \in V_1\right] \le \exp\left(-\frac{\Delta}{s} D_{\mathrm{KL}}\left((1+o(1))s/\sqrt{\Delta}\|1/2 + o(1)\right)\right) = \exp\left(-(1+o(1))\frac{\Delta\ln 2}{s}\right). \tag{4.55}$$

Since **SC1** ensures that the random variables $(\boldsymbol{S}_x[i+j-1])_{j\in[s]}$ are mutually independent, (4.55) yields

$$\mathbb{P}\left[\boldsymbol{S}_x \le \sqrt{\Delta} \mid x \in V_1\right] \le 2^{-(1+o(1))\Delta}. \tag{4.56}$$

Finally, the assumption $m \ge (1+\varepsilon)m_{\inf}$ for a fixed $\varepsilon > 0$ and the choice (4.5) of Δ ensure that $2^{-(1+o(1))\Delta} = o(1/k)$. Thus, the assertion follows from (4.56) by taking a union bound on $x \in V_1$. □

Proof of Proposition 4.7. For $j = 1 \ldots \lceil \ln n \rceil$, let

$$\mathscr{M}_j = \left\{x \in V : \tau_x^{(j)} \ne \boldsymbol{\sigma}_x\right\}$$

contain all individuals that remain misclassified at the j-th iteration of the clean-up step. Proposition 4.6 shows that w.h.p.

$$|\mathscr{M}_1| \le k\exp\left(-\frac{\ln n}{(\ln\ln n)^6}\right). \tag{4.57}$$

Furthermore, in light of Lemma 4.17 we may condition on the event $\mathscr{A} = \{\min_{x\in V_1} \boldsymbol{S}_x \ge \sqrt{\Delta}\}$.

We now claim that given $\mathscr{A}$ for every $j \geq 1$

$$\mathscr{M}_{j+1} \subset \left\{ x \in V : \sum_{a \in \partial x \setminus F[0]} \left| \partial a \cap \mathscr{M}_j \setminus \{x\} \right| \geq \lceil \ln^{1/4} n \rceil \right\}. \tag{4.58}$$

To see this, suppose that $x \in \mathscr{M}_{j+1}$ and recall that the assumption $m \geq m_{\inf}$ and (4.5) ensure that $\Delta = \Omega(\ln n)$. Also recall that SPIV's Step 15 thresholds the number

$$S_x(\tau^{(j)}) = \sum_{a \in \partial x : \hat{\boldsymbol{\sigma}}_a = 1} \mathbf{1}\left\{ \forall y \in \partial a \setminus \{x\} : \tau_y^{(j)} = 0 \right\}$$

of positive tests containing x whose other individuals are deemed uninfected. There are two cases to consider.

- **Case 1:** $x \in V_0$**:** in this case every positive tests $a \in \partial x$ contains an individual that is actually infected. Hence, if $\tau_y^{(j)} = 0$ for all $y \in \partial a \setminus \{x\}$, then $\partial a \cap \mathscr{M}_j \setminus \{x\} \neq \emptyset$. Consequently, since Step 15 of SPIV applies the threshold of $S_x(\tau^{(j)}) \geq \ln^{1/4} n$, there are at least $\ln^{1/4} n$ tests $a \in \partial x$ such that $\partial a \cap \mathscr{M}_j \setminus \{x\} \neq \emptyset$.
- **Case 2:** $x \in V_1$**:** given $\mathscr{A}$ every infected x participates in at least $\boldsymbol{S}_x \geq \sqrt{\Delta} = \Omega(\ln^{1/2} n)$ tests that do not actually contain another infected individual. Hence, if $S_x(\tau^{(j)}) \leq \ln^{1/4} n$, then at least $\sqrt{\Delta} - \ln^{1/4} n \geq \ln^{1/4} n$ tests $a \in \partial x$ contain an individual from $\mathscr{M}_j \setminus \{x\}$.

Thus, we obtain (4.58). Finally, (4.57), (4.58) and Lemma 4.16 show that w.h.p. $|\mathscr{M}_{j+1}| \leq |\mathscr{M}_j|/3$ for all $j \geq 1$. Consequently, $\mathscr{M}_{\lceil \ln n \rceil} = \emptyset$ w.h.p. □

5. Optimal adaptive group testing

In this final section we show how the test design $\boldsymbol{G}$ from Section 4 can be extended into an optimal two-stage adaptive design. The key observation is that Proposition 4.6, which summarises the analysis of the first two phases of SPIV (i.e., steps 1–10) only requires $m \geq (1+\varepsilon)m_{\mathrm{ad}}$ tests. In other words, the excess number $(1+\varepsilon)(m_{\inf} - m_{\mathrm{ad}})$ of tests required for non-adaptive group testing is necessary only to facilitate the clean-up step, namely phase 3 of SPIV.

Replacing phase 3 of SPIV by a second test stage, we obtain an optimal adaptive test design. To this end we follow Scarlett [32], who observed that a single-stage group testing scheme that correctly diagnoses all but $o(k)$ individuals with $(1+o(1))m_{\mathrm{ad}}$ tests could be turned into a two-stage design that diagnoses all individuals correctly w.h.p. with $(1+o(1))m_{\mathrm{ad}}$ tests in total. (Of course, at the time no such optimal single-stage test design and algorithm were known.) The second test stage works as follows. Let τ denote the outcome of phases 1 and 2 of SPIV applied to $\boldsymbol{G}$ with $m = (1+\varepsilon)m_{\mathrm{ad}}$.

- **T1:** Test every individual from the set $V_1(\tau) = \{x \in V : \tau_x = 1\}$ of individuals that SPIV diagnosed as infected separately.
- **T2:** To the individuals $V_0(\tau) = \{x \in V : \tau_x = 0\}$ apply the random d-out design and the DD-algorithm from Section 4.1 with a total of $m = k$ tests and $d = \lceil 10 \ln n \rceil$.

Let $\tau' \in \{0,1\}^V$ be the result of **T1**–**T2**.

Proposition 5.1. *W.h.p. we have $\tau'_x = \boldsymbol{\sigma}_x$ for all $x \in V$.*

As a matter of course **T1** renders correct results, i.e., for all individuals $x \in V_1(\tau)$ we have $\tau'_x = \boldsymbol{\sigma}_x$. Further, to analyse **T2** we use a similar argument as in the analysis of the first phase of SPIV in Section 4.5; we include the analysis for the sake of completeness. We begin by investigating the number of negative tests. Let $\boldsymbol{G}'$ denote the test design set up by **T2**, let $F' = \{b_1, \ldots, b_k\}$ denote its set of tests and let $\hat{\boldsymbol{\sigma}}_{b_1}, \ldots, \hat{\boldsymbol{\sigma}}_{b_k}$ signify the corresponding test results. Further, let $F'_0 = \{b \in F' : \hat{\boldsymbol{\sigma}}_b = 0\}$ and $F'_1 = \{b \in F' : \hat{\boldsymbol{\sigma}}_b = 1\}$ be the set of negative and positive tests, respectively.

Lemma 5.2. *W.h.p. we have $|F'_1| \leq \frac{k}{2}$.*

Proof. Proposition 4.6 implies that w.h.p.

$$|V_0(\tau) \cap V_1| \leq \sum_{x \in V} \mathbf{1}\{\tau_x \neq \boldsymbol{\sigma}_x\} \leq k \exp\left(-\frac{\ln n}{(\ln \ln n)^6} \right). \tag{5.1}$$

Moreover, since every individual $x \in V_0(\tau)$ joins d random tests, for any specific test $b \in F'$ we have

$$\mathbb{P}\left[x \in \partial_{\boldsymbol{G}'} b\right] = 1 - \mathbb{P}\left[x \notin \partial_{\boldsymbol{G}'} b\right] = 1 - \binom{k-1}{d}\binom{k}{d}^{-1} = \frac{d}{k}(1 + O(n^{-\Omega(1)})).$$

Hence, for every test $b \in F'$,

$$\mathbb{E}\left[|\partial b \cap V_1| \,\middle|\, |V_0(\tau) \cap V_1| \le k \exp\left(-\frac{\ln n}{(\ln\ln n)^6}\right)\right] = O(1/\ln n).$$

Consequently,

$$\mathbb{E}\left[|F_1'| \,|\, |V_0(\tau) \cap V_1| \le k/\ln n\right] = O(k/\ln n). \tag{5.2}$$

Finally, combining (5.1) and (5.2) and applying Markov's inequality, we conclude that $|F_1'| \le \frac{k}{2}$ w.h.p. □

Corollary 5.3. *W.h.p. for every $x \in V_0(\tau)$ there is a test $b \in F'$ such that $\partial b \setminus \{x\} \subset V_0$.*

Proof. We construct the random graph $\boldsymbol{G}'$ in two rounds. In the first round we first expose the neighbourhoods $(\partial_{\boldsymbol{G}'} y)_{y \in V_0(\tau)\setminus\{x\}}$. Lemma 5.2 implies that after the first round the number $\boldsymbol{X}$ of tests that do not contain an infected individual $y \in V_0(\tau) \cap V_1$ exceeds $k/2$ w.h.p. In the second round we expose $\partial_{\boldsymbol{G}'} x$. Because $\partial_{\boldsymbol{G}'} x$ is chosen independently of the neighbourhoods $(\partial_{\boldsymbol{G}'} y)_{y \in V_0(\tau)\setminus\{x\}}$, the number of tests $b \in \partial_{\boldsymbol{G}'} x$ that do not contain an infected individual $y \in V_0(\tau) \cap V_1$ has distribution $\mathrm{Hyp}(k, \boldsymbol{X}, d)$. Therefore, since $d \ge 10 \ln n$ we obtain

$$\mathbb{P}\left[\forall b \in \partial x : V_1 \cap \partial b \setminus \{x\} \ne \emptyset \mid \boldsymbol{X} \le k/2\right] \le \mathbb{P}\left[\mathrm{Hyp}(k, k/2, d) = 0\right] \le 2^{-d} = o(1/n). \tag{5.3}$$

Finally, the assertion follows (5.3) and the union bound. □

Proof of Proposition 5.1. Corollary 5.3 shows that we may assume that for every $x \in V_0(\tau)$ there is a test $b_x \in F'$ with $\partial b_x \setminus \{x\} \subset V_0$. As a consequence, upon executing the first step **DD1** of the DD algorithm, **T2** will correctly diagnose all individuals $x \in V_0(\tau) \cap V_0$. Therefore, if $x \in V_0(\tau) \cap V_1$, then **DD2** will correctly identify x as infected because all other individuals $y \in \partial b_x$ were already identified as healthy by **DD1**. Thus, $\tau'_x = \boldsymbol{\sigma}_x$ for all $x \in V$. □

Proof of Theorem 1.3. Proposition 5.1 already establishes that the output of the two-stage adaptive test is correct w.h.p. Hence, to complete the proof we just observe that the total number of tests comes to $(1+\varepsilon)m_{\mathrm{ad}}$ for the first stage plus $|V_1(\tau)| + k$ for the second stage. Furthermore, Proposition 4.6 implies that w.h.p.

$$|V_1(\tau)| \le |V_1| + \sum_{x \in V} \mathbf{1}\{\tau_x \ne \boldsymbol{\sigma}_x\} \le k\left(1 + \exp\left(-\frac{\ln n}{(\ln\ln n)^6}\right)\right) = (1 + o(1))k.$$

Thus, the second stage conducts $O(k) = o(m_{\mathrm{ad}})$ tests. □

Acknowledgment. We thank Arya Mazumdar for bringing the group testing problem to our attention.

References

[1] E. Abbe: Community detection and stochastic block models: recent developments. The Journal of Machine Learning Research **18** (2017) 6446–6531.

[2] A. Alaoui, A. Ramdas, F. Krzakala, L. Zdeborová, M. Jordan: Decoding from pooled data: phase transitions of message passing. IEEE Transactions on Information Theory **65** (2019) 572–585.

[3] A. Alaoui, A. Ramdas, F. Krzakala, L. Zdeborová, M. Jordan: Decoding from pooled data: Sharp information-theoretic bounds. SIAM Journal on Mathematics of Data Science **1** (2019) 161–188.

[4] M. Aldridge, L. Baldassini, O. Johnson: Group testing algorithms: bounds and simulations. IEEE Transactions on Information Theory **60** (2014) 3671–3687.

[5] M. Aldridge: Individual testing is optimal for nonadaptive group testing in the linear regime. IEEE Transactions on Information Theory **65** (2019) 2058–2061.

[6] M. Aldridge, O. Johnson, J. Scarlett: Group testing: an information theory perspective. Foundations and Trends in Communications and Information Theory (2019).

[7] N. Alon, M. Krivelevich, B. Sudakov: Finding a large hidden clique in a random graph. Proc. 9th SODA (1998) 594–598.

[8] T. Berger, V. Levenshtein: Asymptotic efficiency of two-stage disjunctive testing. IEEE Transactions on Information Theory, **48** (2002) 1741–1749.

[9] M. Brennan, G. Bresler: Optimal average-case reductions to sparse pca: From weak assumptions to strong hardness. arXiv:1902.07380.

[10] H. Chen, F. Hwang: A survey on nonadaptive group testing algorithms through the angle of decoding. Journal of Combinatorial Optimization **15** (2008) 49–59.

[11] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, P. Loick: Information-theoretic and algorithmic thresholds for group testing. Proc. 46th ICALP (2019) #43.

[12] A. Decelle, F. Krzakala, C. Moore, L. Zdeborová: Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. Phys. Rev. E **84** (2011) 066106.
[13] D. Donoho: Compressed sensing. IEEE Transactions on Information Theory **52** (2006) 1289–1306.
[14] D. Donoho, A. Javanmard , A. Montanari: Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing. IEEE Transactions on Information Theory **59** (2013) 7434–7464.
[15] R. Dorfman: The detection of defective members of large populations. Annals of Mathematical Statistics **14** (1943) 436–440.
[16] A. D'yachkov, V. Rykov: Bounds on the length of disjunctive codes. Problemy Peredachi Informatsii **18** (1982) 166–171.
[17] P. Erdős, A. Rényi: On Two Problems of Information Theory. Magyar Tud. Akad. Mat. Kutató Int. Közl **8** (1963) 229–243.
[18] A. Felstrom, K. Zigangirov: Time-varying periodic convolutional codes with low-density parity-check matrix. IEEE Transactions on Information Theory **45** (1999) 2181–2191.
[19] W. Hoeffding: Probability inequalities for sums of bounded random variables. In N. Fisher, P. Sen (eds.): The collected works of Wassily Hoeffding. Springer Series in Statistics (Perspectives in Statistics). Springer, New York, NY (1994) 409–426.
[20] F. Hwang: A method for detecting all defective members in a population by group testing. Journal of the American Statistical Association **67** (1972) 605–608.
[21] S. Janson, T. Luczak, A. Rucinski: Random Graphs. John Wiley & Sons (2011).
[22] O. Johnson, M. Aldridge, J. Scarlett: Performance of group testing algorithms with near-constant tests per item. IEEE Transactions on Information Theory **65** (2018) 707–723.
[23] W. Kautz, R. Singleton: Nonrandom binary superimposed codes. IEEE Transactions on Information Theory **10** (1964), 363–377.
[24] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, L. Zdeborová: Statistical-physics-based reconstruction in compressed sensing. Physical Review X **2** (2012) 021005.
[25] S. Kudekar, H. Pfister: The effect of spatial coupling on compressive sensing. Proc. 48th Allerton (2010) 347–353.
[26] S. Kudekar, T. Richardson, R. Urbanke: Threshold saturation via spatial coupling: why convolutional LDPC ensembles perform so well over the BEC. IEEE Transaction on Information Theory **57** (2011) 803–834.
[27] S. Kudekar, T. Richardson, R. Urbanke: Spatially coupled ensembles universally achieve capacity under belief propagation. IEEE Transaction on Information Theory **59** (2013) 7761–7813.
[28] H. Kwang-Ming, D. Ding-Zhu: Pooling designs and nonadaptive group testing: important tools for DNA sequencing. World Scientific (2006).
[29] M. Mézard, M. Tarzia, C. Toninelli: Group testing with random pools: phase transitions and optimal strategy. Journal of Statistical Physics **131** (2008) 783–801.
[30] C. Moore: The computer science and physics of community detection: landscapes, phase transitions, and hardness. Bulletin of the EATCS **121** (2017).
[31] G. Reeves, H. Pfister (2019). Understanding phase transitions via mutual information and MMSE. arXiv:1907.02095.
[32] J. Scarlett: Noisy adaptive group testing: Bounds and algorithms. IEEE Transactions on Information Theory **65** (2018) 3646–3661.
[33] J. Scarlett: An efficient algorithm for capacity-approaching noisy adaptive group testing. Proc. IEEE International Symposium on Information Theory (2019) 2679–2683.
[34] K. Takeuchi, T. Tanaka, T. Kawabata: Improvement of BP-based CDMA multiuser detection by spatial coupling. Proc. IEEE International Symposium on Information Theory Proceedings (2011) 1489–1493.
[35] P. Ungar: The cutoff point for group testing. Communications on Pure and Applied Mathematics **13** (1960) 49–54.
[36] L. Wang, X. Li, Y. Zhang, K. Zhang: Evolution of scaling emergence in large-scale spatial epidemic spreading. PLoS ONE **6** (2011).
[37] Y. Wu, S. Verdu, Renyi information dimension: fundamental limits of almost lossless analog compression. IEEE Transactions on Information Theory **56** (2010) 3721–3748.
[38] L. Zdeborová, F. Krzakala: Statistical physics of inference: thresholds and algorithms. Advances in Physics **65** (2016) 453–552.

AMIN COJA-OGHLAN, acoghlan@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

OLIVER GEBHARD, gebhard@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

MAX HAHN-KLIMROTH, hahnklim@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

PHILIPP LOICK, loick@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

Statistical and Computational Phase Transitions in Group Testing

Amin Coja-Oghlan*[1], Oliver Gebhard†[1], Max Hahn-Klimroth‡[1], Alexander S. Wein§[2], and Ilias Zadik¶[3]

[1]Department of Computer Science, TU Dortmund
[2]Algorithms and Randomness Center, Georgia Tech
[3]Department of Mathematics, MIT

Abstract

We study the *group testing* problem where the goal is to identify a set of k infected individuals carrying a rare disease within a population of size n, based on the outcomes of pooled tests which return positive whenever there is at least one infected individual in the tested group. We consider two different simple random procedures for assigning individuals to tests: the *constant-column design* and *Bernoulli design*. Our first set of results concerns the fundamental *statistical* limits. For the constant-column design, we give a new information-theoretic lower bound which implies that the proportion of correctly identifiable infected individuals undergoes a sharp "all-or-nothing" phase transition when the number of tests crosses a particular threshold. For the Bernoulli design, we determine the precise number of tests required to solve the associated detection problem (where the goal is to distinguish between a group testing instance and pure noise), improving both the upper and lower bounds of Truong, Aldridge, and Scarlett (2020). For both group testing models, we also study the power of *computationally efficient* (polynomial-time) inference procedures. We determine the precise number of tests required for the class of *low-degree polynomial algorithms* to solve the detection problem. This provides evidence for an inherent *computational-statistical* gap in both the detection and recovery problems at small sparsity levels. Notably, our evidence is contrary to that of Iliopoulos and Zadik (2021), who predicted the absence of a computational-statistical gap in the Bernoulli design.[1]

*Email: *amin.coja-oghlan@tu-dortmund.de*. Supported by DFG grant CO 646/3 and DFG grant FOR 2975.

†Email: *oliver.gebhard@tu-dortmund.de*. Supported by DFG grant CO 646/3.

‡Email: *maximilian.hahnklimroth@tu-dortmund.de*. Supported by DFG grant FOR 2975.

§Email: *awein@cims.nyu.edu*. Supported by NSF grants CCF-2007443 and CCF-2106444. Part of this work was done while the author was visiting the Simons Institute for the Theory of Computing. Part of this work was done while the author was with the Courant Institute at NYU, partially supported by NSF grant DMS-1712730 and by the Simons Collaboration on Algorithms and Geometry.

¶Email: *izadik@mit.edu*. Supported by the Simons-NSF grant DMS-2031883 on the Theoretical Foundations of Deep Learning and the Vannevar Bush Faculty Fellowship ONR-N00014-20-1-2826. Part of this work was done while the author was visiting the Simons Institute for the Theory of Computing. Part of this work was done while the author was with the Center for Data Science at NYU, supported by a Moore-Sloan CDS postdoctoral fellowship.

[1]Accepted for presentation at the Conference on Learning Theory (COLT) 2022.

Contents

1 Introduction

Motivated by the ongoing COVID-19 pandemic [MNB+21, MTB12] but also a growing algorithmic and information-theoretic literature [AJS19], in this work we focus on the *group (or pooled) testing model*. Introduced by [Dor43], group testing is concerned with finding a subset of k individuals carrying a rare disease within a population of size n. One is equipped with a procedure that allows for testing groups of individuals such that a test returns positive if (and only if) at least one infected individual is contained in the tested group. The ultimate goal is to find a pooling procedure and a (time-efficient) algorithm such that inference of the infection status of all individuals is conducted with as few tests as possible. Furthermore, group testing has found its way into various real-world applications such as DNA sequencing [KMDZ06, ND00], protein interaction experiments [MDM13, TM06] and machine learning [EVM15].

As carrying out a test is often time-consuming, many real-world applications call for fast identification schemes. As a consequence, recent research focuses on *non-adaptive* pooling schemes, i.e., all tests are conducted in parallel [SC16, Ald19, COGHKL20a, COGHKL20b, IZ21]. On top of this, naturally the testing scheme is required to be simple as well. Two of the most well-established and simple non-adaptive group testing designs are the *Bernoulli design* and the *constant-column design* (for a survey, see [AJS19]). The Bernoulli design is a randomised pooling scheme under which each individual participates in each test with a fixed probability q independently of everything else [SC16]. In the constant-column design [AJS16, COGHKL20a], each individual independently chooses a fixed number Δ of tests uniformly at random. We remark that the *spatially coupled design* of [COGHKL20b] may be an attractive choice in practice because it admits information-theoretically optimal inference with a computationally efficient algorithm. In this paper our focus will be on the two simpler designs (Bernoulli and constant-column), which may be favorable due to their simplicity and also serve as a testbed for studying computational-statistical gaps.

In this work, we take the number of infected individuals to scale *sublinearly* in the population size as is typical in group testing tasks, that is $k = n^{\theta+o(1)}$ for a fixed constant $\theta \in (0,1)$. This regime is mathematically interesting and is also the one most suitable for modelling the early stages of an epidemic in the context of medical testing [WLZ+11]. In the two group testing models, we study two different inference tasks (defined formally in Section 2.1): (a) *approximate recovery*, where the goal is to achieve almost perfect correlation with the set of infected individuals, and (b) *weak recovery*, where the goal is to achieve positive correlation with the set of infected individuals. The task of *exact recovery* has also been studied (see [COGHKL20a]) but will not be our focus here.

Recently, there has been substantial work on the information-theoretic limits of group testing [CCJS11, ABJ14, COGHKL20a, COGHKL20b, TAS20]. An interesting recent discovery is that for the Bernoulli group testing model there exists a critical threshold $m_{\inf} := (\ln 2)^{-1} k \ln(n/k)$ such that when the number of tests m satisfies $m \geq (1+\varepsilon)m_{\inf}$ for any fixed $\varepsilon > 0$ there is a (brute-force) algorithm that can approximately recover the infected individuals, but when $m \leq (1-\varepsilon)m_{\inf}$ no algorithm (efficient or not) can even weakly recover the infected individuals. This sharp phase transition, known as the *All-or-Nothing (AoN) phenomenon*, was first proven by [TAS20] for $\theta = 0$ (that is, $k = n^{o(1)}$) and then proven for all $\theta \in [0,1)$ by [NWZ21]. This sharp phenomenon has been established recently in many other sparse Generalized Linear Models (GLMs), starting with sparse regression [RXZ19b]. *Our first main result* (Theorem 3.1) establishes the AoN phenomenon

in the constant-column group testing model for any $\theta \in (0,1)$, occurring at the same information-theoretic threshold $m_{\inf}$ as in the Bernoulli model. To our knowledge, this is the first instance where AoN has been established for a GLM where the samples (tests) are not independent (see Section 1.1 for further discussion).

An emerging but less understood direction is to study the algorithmic thresholds of the group testing models. In both group testing models, the best known polynomial-time algorithm achieves approximate recovery only under the statistically suboptimal condition $m \geq (1+\varepsilon)m_{\text{alg}}$ where $m_{\text{alg}} := (\ln 2)^{-1} m_{\inf}$. For the constant-column design, the algorithm achieving this is Combinatorial Orthogonal Matching Pursuit (COMP) [CCJS11, CJSA14], which simply outputs all individuals who participate in no negative tests. For the Bernoulli design, the algorithm achieving m_{alg} is called Separate Decoding [SC18], which outputs all individuals who participate in no negative tests and "sufficiently many" positive tests (above some threshold). These results raise the question of whether better algorithms exist, or whether there is an inherent *computational-statistical gap*. Starting from the seminal work of [BR13], conjectured gaps between the power of all estimators and the power of all *polynomial-time* algorithms have appeared recently throughout many high-dimensional statistical inference problems. While we do not currently have tools to prove complexity-theoretic hardness of statistical problems, there are various forms of "rigorous evidence" for hardness that can be used to justify these computational-statistical gaps, including average-case reductions (see e.g. [BB20]), sum-of-squares lower bounds (see e.g. [RSS18]), and others.

In the Bernoulli group testing model, the recent work of [IZ21] suggested (but did not prove) that a polynomial-time Markov Chain Monte Carlo (MCMC) method can achieve approximate recovery all the way down to the information-theoretic threshold (that is, using only $m_{\inf}$ tests). The evidence for this is based on first-moment Overlap Gap Property calculations and numerical simulations. The Overlap Gap Property is a landscape property originating in spin glass theory, which has been repeatedly used to offer evidence for the performance of local search and MCMC methods in inference problems, as initiated by [GZ17]. A significant motivation for the present work is to gain further insight into the existence or not of such a computational-statistical gap for both the constant-column and Bernoulli designs. Our approach is based on the well-studied *low-degree likelihood ratio* (discussed further in Section 2.2), which is another framework for understanding computational-statistical gaps.

In line with most existing results using the low-degree framework, we consider a *detection* (or *hypothesis testing*) formulation of the problem. In our case, this amounts to the task of deciding whether a given group testing instance was actually drawn from the group testing model with k infected individuals, or whether it was drawn from an appropriate "null" model where the test outcomes are random coin flips (containing no information about the infected individuals). *Our second set of results* is that for both the constant-column and Bernoulli designs, we pinpoint the precise low-degree detection threshold $m_{\text{LD}} = m_{\text{LD}}(k,n)$ (which is different for the two designs) in the following sense: when the number of tests exceeds this threshold, there is a polynomial-time algorithm that provably achieves *strong detection* (that is, testing with $o(1)$ error probability); on the other hand, if the number of tests lies below the threshold, all *low-degree algorithms* provably fail to *separate* the two distributions (as defined in Section 2.2). This class of low-degree algorithms captures the best known poly-time algorithms for many high-dimensional testing tasks (including those studied in this paper), and so our result suggests inherent computational hardness of detection below the threshold m_{LD}. For the exact thresholds, see Theorem 3.2 for the constant-

column design and Theorem 3.3 for Bernoulli design.

Since approximate recovery is a harder problem than detection (this is formalized in Appendix C), our results also suggest that approximate recovery is computationally hard below m_{LD}. Since m_{LD} exceeds m_{inf} for sufficiently small θ (see Figure 2), this suggests the presence of a computational-statistical gap for the recovery problem (in both group testing models). Notably, our evidence is contrary to that of [IZ21], who suggested the absence of a comp-stat gap in the Bernoulli model for all $\theta \in (0, 1)$.

Finally, *our third set of results* is to identify the precise *statistical* (information-theoretic) threshold for detection in the Bernoulli design (commonly referred to in the statistics literature as the *detection boundary*); see Theorem 3.4.

Our main results are summarized by the phase diagrams in Figure 2.

1.1 Relation to Prior Work

Detection in the Bernoulli design To our knowledge, the only existing work on the detection boundary in group testing is [TAS20], which focused on the Bernoulli design. They gave a detection algorithm and an information-theoretic lower bound which did not match. In this work we pinpoint the precise information-theoretic detection boundary by improving both the algorithm and lower bound (Theorem 3.4). The new algorithm involves counting the number of individuals who participate in no negative tests and "sufficiently many" positive tests (above some carefully chosen threshold). The lower bound of [TAS20] is based on a second moment calculation, and our improved lower bound uses a *conditional* second moment calculation (which conditions away a rare "bad" event).

Strictly speaking, our detection problem differs from the one studied by [TAS20] because our detection problem takes place on "pre-processed" graphs where the negative tests have been removed (see Section 2.1), but we show in Appendix D that our results can be transferred to their setting.

All-or-Nothing phenomenon The All-or-Nothing (AoN) phenomenon was originally proven in the context of sparse regression with an i.i.d. Gaussian measurement matrix [GZ17, RXZ19a, RXZ19b], and was later established for (a) various other Generalized Linear Models (GLMs) such as Bernoulli group testing [TAS20, NWZ21] and the Gaussian Perceptron [LBM20, NWZ21], (b) variants of sparse principal component analysis [BMR20, NWZ20], and (c) graph matching models [WXY21]. In all of the GLM cases, a key assumption behind all such proofs is that the samples (or tests in the case of Bernoulli group testing) are independent. This sample independence gives rise to properties similar to the I-MMSE formula [GSV05], which can then be used to establish the AoN phenomenon by simply bounding the KL divergence between the planted model and an appropriate null model.

In the present work, we establish AoN for the constant-column group testing model which is a GLM where the samples (tests) are *dependent*. Despite this barrier, we manage to prove this result by following a more involved but direct argument, which employs a careful conditional second moment argument alongside a technique from the study of random CSPs known as the "planting trick" originally used in the context of random k-SAT [ACO08]. A more detailed proof outline is given in Section 5.

Low-degree lower bounds Starting from the work of [BHK+19, Hop18, HKP+17, HS17], lower bounds against the class of "low-degree polynomial algorithms" (defined in Section 2.2) are a common form of concrete evidence for computational hardness of statistical problems (see [KWB19] for a survey). In this paper we apply this framework to the detection problems in both group testing models, with a few key differences from prior work. For the Bernoulli design, the standard tool—the *low-degree likelihood ratio*—does not suffice to establish sharp low-degree lower bounds, and we instead need a *conditional* variant of this argument that conditions away a rare "bad" event. While such arguments are common for information-theoretic lower bounds, this is (to our knowledge) the first setting where a conditional low-degree argument has been needed, along with the concurrent work [BEH+22] on sparse regression. Our result for the constant-column design is (to our knowledge) the first example of a low-degree lower bound where the null distribution does not have independent coordinates. For both group testing models, the key insight to make these calculations tractable is a "low-overlap second moment calculation," which is explained in Section 7 (particularly 7.4).

Comparison with [IZ21] Perhaps the most relevant work, in terms of studying the computational complexity of group testing, is the recent work of [IZ21] which focuses on the Bernoulli design. The authors provide simulations and first-moment Overlap Gap Property (OGP) evidence that a polynomial-time "local" MCMC method can approximately recover the infected individuals for any statistically possible number of tests $m \geq (1+\varepsilon)m_{\inf}$ and any $\theta \in (0,1)$. However, proving this remains open.

In contrast, our present work shows that at least when $\theta > 0$ is small enough no low-degree polynomial algorithm can even solve the easier detection task for some number of tests strictly above $m_{\inf}$. Given the low-degree framework's track record of capturing the best known algorithmic thresholds for a wide variety of statistical problems, this casts some doubts on the prediction of [IZ21]. However, our results do not formally imply failure of the MCMC method (which is not a low-degree algorithm) and the failure of low-degree algorithms is only known to imply the failure of MCMC methods for the class of Gaussian additive models [BEH+22]. Our results "raise the stakes" for proving statistical optimality of the MCMC method, as this would be a significant counterexample to optimality of low-degree algorithms for statistical problems.

Notation

We will consider the limit $n \to \infty$. Some parameters (e.g. θ, c) will be designated as "constants" (fixed, not depending on n) while others (e.g. k) will be assumed to scale with n in a prescribed way. Asymptotic notation $o(\cdot), O(\cdot), \omega(\cdot), \Omega(\cdot)$ pertains to this limit (unless stated otherwise), i.e., this notation may hide factors depending on constants such as θ, c. We use $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$ to hide a factor of $(\ln n)^{O(1)}$. An event is said to occur *with high probability* if it has probability $1 - o(1)$, and *overwhelming probability* if it has probability $1 - n^{-\omega(1)}$.

2 Getting Started

2.1 Group Testing Setup and Objectives

We will consider two different group testing models. The following basic setup pertains to both.

Group testing We first fix two constants $\theta \in (0,1)$ and $c > 0$. A group testing instance is generated as follows. There are n individuals $x_1, \dots, x_n$ out of which exactly $k = n^{\theta+o(1)}$ are infected. There are $m = (c+o(1))k\ln(n/k)$ tests $a_1, \dots, a_m$.

For each test, a particular subset of the individuals is chosen to participate in that test, according to one of the two designs (constant-column or Bernoulli) described below. The assignment of individuals to tests can be expressed by a bipartite graph (see Figure 1). The *ground-truth* $\boldsymbol{\sigma} \in \{0,1\}^n$ is drawn uniformly at random among all binary vectors of length n and Hamming weight k. We say individual x_i is infected if and only if $\boldsymbol{\sigma}_i = 1$. We denote the sequence of test results by $\hat{\boldsymbol{\sigma}} \in \{0,1\}^m$, where $\hat{\boldsymbol{\sigma}}_j$ is equal to one if and only if the j-th test contains at least one infected individual.

We consider two different schemes for assigning individuals to tests, which are defined below.

Constant-column design In the *constant column weight design* (also called the *random regular design*), every individual independently chooses a set of exactly $\Delta = (c+o(1))\ln(2)\ln(n/k)$ tests to participate in, uniformly at random from the $\binom{m}{\Delta}$ possibilities.

Bernoulli design In the *Bernoulli design*, every individual participates in each test independently with probability $q := \nu/k$ where $\nu = \ln 2 + o(1)$ is the solution to $(1-\nu/k)^k = 1/2$ so that each test is positive with probability exactly $1/2$.

We remark that the parameter ν (in the Bernoulli design) and the constant $\ln(2)$ in the definition of Δ (in the constant-column design) could have been treated as free tuning parameters. To simplify matters, we have chosen to fix these values so that roughly half the tests are positive (maximizing the "information content" per test), but we expect our results could be readily extended to the general case.

We will be interested in the task of recovering the ground truth $\boldsymbol{\sigma}$. Two different notions of success are considered, as defined below.

Approximate recovery An algorithm is said to achieve *approximate recovery* if, given input $(\boldsymbol{G}_{GT}, \hat{\boldsymbol{\sigma}}, k)$, it outputs a binary vector $\boldsymbol{\tau} \in \{0,1\}^n$ with the following guarantee: $\frac{\langle \boldsymbol{\tau}, \boldsymbol{\sigma}\rangle}{\|\boldsymbol{\tau}\|_2\|\boldsymbol{\sigma}\|_2} = 1 - o(1)$ with probability $1 - o(1)$.

Equivalently, approximate recovery means the number of false positive and false negatives are both $o(k)$.

Weak recovery An algorithm is said to achieve *weak recovery* if, given input $(\boldsymbol{G}_{GT}, \hat{\boldsymbol{\sigma}}, k)$, it outputs a binary vector $\boldsymbol{\tau} \in \{0,1\}^n$ with the following guarantee: with probability $1 - o(1)$, $\frac{\langle \boldsymbol{\tau}, \boldsymbol{\sigma}\rangle}{\|\boldsymbol{\tau}\|_2\|\boldsymbol{\sigma}\|_2} = \Omega(1)$.

Pre-processing via COMP Note that in both models we can immediately classify any individual who participates in a negative test as uninfected. Therefore, the first step in any recovery algorithm should be to pre-process the graph by removing all negative tests and their adjacent individuals. (We sometimes refer to this pre-processing step as COMP because it is the main step of the COMP algorithm of [CCJS11, CJSA14], which simply performs this pre-processing step and then reports all remaining individuals as infected.) The resulting graph is denoted $\boldsymbol{G}'_{GT}$ (see Figure 1). We let N denote the number of remaining individuals and let M denote the number of remaining tests. We use $\boldsymbol{\sigma}' \in \{0,1\}^N$ to denote the indicator vector for the infected individuals. Note that after pre-processing, all remaining tests are positive and so $\hat{\boldsymbol{\sigma}}$ can be discarded.

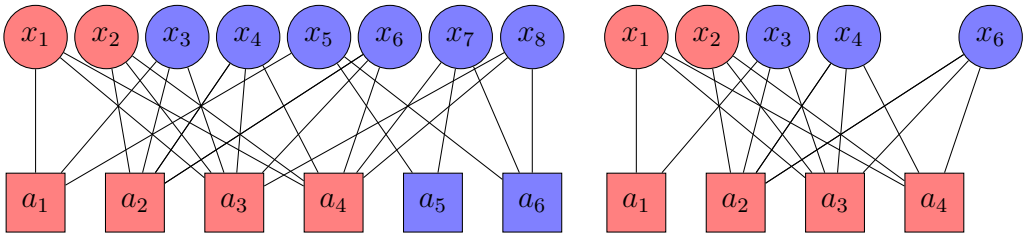

Figure 1: The bipartite factor graph representing a group testing instance. Circles represent individuals while squares represent tests. The colour of circle/square indicates *infected / positive* in red and *uninfected / negative* in blue. The left figure shows an instance of $\boldsymbol{G}_{GT}$ while the right figure shows the corresponding instance of $\boldsymbol{G}'_{GT}$ where individuals in negative tests have already been classified and removed.

In addition to recovery, we will also consider an easier hypothesis testing task. Here the goal is to distinguish between a ("planted") group testing instance and an unstructured ("null") instance. We now define this testing model for both group testing designs. The input is an (N, M)-bipartite graph, representing a group testing instance that has already been pre-processed as described above.

Constant-column design (testing) Let $N = N_n$ and $M = M_n$ scale as $N = n^{1-(1-\theta)c(\ln 2)^2+o(1)}$ and $M = (c/2+o(1))k\ln(n/k)$; this choice is justified below. Consider the following distributions over (N, M)-bipartite graphs (encoding adjacency between N individuals and M tests).

- Under the null distribution $\mathbb{Q}$, each of the N individuals participates in exactly Δ (defined above) tests, chosen uniformly at random.
- Under the planted distribution $\mathbb{P}$, a set of k infected individuals out of N is chosen uniformly at random. Then a graph is drawn from $\mathbb{Q}$ conditioned on having at least one infected individual in every test.

Bernoulli design (testing) Let $N = N_n$ and $M = M_n$ scale as $N = n^{1-(1-\theta)\frac{c}{2}\ln 2+o(1)}$ and $M = (c/2 + o(1))k\ln(n/k)$; this choice is justified below. Consider the following distributions over (N, M)-bipartite graphs (encoding adjacency between N individuals and M tests).

- Under the null distribution $\mathbb{Q}$, each of the N individuals participates in each of the M tests with probability q (defined above) independently.

- Under the planted distribution $\mathbb{P}$, a set of k infected individuals out of N is chosen uniformly at random. Then a graph is drawn from $\mathbb{Q}$ conditioned on having at least one infected individual in every test.

Note that in the pre-processed group testing graph $\boldsymbol{G}'_{GT}$, the dimensions N, M are random variables. For the testing problems above, we will instead think of N, M as deterministic functions of n, which are allowed to vary arbitrarily within some range (due to the $o(1)$ terms). The specific scaling of N, M is chosen so that the actual dimensions of $\boldsymbol{G}'_{GT}$ obey this scaling with high probability (see e.g. [COGHKL20a, IZ21]). Furthermore, the planted distribution $\mathbb{P}$ is precisely the distribution of $\boldsymbol{G}'_{GT}$ conditioned on the dimensions N, M.

We now define two different criteria for success in the testing problem.

Strong detection An algorithm is said to achieve *strong detection* if, given input $(\boldsymbol{G}, k)$ with $\boldsymbol{G}$ drawn from either $\mathbb{Q}$ or $\mathbb{P}$ (each chosen with probability $1/2$), it correctly identifies the distribution ($\mathbb{Q}$ or $\mathbb{P}$) with probability $1 - o(1)$.

Weak detection An algorithm is said to achieve *weak detection* if, given input $(\boldsymbol{G}, k)$ with $\boldsymbol{G}$ drawn from either $\mathbb{Q}$ or $\mathbb{P}$ (each chosen with probability $1/2$), it correctly identifies the distribution ($\mathbb{Q}$ or $\mathbb{P}$) with probability $1/2 + \Omega(1)$.

We will establish a formal connection between the testing and recovery problems: any algorithm for approximate recovery can be used to solve strong detection (see Appendix C for exact statements).

2.2 Hypothesis Testing and the Low-Degree Framework

Following [HS17, HKP+17, Hop18], we will study the class of *low-degree polynomial algorithms* as a proxy for computationally-efficient algorithms (see also [KWB19] for a survey). Considering the hypothesis testing setting, suppose we have two (sequences of) distributions $\mathbb{P} = \mathbb{P}_n$ and $\mathbb{Q} = \mathbb{Q}_n$ over $\mathbb{R}^p$ for some $p = p_n$. Since our testing problems are over (N, M)-bipartite graphs, we will set $p = NM$ and take $\mathbb{P}, \mathbb{Q}$ to be supported on $\{0,1\}^p$ (encoding the adjacency matrix of a graph). A *degree-D polynomial algorithm* is simply a multivariate polynomial $f : \mathbb{R}^p \to \mathbb{R}$ of degree (at most) D with real coefficients (or rather, a sequence of such polynomials $f = f_n$). In our case, since the inputs will be binary, the polynomial can be multilinear without loss of generality. In line with prior work, we define two different notions of "success" for polynomial-based tests as follows.

Strong/weak separation A polynomial $f : \mathbb{R}^p \to \mathbb{R}$ is said to *strongly separate* $\mathbb{P}$ and $\mathbb{Q}$ if

$$\sqrt{\max\left\{\operatorname*{Var}_{\mathbb{P}}[f], \operatorname*{Var}_{\mathbb{Q}}[f]\right\}} = o\left(\left|\operatorname*{\mathbb{E}}_{\mathbb{P}}[f] - \operatorname*{\mathbb{E}}_{\mathbb{Q}}[f]\right|\right). \tag{2.1}$$

Also, a polynomial $f : \mathbb{R}^p \to \mathbb{R}$ is said to *weakly separate* $\mathbb{P}$ and $\mathbb{Q}$ if

$$\sqrt{\max\left\{\operatorname*{Var}_{\mathbb{P}}[f], \operatorname*{Var}_{\mathbb{Q}}[f]\right\}} = O\left(\left|\operatorname*{\mathbb{E}}_{\mathbb{P}}[f] - \operatorname*{\mathbb{E}}_{\mathbb{Q}}[f]\right|\right). \tag{2.2}$$

These are natural sufficient conditions for strong/weak detection: note that by Chebyshev's inequality, strong separation immediately implies that strong detection can be achieved by thresholding the output of f; also, by a less direct argument, weak separation implies that weak detection can be achieved using the output of f [BEH+22, Proposition 6.1].

Perhaps surprisingly, it has now been established that for a wide variety of "high-dimensional testing problems" (including planted clique, sparse PCA, community detection, tensor PCA, and many others), the class of degree-$O(\ln p)$ polynomial algorithms is precisely as powerful as the best known polynomial-time algorithms (e.g. [BKW20, DKWB19, Hop18, HKP+17, HS17, KWB19]). One explanation for this is that such polynomials can capture powerful algorithmic frameworks such as spectral methods (see [KWB19], Theorem 4.4). Also, lower bounds against low-degree algorithms imply failure of all *statistical query algorithms* (under mild assumptions) [BBH+21] and have conjectural connections to the *sum-of-squares hierarchy* (see e.g. [HKP+17, Hop18]). While there is no guarantee that a degree-$O(\ln p)$ polynomial can be computed in polynomial time, the success of such a polynomial still tends to coincide with existence of a poly-time algorithm.

In light of the above, *low-degree lower bounds* (i.e., provable failure of all low-degree algorithms to achieve strong/weak separation) is commonly used as a form of concrete evidence for computational hardness of statistical problems. In line with prior work, we will aim to prove hardness results of the following form.

Low-degree hardness If no degree-D polynomial achieves strong (respectively, weak) separation for some $D = \omega(\ln p)$, we say "strong (resp., weak) detection is low-degree hard"; this suggests that strong (resp., weak) detection admits no polynomial-time algorithm and furthermore requires runtime $\exp(\tilde{\Omega}(D))$ where $\tilde{\Omega}$ hides factors of $\ln p$.

In this paper, we will establish low-degree hardness of group testing models in certain parameter regimes. While the implications for all polynomial-time algorithms are conjectural, these results identify apparent computational barriers in group testing that are analogous to those in many other problems. As a result, we feel there is unlikely to be a polynomial-time algorithm in the low-degree hard regime, at least barring a major algorithmic breakthrough.[2] Throughout the rest of this paper we focus on proving low-degree hardness as a goal of inherent interest, and refer the reader to the references mentioned above for further discussion on how low-degree hardness should be interpreted.

3 Main Results

We now formally state our main results on statistical and computational thresholds in group testing, which are summarized in Figure 2. Throughout, recall that we fix the scaling regime $k = n^{\theta+o(1)}$ and $m = (c+o(1))k\ln(n/k)$ for constants $\theta \in (0,1)$ and $c > 0$. Our objective is to characterize the values of (θ, c) for which various group testing tasks are "easy" (i.e., poly-time solvable), "hard" (in the low-degree framework), and (information-theoretically) "impossible."

[2] Strictly speaking, we should perhaps only conjecture computational hardness for a slightly noisy version of group testing (say where a small constant fraction of test results are changed at random) because some "noiseless" statistical problems admit a poly-time algorithm in regimes where low-degree polynomials fail; see e.g. Section 1.3 of [ZSWB21] for discussion.

3.1 Constant-Column Design

Our first set of results pertains to the constant-column design, as defined in Section 2.1.

Weak recovery: All-or-Nothing phenomenon We start by focusing on the information-theoretic limits of weak recovery in the constant-column design. We show that the AoN phenomenon occurs at the critical constant $c_{\inf} = 1/\ln 2$, i.e., at the critical number of tests $m_{\inf} = (\ln 2)^{-1} k \ln(n/k)$. It was known previously that when $c > 1/\ln 2$, one can approximately recover (as defined in Section 2.1) the infected individuals via a brute-force algorithm [COGHKL20a, COGHKL20b]. It was also known that when $c < 1/\ln 2$, one *cannot* approximately recover the infected individuals (see [AJS19]). We show that in fact a much stronger lower bound holds: when $c < 1/\ln 2$, no algorithm can even achieve *weak* recovery.

Theorem 3.1. *Consider the constant-column design with any fixed $\theta \in (0,1)$. If $c < c_{\inf} := 1/\ln 2$ then every algorithm (efficient or not) taking input $(\boldsymbol{G}_{GT}, \hat{\boldsymbol{\sigma}}, k)$ and returning a binary vector $\boldsymbol{\tau} \in \{0,1\}^n$ must satisfy $\frac{\langle \boldsymbol{\tau}, \boldsymbol{\sigma} \rangle}{\|\boldsymbol{\tau}\|_2 \|\boldsymbol{\sigma}\|_2} = o(1)$ with probability $1 - o(1)$. In particular, weak recovery is impossible.*

Combined with the prior work mentioned above, this establishes the All-or-Nothing phenomenon, namely:

- If $c > c_{\inf}$ and $m = (c + o(1)) k \ln(n/k)$ then *approximate* recovery is *possible*.
- If $c < c_{\inf}$ and $m = (c + o(1)) k \ln(n/k)$ then *weak* recovery is *impossible*.

As mentioned in the Introduction, the only algorithms known to achieve approximate recovery with the statistically optimal number of tests $m_{\inf}$ do not have polynomial runtime [COGHKL20a, COGHKL20b]. As a tool for studying this potential computational-statistical gap (and out of independent interest), we next turn our attention to the easier *detection* task. We will return to discuss the implications for hardness of the recovery problem later.

Detection boundary and low-degree methods We first pinpoint the precise "low-degree" threshold $c_{\mathrm{LD}}^{\mathrm{CC}} = c_{\mathrm{LD}}^{\mathrm{CC}}(\theta)$ (where the superscript indicates "constant-column") for detection: above this threshold we prove that a new poly-time algorithm achieves strong detection; below this threshold we prove that all low-degree polynomial algorithms fail to achieve weak separation, giving concrete evidence for hardness (see Section 2.2). As a sanity check for the low-degree lower bound, we also verify that low-degree algorithms indeed succeed at strong separation above the threshold (specifically, this is achieved by a degree-2 polynomial that computes the empirical variance of the test degrees).

Theorem 3.2. *Consider the constant-column design (testing variant) with parameters $\theta \in (0,1)$ and $c > 0$. Define*

$$c_{\mathrm{LD}}^{\mathrm{CC}} = \begin{cases} \frac{1}{(\ln 2)^2} \left(1 - \frac{\theta}{2(1-\theta)}\right) & \text{if } 0 < \theta < 2/3, \\ 0 & \text{if } 2/3 \le \theta < 1. \end{cases} \tag{3.1}$$

(a) (Easy) If $c > c_{\mathrm{LD}}^{\mathrm{CC}}$, there is a degree-2 polynomial achieving strong separation, and a polynomial-time algorithm achieving strong detection.

(b) (Hard) If $c < c_{\mathrm{LD}}^{\mathrm{CC}}$ then there is a $D = n^{\Omega(1)}$ such that any degree-D polynomial fails to achieve weak separation. (This suggests that weak detection requires runtime $\exp(n^{\Omega(1)})$.)

We remark that when $\theta \geq 2/3$, the problem is "easy" for any constant $c > 0$ (and perhaps even for some sub-constant scalings for c, although we have not attempted to investigate this).

Hardness of Recovery Above, we have given evidence for hardness of detection below the threshold $c_{\mathrm{LD}}^{\mathrm{CC}}$. We also show in Appendix C that recovery is a formally harder problem than detection: any poly-time algorithm for approximate recovery can be made into a poly-time algorithm for strong detection, succeeding for the same parameters θ, c. These two results together give evidence for hardness of *recovery* below $c_{\mathrm{LD}}^{\mathrm{CC}}$ via a two-step argument: our low-degree hardness for detection leads us to conjecture that there is no poly-time algorithm for detection below $c_{\mathrm{LD}}^{\mathrm{CC}}$, and this conjecture (if true) formally implies that there is no poly-time algorithm for approximate recovery below $c_{\mathrm{LD}}^{\mathrm{CC}}$. (However, our results do not formally imply failure of *low-degree* algorithms for *recovery*.) Notably, it turns out that $c_{\mathrm{LD}}^{\mathrm{CC}}$ exceeds c_{inf} for some values of θ (namely $0 < \theta < 1 + \frac{1}{2\ln 2 - 3} \approx 0.38$), revealing a possible-but-hard regime for recovery (Region I in Figure 2).

Since the recovery problem might be strictly harder than testing, our results do not pinpoint a precise computational threshold for recovery (even conjecturally). However, one case where we do pinpoint the computational recovery threshold is in the limit $\theta \to 0$: here, the thresholds $c_{\mathrm{LD}}^{\mathrm{CC}}$ and c_{alg} coincide, that is, our low-degree hardness result for detection matches the best known poly-time algorithm for recovery (COMP). This suggests that for small θ, the COMP algorithm is optimal among poly-time methods (for approximate recovery).

An interesting open question is to resolve the low-degree threshold for *recovery*, in the style of [SW20]. However, it is not clear that their techniques immediately apply here.

3.2 Bernoulli Design

Our second set of our results pertains to the Bernoulli design as defined in Section 2.1. As always, we fix the scaling regime $k = n^{\theta + o(1)}$ and $m = (c + o(1))k\ln(n/k)$ for constants $\theta \in (0, 1)$ and $c > 0$.

Detection boundary and low-degree methods We will determine both the statistical and low-degree thresholds for detection. The thresholds are more complicated than in the constant-column design and involve the *Lambert W function*: for $x \geq -\frac{1}{e}$, define $W_0(x)$ to be the unique $y \geq -1$ satisfying $ye^y = x$. We begin with the low-degree threshold.

Theorem 3.3. *Consider the Bernoulli design (testing variant) with parameters $\theta \in (0, 1)$ and $c > 0$. Define*

$$c_{\mathrm{LD}}^{\mathrm{B}} = \begin{cases} -\frac{1}{\ln^2 2} W_0(-\exp(-\frac{\theta}{1-\theta}\ln 2 - 1)) & \text{if } 0 < \theta < \frac{1}{2}(1 - \frac{1}{4\ln 2 - 1}), \\ \frac{1}{\ln 2} \cdot \frac{1-2\theta}{1-\theta} & \text{if } \frac{1}{2}(1 - \frac{1}{4\ln 2 - 1}) \leq \theta < \frac{1}{2}, \\ 0 & \text{if } \frac{1}{2} \leq \theta < 1. \end{cases} \tag{3.2}$$

(a) (Easy) If $c > c_{\mathrm{LD}}^{\mathrm{B}}$, there is a degree-$O(\ln n)$ polynomial achieving strong separation, and a polynomial-time algorithm achieving strong detection.

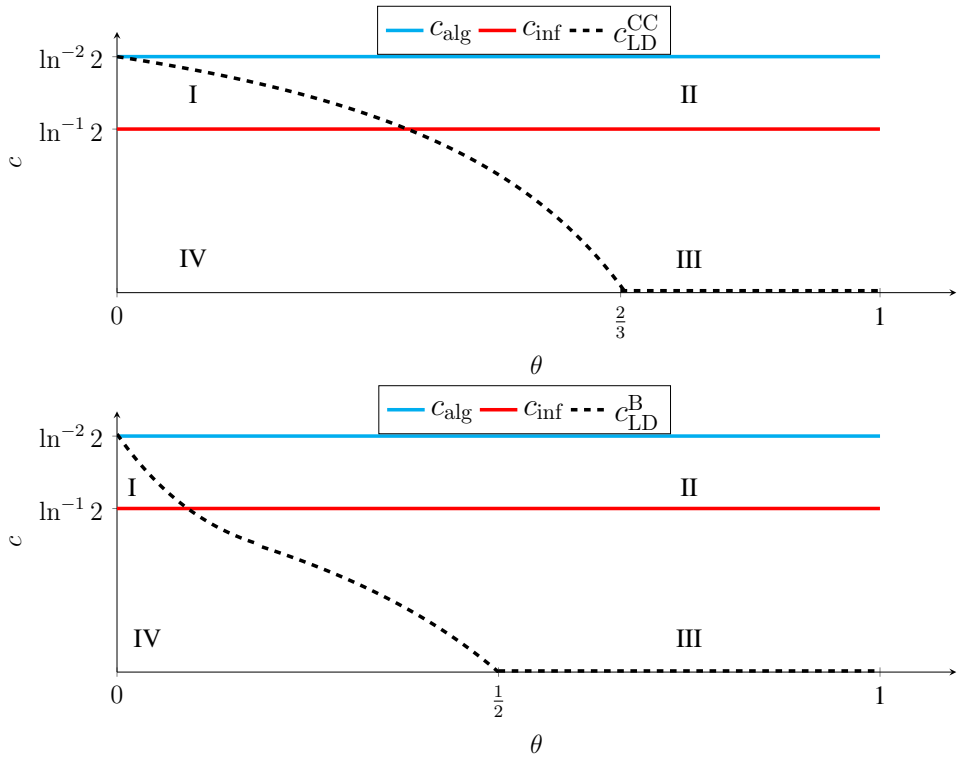

Figure 2: Phase transitions in the constant-column (left) and Bernoulli (right) designs, in (θ, c) space where $k = n^{\theta+o(1)}$ and $m = (c + o(1))k \ln(n/k)$. Recovery is possible above the red line and impossible below it. Polynomial-time recovery is only known above the blue line. Detection is achievable in polynomial time above the dotted line and (low-degree) hard below it. In Region I, detection and recovery are both possible-but-hard. In Region II, detection is easy and recovery is possible, but it is open whether recovery is easy or hard. In Region III, detection is easy and recovery is impossible. In Region IV, recovery is impossible; we expect detection is also impossible, and this is proven for the Bernoulli design only. Above the blue line, detection and recovery are both easy. See Section 3 for the formal statements.

(b) (Hard) If $c < c_{\mathrm{LD}}^{\mathrm{B}}$ then any degree-$o(k)$ polynomial fails to achieve weak separation. (This suggests that weak detection requires runtime $\exp(\tilde{\Omega}(k))$.)

We remark that $c_{\mathrm{LD}}^{\mathrm{B}}$ is a continuous function of θ (see Figure 2). The new algorithm that succeeds in the "easy" regime is based on counting the number of individuals whose degree (in the graph-theoretic sense) exceeds a particular threshold. For θ in the first case of (3.2), the low-degree hardness result requires a conditional argument that conditions away a certain rare "bad" event; for θ in the second case of (3.2), no conditioning is required and the resulting threshold matches the information-theoretic detection lower bound of [TAS20]. We remark that the predicted runtime $\exp(\tilde{\Omega}(k))$ in the "hard" regime is essentially tight, matching the runtime of the brute-force algorithm up to log factors in the exponent.

Next, we determine the precise information-theoretic detection boundary. One (inefficient) detection algorithm is the brute-force algorithm for optimal *recovery* (which can be made into a detection algorithm per Proposition C.1 in Appendix C). Another (efficient) detection algorithm is the low-degree algorithm from Theorem 3.3 above. We show that for each $\theta \in (0,1)$, statistically optimal detection is achieved by the better of these two algorithms. Brute-force is better when $\theta < 1 - \frac{\ln 2}{2\ln 2 - \ln\ln 2 - 1} \approx 0.079$, and otherwise low-degree is better.

Theorem 3.4. *Consider the Bernoulli design (testing variant) with parameters $\theta \in (0,1)$ and $c > 0$. Let $c_{\mathrm{inf}} := 1/\ln 2$ and define $c_{\mathrm{LD}}^{\mathrm{B}}$ as in (3.2).*

(a) (Possible) If $c > \min\{c_{\mathrm{inf}}, c_{\mathrm{LD}}^{\mathrm{B}}\}$ then strong detection is possible.

(b) (Impossible) If $c < \min\{c_{\mathrm{inf}}, c_{\mathrm{LD}}^{\mathrm{B}}\}$ then weak detection is impossible.

Hardness of Recovery Similarly to the constant-column design, our low-degree hardness results suggest hardness of recovery below the threshold $c_{\mathrm{LD}}^{\mathrm{B}}$ (see the discussion in Section 3.1). This suggests a possible-but-hard regime for recovery (namely Region I in Figure 2) in the Bernoulli design, for sufficiently small θ (namely $\theta < 1 - \frac{\ln 2}{2\ln 2 - \ln\ln 2 - 1} \approx 0.079$). As discussed in the Introduction, this is contrary to the evidence of [IZ21], who predicted the absence of a computational-statistical gap for all $\theta \in (0,1)$.

4 Background on Constant-Column Group Testing

4.1 General Setting

Recall that, in the underlying group testing instance, we start with n individuals out of which $k = n^{\theta}$ for fixed $\theta \in (0,1)$ are infected, and conduct

$$m = ck \ln\left(\frac{n}{k}\right) = ck(1-\theta)\ln n$$

parallel tests. We assume throughout that c is fixed with $0 < c < \ln^{-2}(2)$. (Strictly speaking we should write e.g. $k = n^{\theta + o(1)}$ due to integrality concerns, but for ease of notation we will drop these $o(1)$ terms.)

Let $\boldsymbol{G}_{GT} = (V_{GT} \cup F_{GT}, E_{GT})$ be a random bipartite graph with $|F_{GT}| = m$ *factor* nodes $(a_1, ..., a_m)$ representing the tests and $|V_{GT}| = n$ *variable* nodes $(x_1, ..., x_n)$ representing the individuals. Each individual independently chooses to participate in exactly $\Delta = c\ln(2)\ln(n/k)$ tests, chosen uniformly at random from the $\binom{m}{\Delta}$ possibilities. If x_i participates in test a_j, this is indicated by an edge between x_i and a_j. As usual, ∂a_j or ∂x_i denotes the neighbourhood of a vertex in $\boldsymbol{G}_{GT}$.

We let $\boldsymbol{\sigma} \in \{0,1\}^n$ denote the ground-truth vector encoding the infection status of each individual, uniformly chosen from all binary vectors of length n and Hamming weight k. Given $\boldsymbol{G}_{GT}$, we let $\hat{\boldsymbol{\sigma}} \in \{0,1\}^m$ denote the sequence of test results, that is

$$\hat{\boldsymbol{\sigma}}_a = \mathbb{1}\left\{\partial a \cap \{x : \boldsymbol{\sigma}(x) = 1\} \neq \emptyset\right\}.$$

We introduce a partition of the set of individuals into the following parts. We denote by $V_0(\boldsymbol{G}_{GT})$ the set of uninfected and by $V_1(\boldsymbol{G}_{GT})$ the set of infected individuals, formally

$$V_0(\boldsymbol{G}_{GT}) = \{x \in V_{GT} : \boldsymbol{\sigma}(x) = 0\} \quad \text{and} \quad V_1(\boldsymbol{G}_{GT}) = \{x \in V_{GT} : \boldsymbol{\sigma}(x) = 1\}.$$

Those individuals appearing in a negative test are *hard fields* and denoted by $V_0^-(\boldsymbol{G}_{GT})$ while the set $V_0^+(\boldsymbol{G}_{GT})$ consists of *disguised* uninfected individuals, that is uninfected individuals that only appear in positive tests:

$$\begin{aligned} & V_0^-(\boldsymbol{G}_{GT}) = \{x \in V_0(\boldsymbol{G}_{GT}) : \exists a \in \partial x : \hat{\boldsymbol{\sigma}}_a = 0\} \\ \text{and} \quad & V_0^+(\boldsymbol{G}_{GT}) = V_0(\boldsymbol{G}_{GT}) \setminus V_0^-(\boldsymbol{G}_{GT}). \end{aligned}$$

As previously mentioned, it is a straightforward task to identify those individuals that participate in a negative test and classify them as non-infected. Let $\boldsymbol{m}_0$ denote the number of tests rendering a negative result.

Lemma 4.1 (see [GJLR21], Lemmas A.4 & B.4)**.** *With high probability* $1 - o(1)$*, we have*

$$\boldsymbol{m_0} = \frac{m}{2} \pm O(\sqrt{m}\ln^2(n)) \quad \textit{and} \quad \left|V_0^+(\boldsymbol{G}_{GT})\right| = \left(1 \pm n^{-\Omega(1)}\right) n^{1-(1-\theta)c\ln^2(2)}.$$

Observe that as long as $c < \ln^{-2}(2)$, the number of disguised uninfected individuals clearly exceeds the number of infected individuals.

4.2 Reduced Setting

Now, we remove all $\boldsymbol{m}_0$ negative tests and their adjacent individuals from $\boldsymbol{G}_{GT}$ and are left with an reduced group testing instance $\boldsymbol{G}'_{GT}$ on $M = m - \boldsymbol{m_0}$ tests and $N = \left|V_0^+(\boldsymbol{G}_{GT})\right| + k$ individuals. Using Lemma 4.1 and the scaling of m, k, Δ we have with high probability,

$$M = \left(1 \pm n^{-\Omega(1)}\right)\frac{k\Delta}{2\ln 2} \qquad \text{and} \qquad N = \left(1 \pm n^{-\Omega(1)}\right) n^{1-(1-\theta)c\ln^2(2)}. \tag{4.1}$$

Let $\boldsymbol{\sigma}' \in \{0,1\}^N$ denote the restriction of $\boldsymbol{\sigma}$ to this reduced instance and observe that there are only positive tests remaining, which we re-label as $a_1, \ldots, a_M$.

5 Proof Roadmap for Theorem 3.1: "All-or-Nothing"

5.1 First Steps

We recall the setting of the theorem. Fix $\theta \in (0,1)$ and $c > 0$. Given n individuals $x_1, \ldots, x_n$, out of which $k = n^\theta$ are infected, and $m = ck\ln(n/k)$ tests $a_1, \ldots, a_m$, we denote by $\boldsymbol{\sigma} \in \{0,1\}^n$ the ground truth that encodes the infection status of the individuals. We create an instance of the constant-column pooling design $\boldsymbol{G}_{GT}$ as described in the previous section: each of the individuals independently chooses exactly $\Delta = c\ln(2)\ln(n/k)$ tests.

Suffices to study the posterior As described in the Introduction, it is known that if $c > 1/\ln(2)$ then approximate recovery is possible. For this reason, we focus here solely on the case $c < 1/\ln(2)$ with the goal of proving the "nothing" part of the all-or-nothing phenomenon, that is for any estimator $\boldsymbol{\tau} = \boldsymbol{\tau}(\boldsymbol{G}_{GT}) \in \{0,1\}^n$ it holds that $\langle \boldsymbol{\tau}, \boldsymbol{\sigma} \rangle = o(\|\boldsymbol{\tau}\|_2 \|\boldsymbol{\sigma}\|_2)$ with probability $1 - o(1)$. Our first observation is that it suffices to prove that the inner product between a draw from the posterior distribution $\boldsymbol{\sigma}|\boldsymbol{G}_{GT}$ and the ground truth $\boldsymbol{\sigma}$ is $o(k)$ in expectation, that is it suffices to prove

$$\mathop{\mathbb{E}}_{(\boldsymbol{\sigma}, \boldsymbol{G}_{GT})} \mathop{\mathbb{E}}_{\boldsymbol{\tau} \sim \boldsymbol{\sigma}|\boldsymbol{G}_{GT}} [\langle \boldsymbol{\tau}, \boldsymbol{\sigma} \rangle] = o(k). \tag{5.1}$$

Indeed, under (5.1) using the so-called "Nishimori identity" (see e.g. [NWZ21, Lemma 2]) and the Bayes optimality of the posterior mean, we have that *for any estimator* (with no norm restriction) $\boldsymbol{\tau} = \boldsymbol{\tau}(\boldsymbol{G}_{GT})$ it holds $\mathbb{E}[\|\boldsymbol{\tau} - \boldsymbol{\sigma}\|_2^2] = k(1 - o(1))$. The following lemma then gives the desired result.

Lemma 5.1. *Under our above assumptions, suppose that for any estimator $\boldsymbol{\tau} = \boldsymbol{\tau}(\boldsymbol{G}_{GT})$ it holds $\mathbb{E}[\|\boldsymbol{\tau} - \boldsymbol{\sigma}\|_2^2] = k(1 - o(1))$. Then for any estimator $\boldsymbol{\tau} = \boldsymbol{\tau}(\boldsymbol{G}_{GT})$ with $\|\boldsymbol{\tau}\|_2 = 1$ almost surely, it holds $\mathbb{E}[\langle \boldsymbol{\tau}, \boldsymbol{\sigma} \rangle]^2 = o(k) = o(\|\boldsymbol{\sigma}\|_2^2)$. In particular, for any estimator $\boldsymbol{\tau} = \boldsymbol{\tau}(\boldsymbol{G}_{GT}) \in \{0,1\}^n$ it holds that $\langle \boldsymbol{\tau}, \boldsymbol{\sigma} \rangle = o(\|\boldsymbol{\tau}\|_2 \|\boldsymbol{\sigma}\|_2)$ with probability $1 - o(1)$.*

Proof of Lemma 5.1. Fix any $\boldsymbol{\tau} = \boldsymbol{\tau}(\boldsymbol{G}_{GT})$ with $\|\boldsymbol{\tau}\|_2 = 1$ almost surely. Then for $\alpha := \mathbb{E}[\langle \boldsymbol{\tau}, \boldsymbol{\sigma} \rangle]$ we have that it must hold

$$\mathbb{E}[\|\alpha\boldsymbol{\tau} - \boldsymbol{\sigma}\|^2] = k(1 - o(1))$$

which implies,

$$\alpha^2 + k - 2\alpha\,\mathbb{E}[\langle \boldsymbol{\tau}, \boldsymbol{\sigma} \rangle] = k(1 - o(1))$$

and using the value of α we conclude

$$\mathbb{E}[\langle \boldsymbol{\tau}, \boldsymbol{\sigma} \rangle]^2 = o(k),$$

as we wanted. The lemma's final claim follows by normalizing $\boldsymbol{\tau}$ and using Markov's inequality. □

The posterior is uniform among "solutions" Now an easy computation using Bayes' rule gives that the posterior distribution is simply the uniform distribution over vectors $\sigma \in \{0,1\}^n$ with Hamming weight k that are *solutions* in the sense that every positive test contains at least one individual in the support of σ and none of the individuals in the support of σ participate in any negative tests. Therefore to prove (5.1), it suffices to show the following statement: with probability $1 - o(1)$ over $\boldsymbol{G}_{GT}$, a uniformly random solution for $\boldsymbol{G}_{GT}$ overlaps with the ground truth in at most $o(k)$ individuals.

Reducing the instance by removing negative tests We can simplify the problem by working with the reduced instance $\boldsymbol{G}'_{GT}$ defined in Section 4, where we have removed the negative tests and their adjacent individuals (so that only the positive tests remain). For simplicity in what follows, we re-label the individuals in $\boldsymbol{G}'_{GT}$ by $x_1, \dots, x_N$ and the tests by $a_1, \dots, a_M$. Recall that $\boldsymbol{\sigma}' \in \{0,1\}^N$ denotes the ground truth restricted to the individuals in $\boldsymbol{G}'_{GT}$. To show (5.1) it suffices to show that if $c < 1/\ln(2)$, a uniformly random "solution" *in the reduced model* overlaps with $\boldsymbol{\sigma}'$ in at most $o(k)$ individuals, with probability $1 - o(1)$. Here, with a slight abuse of notation, we define from now on a "solution" in $\boldsymbol{G}'_{GT}$ to be a vector $\sigma \in \{0,1\}^N$ of Hamming weight k with the property that each of the M (positive) tests in $\boldsymbol{G}'_{GT}$ contains at least one individual in the support of σ. Formally, we define the set of solutions $\boldsymbol{S} = \boldsymbol{S}(\boldsymbol{G}'_{GT})$ by

$$\boldsymbol{S} = \left\{ \sigma \in \binom{[N]}{k} \; : \; \max_{x \in \partial a_j} \sigma_x = 1 \text{ for all } j = 1, \dots, M \right\}. \tag{5.2}$$

As discussed above, (5.1), which implies the desired "nothing" result, follows by showing that almost all elements of $\boldsymbol{S}$ have a small *overlap*, in expectation, with the ground truth. In other words, since convergence in expectation and in probability are equivalent for bounded random variables, our new goal is to prove the following result.

Proposition 5.2. *Fix constants $0 < c < \ln^{-1}(2)$ and $\theta \in (0,1)$. Fix any constant $\delta > 0$ and let $\boldsymbol{\tau} \in \{0,1\}^N$ be uniformly sampled from $\boldsymbol{S}$. Then*

$$\Pr\left(\langle \boldsymbol{\sigma}', \boldsymbol{\tau} \rangle \geq \delta k\right) = o(1).$$

Here the probability is over both $\boldsymbol{G}'_{GT}$ and τ.

By the above discussion, Theorem 3.1 follows as a corollary of Proposition 5.2.

5.2 Proof Roadmap for Proposition 5.2: Two Null Models and their Roles

Now we describe the proof roadmap for Proposition 5.2 which completes the proof of Theorem 3.1. Here and in the following, we treat N, M as deterministic quantities lying in the "typical" range (4.1). We let $\mathbb{P}_\Delta$ denote the ("planted") distribution of the reduced instance $\boldsymbol{G}'_{GT}$ described in the previous section, conditioned on our chosen values of N, M. For an (N, M)-bipartite graph G, we let $\boldsymbol{Z}(G) := |\boldsymbol{S}(G)|$ denote the number of solutions in G as defined in (5.2). Furthermore, for the ground truth set of infected individuals $\boldsymbol{\sigma} \in \{0,1\}^N$ (since we will work exclusively in the reduced instance from now on, we simply write $\boldsymbol{\sigma}$ instead of $\boldsymbol{\sigma}'$) and some $\alpha \in (0,1]$, we let $\boldsymbol{Z}_{\boldsymbol{\sigma}}(G, \alpha)$ denote the number of solutions $\boldsymbol{\tau} \in \boldsymbol{S}$ with $\langle \boldsymbol{\tau}, \boldsymbol{\sigma} \rangle = \lfloor \alpha k \rfloor$.

First step In this notation, Proposition 5.2 asks that with probability $1 - o(1)$ over $G \sim \mathbb{P}_\Delta$,

$$\sum_{\delta k \le \ell \le k} \boldsymbol{Z}_{\boldsymbol{\sigma}}(G, \ell/k) = o(\boldsymbol{Z}(G)).$$

Notice that by Markov's inequality, it suffices to show that with probability $1 - o(1)$ over $G \sim \mathbb{P}_\Delta$,

$$\sum_{\delta k \le \ell \le k} \mathop{\mathbb{E}}_{\mathbb{P}_\Delta}[\boldsymbol{Z}_{\boldsymbol{\sigma}}(G, \ell/k)] = o(\boldsymbol{Z}(G)). \tag{5.3}$$

Unfortunately, direct calculations in the planted model $\mathbb{P}_\Delta$ are challenging. Towards establishing (5.3), we make use of two different "null" distributions over bipartite graphs with N individuals and M tests which are Δ-regular on the individuals side.

The Δ-Null Model First, we consider the Δ-null model $\mathbb{Q}_\Delta$ which is simply the measure on bipartite graphs with N individuals and M tests where each individual independently chooses exactly Δ tests uniformly at random (in particular, notice that no individual is assumed to be "infected").

The reason we introduce this model is because *the expected number of solutions of a graph G drawn from $\mathbb{Q}_\Delta$* offers a very simple high-probability lower bound on $\boldsymbol{Z}(G)$ for $G \sim \mathbb{P}_\Delta$. This is based on an application of the so-called *planting trick* introduced in the context of random k-SAT [ACO08]. The following lemma holds.

Lemma 5.3. *For any $\varepsilon > 0$,*

$$\mathbb{P}_\Delta \left\{ \boldsymbol{Z}(G) \le \varepsilon \mathop{\mathbb{E}}_{\mathbb{Q}_\Delta}[\boldsymbol{Z}(G)] \right\} \le \varepsilon.$$

In light of Lemma 5.3, to prove (5.3) it suffices to show

$$\sum_{\delta k \le \ell \le k} \mathop{\mathbb{E}}_{\mathbb{P}_\Delta}[\boldsymbol{Z}_{\boldsymbol{\sigma}}(G, \ell/k)] = o\left(\mathop{\mathbb{E}}_{\mathbb{Q}_\Delta}[\boldsymbol{Z}(G)] \right). \tag{5.4}$$

But now notice the following relation between $\mathbb{P}_\Delta$ and $\mathbb{Q}_\Delta$.

Fact 5.4. *One can generate a valid sample $(\boldsymbol{\sigma}, \boldsymbol{G}) \sim \mathbb{P}_\Delta$ by first choosing $\boldsymbol{\sigma} \in \{0,1\}^N$ uniformly from binary vectors of Hamming weight k, and then drawing $\boldsymbol{G}$ from $\mathbb{Q}_\Delta | \boldsymbol{\sigma}$, that is $\mathbb{Q}_\Delta$ conditioned on $\boldsymbol{\sigma}$ being a solution.*

Introducing the notation that for some $\alpha \in (0, 1]$ and a graph G we call $\boldsymbol{Z}(G, \alpha)$ the number of *pairs of solutions* $\boldsymbol{\tau}, \boldsymbol{\sigma} \in \boldsymbol{S}$ with $\langle \boldsymbol{\tau}, \boldsymbol{\sigma} \rangle = \lfloor \alpha k \rfloor$, we will use Fact 5.4 to prove the following "change-of-measure" lemma.

Lemma 5.5. *For any $\alpha \in (0, 1]$,*

$$\mathop{\mathbb{E}}_{\mathbb{P}_\Delta}[\boldsymbol{Z}_{\boldsymbol{\sigma}}(G, \alpha)] = \frac{\mathbb{E}_{\mathbb{Q}_\Delta}[\boldsymbol{Z}(G, \alpha)]}{\mathbb{E}_{\mathbb{Q}_\Delta}[\boldsymbol{Z}(G)]}.$$

Therefore, to prove (5.4) it suffices to show to Δ-null model property,

$$\sum_{\delta k \le \ell \le k} \mathop{\mathbb{E}}_{\mathbb{Q}_\Delta}[\boldsymbol{Z}(G, \ell/k)] = o\left(\mathop{\mathbb{E}}_{\mathbb{Q}_\Delta}[\boldsymbol{Z}(G)]^2 \right). \tag{5.5}$$

The (Δ,Γ)-Null Model Now, unfortunately it turns out that establishing (5.5) remains a highly technical task. Our way of establishing it is by considering another null model where the computations are easier, which we call the (Δ,Γ)-null model $\mathbb{Q}^{\star}_{\Delta,\Gamma}$. Here, instead of choosing Δ distinct tests (without replacement), each individual chooses Δ tests *with replacement*. Thus, under $\mathbb{Q}^{\star}_{\Delta,\Gamma}$ we allow (for technical reasons) the existence of *multi-edges*, as opposed to $\mathbb{P}_{\Delta}$ or $\mathbb{Q}_{\Delta}$. (Throughout, we will use an asterisk to signify models with multi-edges.) Also, we condition on every test having degree exactly $\Gamma = N\Delta/M$. Formally, $\mathbb{Q}^{\star}_{\Delta,\Gamma}$ is generated from the configuration model (see e.g. [JLR11]) over bipartite (multi-)graphs with N individuals, M tests, Δ degree for the individuals, and $\Gamma = N\Delta/M$ degree for the tests. Under $\mathbb{Q}_{\Delta}$, the test degrees concentrate tightly around Γ, and as a result we will be able to show that the models $\mathbb{Q}_{\Delta}$ and $\mathbb{Q}^{\star}_{\Delta,\Gamma}$ are "close." Specifically, this is formalized as follows.

Lemma 5.6. *For any fixed $0 < c < \ln^{-1}(2)$, $0 < \theta < 1$, and $\delta > 0$, it holds for all $\delta \le \alpha \le 1$ that*

$$\mathop{\mathbb{E}}_{\mathbb{Q}^{\star}_{\Delta,\Gamma}}[\boldsymbol{Z}(G)] \le \mathop{\mathbb{E}}_{\mathbb{Q}_{\Delta}}[\boldsymbol{Z}(G)]\exp\left(o(k\Delta)\right) \quad \text{and}$$
$$\mathop{\mathbb{E}}_{\mathbb{Q}^{\star}_{\Delta,\Gamma}}[\boldsymbol{Z}(G,\alpha)] \ge \mathop{\mathbb{E}}_{\mathbb{Q}_{\Delta}}[\boldsymbol{Z}(G,\alpha)]\exp\left(-o(k\Delta)\right).$$

Calculations in the configuration model are easier, yet still delicate, and allow us to prove the following result which given the above, concludes the proof of (5.5) and therefore of Proposition 5.2.

Proposition 5.7. *For any fixed $0 < c < \ln^{-1}(2)$, $0 < \theta < 1$, and $\delta > 0$, there exists $\varepsilon > 0$ such that the following holds for sufficiently large N. For all $\delta \le \alpha \le 1$,*

$$\frac{\mathbb{E}_{\mathbb{Q}^{\star}_{\Delta,\Gamma}}[\boldsymbol{Z}(G,\alpha)]}{\mathbb{E}_{\mathbb{Q}^{\star}_{\Delta,\Gamma}}[\boldsymbol{Z}(G)]^2} \le \exp(-\varepsilon k\Delta).$$

5.3 Proof of Lemmas 5.3 and 5.5

Proof of Lemma 5.3. Using Fact 5.4, note that $\mathbb{P}_{\Delta}(G)$ is proportional to $\boldsymbol{Z}(G)$, i.e.,

$$\mathbb{P}_{\Delta}(G) = \frac{\boldsymbol{Z}(G)\mathbb{Q}_{\Delta}(G)}{\mathbb{E}_{\mathbb{Q}_{\Delta}}[\boldsymbol{Z}(G)]}. \tag{5.6}$$

Set for simplicity $\lambda = \mathbb{E}_{\mathbb{Q}_{\Delta}}[\boldsymbol{Z}(G)]$. Using (5.6), we find

$$\begin{aligned}
\mathbb{P}_{\Delta}(\boldsymbol{Z}(G) \le \varepsilon\lambda) &= \sum_{G} \mathbb{1}\left\{\boldsymbol{Z}(G) \le \varepsilon\mathbb{E}_{\mathbb{Q}_{\Delta}}[\boldsymbol{Z}(G)]\right\} \frac{\boldsymbol{Z}(G)\mathbb{Q}_{\Delta}(G)}{\mathbb{E}_{\mathbb{Q}_{\Delta}}[\boldsymbol{Z}(G)]} \\
&\le \sum_{G} \mathbb{1}\left\{\boldsymbol{Z}(G) \le \varepsilon\mathbb{E}_{\mathbb{Q}_{\Delta}}[\boldsymbol{Z}(G)]\right\} \frac{\varepsilon\mathbb{E}_{\mathbb{Q}_{\Delta}}[\boldsymbol{Z}(G)]\mathbb{Q}_{\Delta}(G)}{\mathbb{E}_{\mathbb{Q}_{\Delta}}[\boldsymbol{Z}(G)]} \\
&\le \varepsilon \sum_{G} \mathbb{1}\left\{\boldsymbol{Z}(G) \le \varepsilon\lambda\right\} \mathbb{Q}_{\Delta}(G) \\
&= \varepsilon\, \mathbb{Q}_{\Delta}(\boldsymbol{Z}(G) \le \varepsilon\lambda) \\
&\le \varepsilon.
\end{aligned}$$

This concludes the proof. □

Proof of Lemma 5.5. Given Fact 5.4 and the symmetry of the individuals we have

$$\mathop{\mathbb{E}}_{\mathbb{P}_\Delta}[\boldsymbol{Z}_{\boldsymbol{\sigma}}(G,\alpha)] = \frac{1}{\binom{N}{k}} \sum_{\sigma,\sigma'} \mathbb{Q}_\Delta(\sigma' \in \boldsymbol{S}(G) \mid \sigma \in \boldsymbol{S}(G))$$

where the sum is over σ, σ' pairs with $\langle \sigma, \sigma' \rangle = \lfloor \alpha k \rfloor$

$$\begin{aligned} &= \frac{1}{\binom{N}{k}\mathbb{Q}_\Delta(\sigma \in \boldsymbol{S}(G))} \sum_{\sigma,\sigma'} \mathbb{Q}_\Delta(\sigma' \in \boldsymbol{S}(G), \sigma \in \boldsymbol{S}(G)) \\ &= \frac{\mathbb{E}_{\mathbb{Q}_\Delta}[\boldsymbol{Z}(G,\alpha)]}{\mathbb{E}_{\mathbb{Q}_\Delta}[\boldsymbol{Z}(G)]}. \end{aligned}$$

Note that with some abuse of notation we have pulled a term involving σ outside the sum; this is okay because (by symmetry) this term does not actually depend on σ. The proof is complete. □

6 Remaining Proofs from Section 5: The $\mathbb{Q}^\star_{\Delta,\Gamma}$ Model

6.1 Preliminaries: First and Second Moment under $\mathbb{Q}^\star_{\Delta,\Gamma}$

In this section we consider a bipartite graph drawn from $\mathbb{Q}^\star_{\Delta,\Gamma}$ on M tests $a_1, \ldots, a_M$ of size exactly Γ each and N individuals $x_1, \ldots, x_N$ of degree exactly Δ. Recall that this graph is generated from the configuration model and may feature multi-edges.

Our first result is about the first moment of the number of solutions.

Lemma 6.1. *Let $q \in (0,1)$ be the solution to the equation*

$$\frac{q}{1-(1-q)^\Gamma} = \frac{\Delta k}{\Gamma M}. \tag{6.1}$$

Then

$$\mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[\boldsymbol{Z}(G)] = N^{-O(1)} \binom{N}{k} \frac{(1-(1-q)^\Gamma)^M}{\binom{\Gamma M}{\Delta k} q^{\Delta k}(1-q)^{\Gamma M - \Delta k}}. \tag{6.2}$$

We now present in some detail the proof of Lemma 6.1 since it is a good first example of the technique we follow for the computations in this section.

Proof. By linearity of expectation and symmetry, notice that for any fixed configuration $\sigma \in \{0,1\}^N$ with Hamming weight k, it holds that

$$\mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[\boldsymbol{Z}(G)] = \binom{N}{k} \mathbb{Q}^\star_{\Delta,\Gamma}[\sigma \in \boldsymbol{S}(G)].$$

We now calculate the probability $\mathbb{Q}^\star_{\Delta,\Gamma}[\sigma \in \boldsymbol{S}(G)]$ as follows. We first set up an auxiliary product probability space. Fix *any* parameter $q \in (0,1)$. Construct a product probability space with measure $\mathbb{P}_q$ where we choose ΓM bits $(\boldsymbol{\omega}_{ij})_{i\in[M], j\in[\Gamma]}$ independently such that $\boldsymbol{\omega}_{ij} \sim \mathrm{Ber}(q)$ for

all i, j. (It may help to think of ω_{ij} as representing the infection status of the jth individual in the ith test.) Let $\boldsymbol{R} = \sum_{i,j} \boldsymbol{\omega}_{ij}$ be the total number of ones. Let us define

$$\mathcal{S} = \left\{ \forall i \in [M] : \max_j \boldsymbol{\omega}_{ij} = 1 \right\} \qquad \mathcal{R} = \{\boldsymbol{R} = k\Delta\}. \tag{6.3}$$

But then notice that in this notation the symmetry of the product space gives that *for any* $q \in (0,1)$,

$$\mathbb{Q}^{\star}_{\Delta,\Gamma}[\sigma \in \boldsymbol{S}(G)] = \mathbb{P}_q[\mathcal{S} \mid \mathcal{R}].$$

One can then calculate this conditional probability via Bayes. The unconditional probabilities are easy to compute:

$$\mathbb{P}_q[\mathcal{S}] = (1-(1-q)^{\Gamma})^M, \qquad \mathbb{P}_q[\mathcal{R}] = \binom{\Gamma M}{\Delta k} q^{\Delta k}(1-q)^{\Gamma M - \Delta k}.$$

A priori, the conditional probability $\mathbb{P}_q[\mathcal{R} \mid \mathcal{S}]$ may be difficult to compute and this is where our freedom to choose q becomes important. Specifically, we pick q as in (6.1). By the local limit theorem for sums of independent random variables (see for instance [COHKL+21, Section 6]), this choice ensures that

$$\mathbb{E}[\boldsymbol{R} \mid \mathcal{S}] = \Gamma M \frac{q}{1-(1-q)^{\Gamma}} = \Delta k \qquad \text{and therefore} \qquad \mathbb{P}[\mathcal{R} \mid \mathcal{S}] = N^{-O(1)}.$$

Bayes' theorem now completes the proof of the lemma. □

Using a multidimensional version of the idea that allowed us to calculate the first moment bound we develop the second moment bound by modelling the pairs of configurations via independent random variables. We derive the appropriate probabilities for an "independent" problem setting and then tackle the dependencies afterwards by applying Bayes' formula.

Recall the definition

$$\boldsymbol{Z}(G,\alpha) = |\{\sigma, \tau \in \boldsymbol{S}(G) : \langle \sigma, \tau \rangle = \alpha k\}|$$

denote the number of pairs of solutions that overlap on an α-fraction of entries. We are able to obtain the following sharp bound on the expectation of $\boldsymbol{Z}(G,\alpha)$.

Lemma 6.2. *For any* $\alpha \in (0,1]$ *and any* $(q_{00}, q_{01}, q_{10}, q_{11}) \in [0,1]^4$,

$$\begin{aligned} \mathop{\mathbb{E}}_{\mathbb{Q}^{\star}_{\Delta,\Gamma}}[\boldsymbol{Z}(G,\alpha)] \leq & \binom{N}{\alpha k,\ (1-\alpha)k,\ (1-\alpha)k} \\ & \cdot \frac{\left(1-2(1-q_{01}-q_{11})^{\Gamma}+q_{00}^{\Gamma}\right)^M}{\binom{N\Delta}{\alpha k\Delta,\,(1-\alpha)k\Delta,\,(1-\alpha)k\Delta,\,(N-2k+\alpha k)\Delta} q_{11}^{\alpha k\Delta} q_{10}^{2(k-\alpha k)\Delta} q_{00}^{N\Delta-2k\Delta+\alpha k\Delta}}. \end{aligned} \tag{6.4}$$

Furthermore, if $(q_{00}, q_{01}, q_{10}, q_{11}) \in [0,1]^4$ *is the solution to the system*

$$q_{00}+q_{01}+q_{10}+q_{11} = 1 \qquad q_{01} = q_{10} \tag{6.5}$$

$$\frac{q_{11}}{1-2(1-q_{10}-q_{11})^{\Gamma}+q_{00}^{\Gamma}} = \alpha\frac{k\Delta}{\Gamma M} \qquad \frac{q_{01}\left(1-(q_{00}+q_{10})^{\Gamma-1}\right)}{1-2(1-q_{01}-q_{11})^{\Gamma}+q_{00}^{\Gamma}} = (1-\alpha)\frac{k\Delta}{\Gamma M} \tag{6.6}$$

then

$$\mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[\boldsymbol{Z}(G,\alpha)] = N^{-O(1)}\binom{N}{\alpha k,\,(1-\alpha)k,\,(1-\alpha)k} \cdot \frac{\left(1-2(1-q_{01}-q_{11})^\Gamma+q_{00}^\Gamma\right)^M}{\binom{N\Delta}{\alpha k\Delta,\,(1-\alpha)k\Delta,\,(1-\alpha)k\Delta,\,(N-2k+\alpha k)\Delta} q_{11}^{\alpha k\Delta} q_{10}^{2(k-\alpha k)\Delta} q_{00}^{N\Delta-2k\Delta+\alpha k\Delta}}. \tag{6.7}$$

Proof. The multinomial coefficient simply counts assignments so that the pair of configurations has the correct overlap. Hence, let us fix a pair (σ,τ) with overlap α. As before we employ an auxiliary probability space $(\boldsymbol{\omega}_{ij},\boldsymbol{\omega}'_{ij})_{i\in[M],\,j\in[\Gamma]}$ with independent entries drawn from the distribution $(q_{00},\dots,q_{11})$, e.g., q_{01} is the probability that $\boldsymbol{\omega}_{ij}=0$ and $\boldsymbol{\omega}'_{ij}=1$. (We think of $\boldsymbol{\omega}_{ij}$ as the infection status of the jth individual in the ith test under σ, and $\boldsymbol{\omega}'_{ij}$ is the same for τ.) Let $\mathcal{S}$ be the event that all tests are positive under both assignments and let $\mathcal{R}$ be the event that

$$\sum_{i,j}\boldsymbol{\omega}_{ij}=\sum_{i,j}\boldsymbol{\omega}'_{ij}=k\Delta \qquad \text{and} \qquad \sum_{i,j}\boldsymbol{\omega}_{ij}\boldsymbol{\omega}'_{ij}=\alpha k\Delta.$$

Then

$$\begin{aligned}\mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[\boldsymbol{Z}(G,\alpha)] &= \binom{N}{\alpha k,\,(1-\alpha)k,\,(1-\alpha)k}\mathbb{P}\left[\mathcal{S}\mid\mathcal{R}\right]\\ &= \binom{N}{\alpha k,\,(1-\alpha)k,\,(1-\alpha)k}\frac{\mathbb{P}\left[\mathcal{S}\right]\mathbb{P}\left[\mathcal{R}\mid\mathcal{S}\right]}{\mathbb{P}\left[\mathcal{R}\right]}.\end{aligned}$$

Once again we use Bayes' rule. The unconditional probabilities are easy:

$$\begin{aligned}\mathbb{P}\left[\mathcal{R}\right] &= \binom{N\Delta}{\alpha k\Delta,\,(1-\alpha)k\Delta,\,(1-\alpha)k\Delta} q_{11}^{\alpha k\Delta} q_{10}^{2(k-\alpha k)\Delta} q_{00}^{N\Delta-2k\Delta+\alpha k\Delta},\\ \mathbb{P}\left[\mathcal{S}\right] &= \left(1-2(1-q_{01}-q_{11})^\Gamma+q_{00}^\Gamma\right)^M.\end{aligned}$$

Using the fact $\mathbb{P}\left[\mathcal{R}\mid\mathcal{S}\right]\le 1$, we can conclude (6.4). Now we also claim that with the choice (6.5)-(6.6),

$$\mathbb{P}\left[\mathcal{R}\mid\mathcal{S}\right]=N^{-O(1)}.$$

As before, this follows from the local limit theorem for sums of independent random variables, provided we can show

$$\mathbb{E}\left[\sum_{i,j}\boldsymbol{\omega}_{ij}\,\middle|\,\mathcal{S}\right]=\mathbb{E}\left[\sum_{i,j}\boldsymbol{\omega}'_{ij}\,\middle|\,\mathcal{S}\right]=k\Delta, \qquad \mathbb{E}\left[\sum_{i,j}\boldsymbol{\omega}_{ij}\boldsymbol{\omega}'_{ij}\,\middle|\,\mathcal{S}\right]=\alpha k\Delta. \tag{6.8}$$

The second equation in (6.8) is easy to compute because any test that contains a $(1,1)$ will instantly be satisfied under both assignments:

$$\mathbb{E}\left[\sum_{i,j}\boldsymbol{\omega}_{ij}\boldsymbol{\omega}'_{ij}\,\middle|\,\mathcal{S}\right]=\frac{\Gamma M q_{11}}{1-2(1-q_{01}-q_{11})^\Gamma+q_{00}^\Gamma}.$$

For the first equation in (6.8), it suffices to show

$$\mathbb{E}\left[\sum_{i,j} \boldsymbol{\omega}_{ij} - \boldsymbol{\omega}_{ij}\boldsymbol{\omega}'_{ij} \;\middle|\; \mathcal{S}\right] = (1-\alpha)k\Delta.$$

If a test contains a $(1,0)$ then it still requires either a $(1,1)$ or a $(0,1)$ to be satisfied under the other assignment as well:

$$\mathbb{E}\left[\sum_{i,j} \boldsymbol{\omega}_{ij} - \boldsymbol{\omega}_{ij}\boldsymbol{\omega}'_{ij} \;\middle|\; \mathcal{S}\right] = \frac{\Gamma M q_{10}\left(1-(q_{00}+q_{01})^{\Gamma-1}\right)}{1-2(1-q_{10}-q_{11})^{\Gamma}+q_{00}^{\Gamma}}.$$

In any case, the choice (6.5)-(6.6) gives what we want. □

6.2 Proof of Proposition 5.7

To prove Proposition 5.7, we need to compare the first moment squared and (part of) the second moment expansion under $\mathbb{Q}^{\star}_{\Delta,\Gamma}$. We begin with a bound on the first moment.

6.2.1 Bound on First Moment

As we have a multiplicative factor $\exp\left(o\left(k\Delta\right)\right)$ of freedom, the result of the following proposition will suffice.

Proposition 6.3. *It holds that*

$$\mathop{\mathbb{E}}_{\mathbb{Q}^{\star}_{\Delta,\Gamma}}\left[\boldsymbol{Z}(G)\right] = \exp\left(o\left(k\Delta\right)\right)\exp\left(k\Delta\frac{1-c\ln(2)}{c\ln(2)}\right).$$

Proof. Our starting point is Lemma 6.1. Recall $\Gamma M = N\Delta$. Define $d > 0$ such that $q = d\frac{k}{N}$ and recall that $\Gamma = \left(2\ln 2 \pm n^{-\Omega(1)}\right)\frac{N}{k}$. Therefore (6.1) is equivalent to

$$1-\exp\left(-2d\ln 2\left(1\pm n^{-\Omega(1)}\right)\right) = d.$$

Therefore, the unique solution $\hat{q}$ to (6.1) turns out to be

$$\hat{q} = \left(1\pm n^{-\Omega(1)}\right)\frac{k}{2N}. \tag{6.9}$$

Furthermore observe for the binomial coefficients needed in Lemma 6.1 that Stirling's formula (Lemma A.1) implies

$$\binom{N\Delta}{k\Delta} = (1+o(1))\frac{1}{\sqrt{2\pi k\Delta}}\left(\frac{Ne}{k}\right)^{k\Delta} \quad \text{and} \quad \binom{N}{k} = (1+o(1))\frac{1}{\sqrt{2\pi k}}\left(\frac{Ne}{k}\right)^{k}. \tag{6.10}$$

Finally, recall the scaling

$$M = \left(1\pm N^{-\Omega(1)}\right)\frac{k\Delta}{2\ln(2)}. \tag{6.11}$$

The proposition follows from plugging (6.9), (6.10) and (6.11) into (6.2) from Lemma 6.1. □

6.2.2 Bound on Second Moment

We will bound the expression for $\mathbb{E}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[\boldsymbol{Z}(G,\alpha)]$ given in Lemma 6.2. Lemma 6.2 yields

$$\mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[\boldsymbol{Z}(G,\alpha)] \leq \binom{N}{\alpha k,\,(1-\alpha)k,\,(1-\alpha)k} \cdot \frac{\left(1-2(1-q_{01}-q_{11})^\Gamma+q_{00}^\Gamma\right)^M}{\binom{N\Delta}{\alpha k\Delta,\,(1-\alpha)k\Delta,\,(1-\alpha)k\Delta,\,(N-2k+\alpha k)\Delta} q_{11}^{\alpha k\Delta} q_{10}^{2(k-\alpha k)\Delta} q_{00}^{N\Delta-2k\Delta+\alpha k\Delta}}.$$

For $\alpha \in (0,1]$, define

$$(q_{00}=q_{00}(\alpha),\, q_{01}=q_{01}(\alpha),\, q_{10}=q_{10}(\alpha),\, q_{11}=q_{11}(\alpha)) \in [0,1]^4$$

to be the solution of (6.5)-(6.6). Using the first two equations of (6.5)-(6.6) it suffices to only keep track of q_{01}, q_{11} because q_{00}, q_{10} are simple linear functions of them.

To this end, define

$$\begin{aligned} G(\alpha,q_{01},q_{11}) = k\Delta\Bigg(&\alpha\ln(\alpha)+2(1-\alpha)\ln(1-\alpha)-(2-\alpha)+(2-\alpha)\frac{1-c\ln^2(2)}{c\ln(2)} \\ &+\frac{1}{2\ln(2)}\ln\left(1-2(1-q_{01}-q_{11})^\Gamma+(1-2q_{01}-q_{11})^\Gamma\right) \\ &-\alpha q_{11}-(2-\alpha)q_{01}\Bigg) \\ &-(N\Delta-2k\Delta+\alpha k\Delta)\ln(1-2q_{01}-q_{11}). \end{aligned}$$

By Stirling's formula this is, up to $o(k\Delta)$ additive error terms, equal to the exponential part of $\mathbb{E}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[\boldsymbol{Z}(G,\alpha)]$ from Lemma 6.2. Indeed,

$$\begin{aligned} G(\alpha,q_{01},q_{11}) = o(\Delta k) + \ln\Bigg(&\binom{N}{\alpha k,\,(1-\alpha)k,\,(1-\alpha)k} \\ &\cdot\frac{\left(1-2(1-q_{01}-q_{11})^\Gamma+q_{00}^\Gamma\right)^M}{\binom{N\Delta}{\alpha k\Delta,\,(1-\alpha)k\Delta,\,(1-\alpha)k\Delta,\,(N-2k+\alpha k)\Delta} q_{11}^{\alpha k\Delta} q_{10}^{2(k-\alpha k)\Delta} q_{00}^{N\Delta-2k\Delta+\alpha k\Delta}}\Bigg). \end{aligned} \tag{6.12}$$

The purpose of this approximation is that the function G can be analysed analytically.

Lemma 6.4. *For any $c < \ln^{-1}(2)$ and any $\theta \in (0,1)$, there exists $\varepsilon > 0$ such that for all $\dot\alpha \in (0,1]$,*

$$G\Big(\dot\alpha, q_{01}(\dot\alpha), q_{11}(\dot\alpha)\Big) < (1-\varepsilon)k\Delta\frac{2(1-c\ln(2))}{c\ln(2)}.$$

Proof. As a first step, we need to determine q_{01}, q_{11} from (6.5)-(6.6) for a general $\dot\alpha \in (0,1]$. We define $x_0, x_1 > 0$ such that

$$q_{01} = x_0\frac{k}{N} \qquad \text{and} \qquad q_{11} = x_1\frac{k}{N}$$

and define

$$\mathcal{W}(x_0, x_1) = 1 - 2\exp\left(-2\ln(2)(x_0 + x_1)\right) + \exp\left(-2\ln(2)(2x_0 + x_1)\right).$$

This allows us to simplify (6.6) to

$$\alpha = \frac{x_1}{\mathcal{W}(x_0, x_1)} \qquad \text{and} \qquad 1 - \alpha = \frac{x_0\left(1 - \exp\left(-2\ln(2)(x_0 + x_1)\right)\right)}{\mathcal{W}(x_0, x_1)}. \tag{6.13}$$

If we plug in (6.13) into the definition of G, we get

$$\begin{aligned} &G(\alpha, q_{01}, q_{11}) \\ &\quad = (1 + o(1))k\Delta\left(\alpha\ln\left(\frac{\alpha}{x_1}\right) + 2(1-\alpha)\ln\left(\frac{1-\alpha}{x_0}\right) + (2-\alpha)\frac{1 - c\ln^2(2)}{c\ln(2)}\right) \\ &\qquad + k\Delta\left(\frac{1}{2\ln(2)}\ln\left(\mathcal{W}(x_0, x_1)\right) + (2x_0 + x_1) - (2-\alpha)\right). \end{aligned} \tag{6.14}$$

While it is easy for a given $\dot{\alpha}$ to determine the solution $(\dot{x}_0, \dot{x}_1)$ of (6.13) numerically, it seems impossible to come up with an analytic closed form expression. Fortunately, by the first part of Lemma 6.2 this is not necessary. Indeed, *any* choice (x_0, x_1) for a given $\dot{\alpha}$ renders an upper bound on (6.14) as this is the leading order part of $\mathbb{E}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[Z(\boldsymbol{G}, \alpha)]$. Specifically, recall from (6.12) that $G(\alpha, q_{01}, q_{11})$ approximates the exponential part of $\mathbb{E}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[Z(\boldsymbol{G})]$ up to an additive error of $o\left(k\Delta\right)$.

We approximate $(\dot{x}_0, \dot{x}_1)$ by a piecewise linear function. Define the following partition of $(0, 1)$:

$$I_1 = \left(0, \frac{1}{4}\right], \; I_2 = \left(\frac{1}{4}, \frac{85}{100}\right), \; I_3 = \left[\frac{85}{100}, 1\right). \tag{6.15}$$

We define

$$x_0(\alpha) = \mathbb{1}_{\{\alpha \in I_1\}} \cdot \left(-\frac{3}{5}\alpha + \frac{1}{2}\right) + \mathbb{1}_{\{\alpha \in I_2\}} \cdot \left(\frac{1}{2} - \frac{3}{10\ln 2}\alpha\right) + \mathbb{1}_{\{\alpha \in I_3\}} \cdot (1 - \alpha), \tag{6.16}$$

$$x_1(\alpha) = \mathbb{1}_{\{\alpha \in I_1\}} \cdot \frac{\alpha}{5} + \mathbb{1}_{\{\alpha \in I_2\}} \cdot \frac{\alpha}{5\ln 2} - \mathbb{1}_{\{\alpha \in I_3\}} \cdot \frac{16\alpha - 11}{10}. \tag{6.17}$$

For brevity, let

$$\begin{aligned} F(\alpha) &= \left(\alpha\ln\left(\frac{\alpha}{x_1}\right) + 2(1-\alpha)\ln\left(\frac{1-\alpha}{x_0}\right) + (2-\alpha)\frac{1 - c\ln^2(2)}{c\ln(2)}\right) \\ &\quad + \left(\frac{1}{2\ln(2)}\ln\left(\mathcal{W}(x_0, x_1)\right) + (2x_0 + x_1) - (2-\alpha)\right) \qquad (6.18) \\ &= G\left(\alpha, x_0\frac{k}{N}, x_1\frac{k}{N}\right)\frac{1 + o(1)}{k\Delta}. \qquad (6.19) \end{aligned}$$

We will bound each piece of F separately, with the goal of establishing the bound

$$F(\alpha) < \frac{2(1 - c\ln(2))}{c\ln(2)} \qquad \text{for all } \alpha \in (0, 1]. \tag{6.20}$$

An illustration of the result of the considered cases can be found in Figure 3.

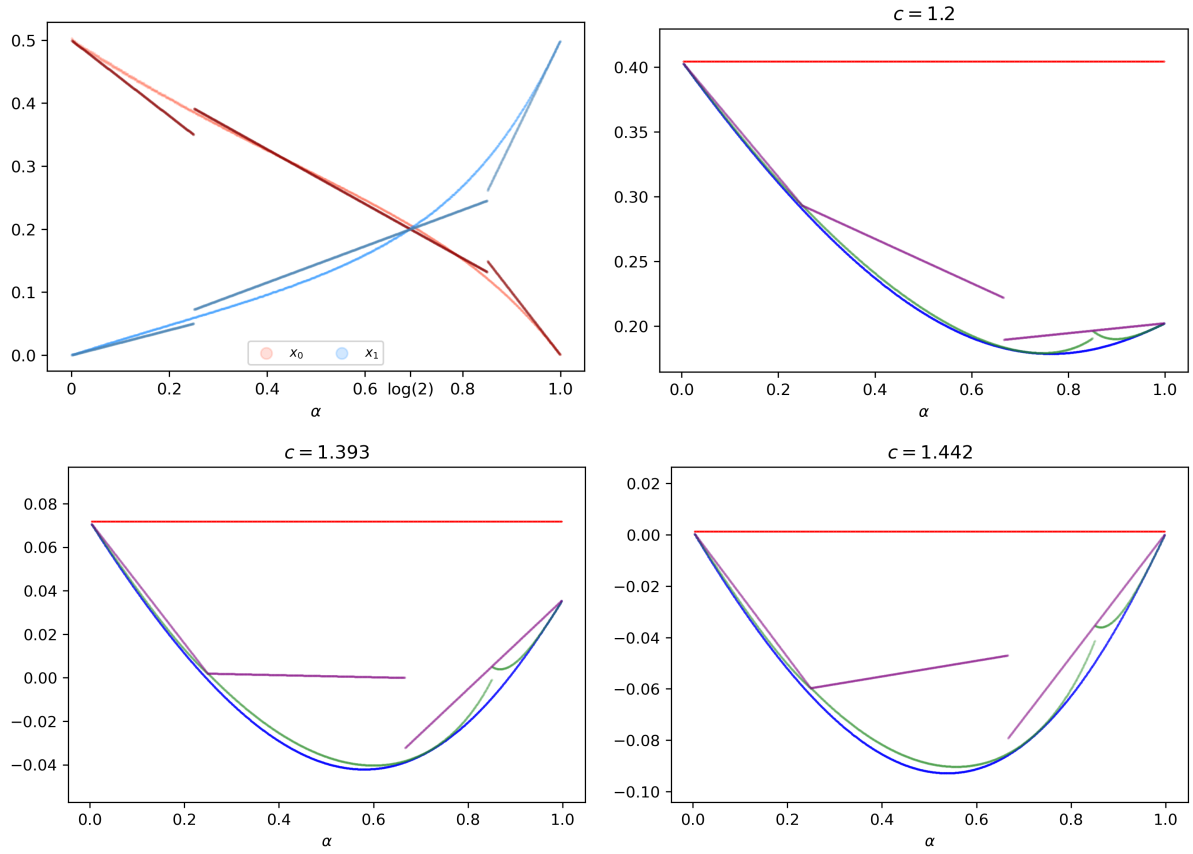

Figure 3: The first plot shows a numerical comparison between the optimal choices (x_0, x_1) and our piece-wise linear approximation. The other plots show how the evaluation of $G\left(\alpha, x_0\frac{k}{N}, x_1\frac{k}{N}\right)$ varies between the numerically calculated optimal values (blue), the linear approximation of (x_0, x_1) applied to $G\left(\alpha, x_0\frac{k}{N}, x_1\frac{k}{N}\right)$ (green) and the easily established upper bound on this quantity through convexity (purple) for different values of $c \in (0, \ln^{-1}(2)]$. The red line equals $\frac{2(1-c\ln(2))}{c\ln(2)}$.

Case $\alpha \in I_1$ **:** In this case, (6.19) reads as

$$F(\alpha) = \alpha\ln(5) + 2(1-\alpha)\ln(1-\alpha) - 2(1-\alpha)\ln\left(\frac{1}{2} - \frac{3}{5}\alpha\right) + (2-\alpha)\frac{1-c\ln^2(2)}{c\ln(2)}$$
$$+ \frac{1}{2\ln(2)}\ln\left(1 - 2\exp\left(-2\ln(2)\left(-\frac{2}{5}\alpha + \frac{1}{2}\right)\right) + \exp\left(-2\ln(2)(1-\alpha)\right)\right) - 1.$$

We find for any $c \in (0, \ln^{-1}(2))$ that

$$\frac{\partial^2 F}{\partial\alpha^2} = \frac{2}{1-\alpha} + \frac{0.72(1-\alpha)}{(-0.6\alpha+0.5)^2} + \frac{2.4}{0.6\alpha-0.5} - \frac{1}{2}\frac{(2^{2\alpha-1} - 1.6\cdot 2^{0.8\alpha-1.0})^2\ln(2)}{(2^{0.8\alpha} - 2^{2\alpha-2} - 1)^2}$$
$$- \frac{\ln(2)}{2}\cdot\frac{2^{2\alpha} - 1.28\cdot 2^{0.8\alpha-1}}{(2^{0.8\alpha} - 2^{2\alpha-2} - 1)} > 0$$

which can be verified analytically (for illustration see Figure 4). To see this we analyse two separate parts. On the one hand,

$$\frac{2}{1-\alpha}+\frac{0.72(1-\alpha)}{(-0.6\alpha+0.5)^2}+\frac{2.4}{0.6\alpha-0.5}>0.$$

On the other hand one can verify that the remainder satisfies

$$-\frac{\ln(2)}{2}\left(\frac{(2^{2\alpha-1}-1.6\cdot 2^{0.8\alpha-1.0})^2}{(2^{0.8\alpha}-2^{2\alpha-2}-1)^2}+\frac{2^{2\alpha}-1.28\cdot 2^{0.8\alpha-1}}{(2^{0.8\alpha}-2^{2\alpha-2}-1)}\right)>0,$$

as

$$(2^{2\alpha-1}-1.6\cdot 2^{0.8\alpha-1.0})^2+\left(2^{2\alpha}-1.28\cdot 2^{0.8\alpha-1}\right)\left(2^{0.8\alpha}-2^{2\alpha-2}-1\right)<-\frac{1}{3}\alpha<0.$$

In particular, $\frac{\partial^2 F}{\partial \alpha^2}$ does not depend on c and is monotonically increasing on I_1. Therefore, F is strictly convex on I_1, and so it suffices to verify (6.20) at the endpoints of I_1. We will apply a first-order Taylor approximation to F at $\alpha=0$. Let $\tilde{F}$ be this approximation. The following holds by Taylor's theorem. For any $\varepsilon>0$ there is $\delta>0$ with the property that

$$F(\alpha)\leq(1+\delta)\tilde{F}(\alpha)\qquad\text{for all }\alpha\in(0,\varepsilon).\tag{6.21}$$

We have

$$\tilde{F}(\alpha)=\frac{\left(\left(5\,\ln(5)\ln(2)-5\,\ln(2)^2-\ln(2)\right)\alpha-10\,\ln(2)\right)c-5\,\alpha+10}{5\,c\ln(2)}.$$

Therefore,

$$\tilde{F}(\alpha)-\frac{2(1-c\ln(2))}{c\ln(2)}=\frac{\left(5\,\ln(5)\ln(2)-5\,\ln(2)^2-\ln(2)\right)\alpha c-5\,\alpha}{5\,c\ln(2)}.$$

Therefore, by (6.21) we only need to verify that there is that there is $\delta'>0$ and $\alpha^\star>0$ such that for all $\alpha\in(0,\alpha^\star)$ and $c<\ln^{-1}(2)$, we have

$$\left(5\,\ln(5)\ln(2)-5\,\ln(2)^2-\ln(2)\right)c-5<-\delta'(\alpha)^{-1}.$$

As $\left(5\,\ln(5)\ln(2)-5\,\ln(2)^2-\ln(2)\right)\approx 2.48$, the strongest requirement is given for $c=\ln^{-1}(2)$ and is satisfied if $\alpha^\star>\delta'/1.4$. Furthermore, it can be verified that

$$\lim_{\alpha\to 0.25}F(\alpha)=\frac{\ln\left(-\frac{1}{4}\sqrt{2}\left(2\sqrt{2}\left(2^{\frac{1}{5}}-1\right)-1\right)\right)}{2\,\ln(2)}+\frac{7}{4\,c\ln(2)}+\frac{1}{4}\ln(5)-\frac{7}{4}\ln(2)+\frac{3}{2}\ln\left(\frac{3}{4}\right)$$
$$-\frac{3}{2}\ln\left(\frac{7}{20}\right)-1<\frac{2(1-c\ln(2))}{c\ln(2)}$$

for any $c\in(0,\ln^{-1}(2))$, thus, (6.20) is satisfied on I_1.

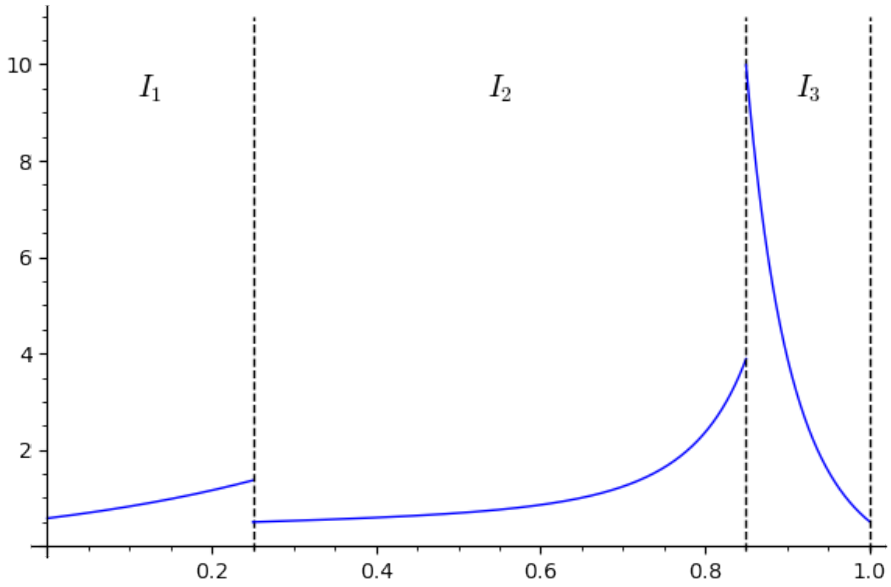

Figure 4: The piece-wise defined second derivative $\frac{\partial^2 F}{\partial \alpha^2}$ on the three intervals I_1, I_2, I_3. As could be seen analytically, it does not depend on c but is a (piece-wise) continuous mapping of α.

Case $\alpha \in I_2$: We have

$$\begin{aligned} F(\alpha) = {} & \alpha \ln(\alpha) - \alpha \ln(\alpha) + \alpha \ln(5 \ln(2)) \\ & + 2(1-\alpha)\ln(1-\alpha) - 2(1-\alpha)\ln\left(0.5 - 0.3 \cdot \frac{1}{\ln(2)}\alpha\right) \\ & + \frac{1}{2\ln(2)} \ln\left(1 - 2\exp\left(-2\ln(2)\left(\frac{1}{2} - \frac{1}{10\ln(2)}\alpha\right)\right)\right. \\ & \qquad \left. + \exp\left(-2\ln(2)\left(1 - \frac{2}{5\ln(2)}\alpha\right)\right)\right) \\ & + (2-\alpha)\frac{1 - c\ln^2(2)}{c\ln(2)} + 1 - \frac{2}{5\ln(2)}\alpha - 2 + \alpha. \end{aligned}$$

In this case,

$$\begin{aligned} \frac{\partial^2 F}{\partial \alpha^2} = {} & \frac{2}{1-\alpha} - \frac{1}{2} \cdot \frac{(0.8 \cdot 2^{0.8\alpha/\ln(2)-2} - 0.4 \cdot 2^{0.2\alpha/\ln(2)-1})^2}{(2^{0.8\alpha/\ln(2)-2} - \exp(0.2\alpha) + 1)^2 \ln(2)} \\ & + \frac{1}{2} \cdot \frac{0.64 \cdot 2^{0.8\alpha/\ln(2)-2} - 0.08 \cdot 2^{0.2\alpha/\ln(2)-1}}{(2^{0.8\alpha/\ln(2)-2} - \exp(0.2\alpha) + 1)\ln(2)} \\ & - \frac{1.2}{(-0.3\alpha/\ln(2) + 0.5)\ln(2)} - \frac{0.18\alpha - 0.18}{(-0.3\alpha/\ln(2) + 0.5)^2 \ln(2)^2)} > 0. \end{aligned}$$

We again verify this by analysing two separate parts. On the one hand one can verify that

$$\frac{2}{1-\alpha} - \frac{1.2}{(-0.3\alpha/\ln(2) + 0.5)\ln(2)} - \frac{0.18\alpha - 0.18}{(-0.3\alpha/\ln(2) + 0.5)^2 \ln(2)^2)} > 0, \tag{6.22}$$

as this can be rearranged to

$$\frac{9}{50}\alpha^2 + \frac{1}{2}\left(\ln(2) + \frac{3}{5}\right)^2 > 0.$$

Now we turn to the second part which reads as follows:

$$-\frac{1}{2\ln(2)} \cdot \left(\frac{(0.8 \cdot 2^{0.8\alpha/\ln(2)-2} - 0.4 \cdot 2^{0.2\alpha/\ln(2)-1})^2}{(2^{0.8\alpha/\ln(2)-2} - \exp(0.2\alpha) + 1)^2} - \frac{0.64 \cdot 2^{0.8\alpha/\ln(2)-2} - 0.08 \cdot 2^{0.2\alpha/\ln(2)-1}}{(2^{0.8\alpha/\ln(2)-2} - \exp(0.2\alpha) + 1)} \right) \tag{6.23}$$

Thus, we show that

$$\Bigg((0.8 \cdot 2^{0.8\alpha/\ln(2)-2} - 0.4 \cdot 2^{0.2\alpha/\ln(2)-1})^2 \\ - \left(0.64 \cdot 2^{0.8\alpha/\ln(2)-2} - 0.08 \cdot 2^{0.2\alpha/\ln(2)-1}\right)\left(2^{0.8\alpha/\ln(2)-2} - \exp(0.2\alpha) + 1\right) \Bigg) < 0.$$

The assertion immediately follows as the latter product exceeds the quadratic expression for all $\alpha \in \left(\frac{1}{4}, \frac{85}{100}\right]$ and all three parts are positive. Thus (6.23) is positive.

It follows that $\frac{\partial^2 F}{\partial \alpha^2}$ is positive by combining our results of (6.22) and (6.23). Thus we find $F(\alpha)$ to be strictly convex on I_2. Furthermore, for $c \in (0, \ln^{-1}(2))$, we find

$$\begin{aligned} \lim_{\alpha \to 0.25} F(\alpha) &\leq -0.785 \ln^{-1}(2) + 1.75/(c\ln(2)) - 1.75\ln(2) + 0.25\ln(5\ln(2)) \\ &\quad - 1.5\ln((0.5\ln(2) - 0.075)/\ln(2)) - 1.18 < \frac{2(1 - c\ln(2))}{c\ln(2)}, \qquad \text{and} \\ \lim_{\alpha \to 0.85} F(\alpha) &\leq -0.92856/\ln(2) + 1.15/(c\ln(2)) - 1.15\ln(2) + 0.85\ln(5\ln(2)) \\ &\quad - 0.3\ln((0.5\ln(2) - 0.255)/\ln(2)) - 0.7191 < \frac{2(1 - c\ln(2))}{c\ln(2)}. \end{aligned}$$

Case $\alpha \in I_3$: In this case, F evaluates to

$$\begin{aligned} F(\alpha) = \alpha \ln\left(\frac{10\alpha}{16\alpha - 11}\right) + \frac{3}{5}\alpha - (2 - \alpha)\frac{c\ln(2)^2 - 1}{c\ln(2)} \\ + \frac{1}{2}\ln\left(2^{4/5\alpha - 9/5} - 2^{-6/5\alpha + 6/5} + 1\right)\ln^{-1}(2) - \frac{11}{10}. \end{aligned}$$

Then we find the following for all $\alpha \in I_3$, which is easy to verify computationally (see Figure 4):

$$\begin{aligned} \frac{\partial^2 F}{\partial \alpha^2} &= -32\,(16\,\alpha - 11)\left(\frac{1}{(16\,\alpha - 11)^2} - \frac{16\,\alpha}{(16\,\alpha - 11)^3}\right) \\ &\quad + \frac{(16\,\alpha - 11)\left(\frac{1}{16\,\alpha - 11} - \frac{16\,\alpha}{(16\,\alpha - 11)^2}\right)}{\alpha} + \frac{16}{16\,\alpha - 11} - \frac{256\,\alpha}{(16\,\alpha - 11)^2} \\ &\quad - \frac{2\left(2^{\frac{4}{5}\alpha - \frac{4}{5}} + 3 \cdot 2^{-\frac{6}{5}\alpha + \frac{6}{5}}\right)^2 \ln(2)}{25\left(2^{\frac{4}{5}\alpha - \frac{9}{5}} - 2^{-\frac{6}{5}\alpha + \frac{6}{5}} + 1\right)^2} + \frac{2\left(2^{\frac{4}{5}\alpha + \frac{1}{5}}\ln(2) - 9 \cdot 2^{-\frac{6}{5}\alpha + \frac{6}{5}}\ln(2)\right)}{25\left(2^{\frac{4}{5}\alpha - \frac{9}{5}} - 2^{-\frac{6}{5}\alpha + \frac{6}{5}} + 1\right)} > 0. \end{aligned}$$

We now check that this inequality holds. First we simplify the polynomial part to

$$\frac{176}{(16\alpha-11)^2}-\frac{11}{\alpha(16\alpha-11)}.$$

Now we lower bound the non-polynomial part

$$h(\alpha)=-\frac{2\left(2^{\frac{4}{5}\alpha-\frac{4}{5}}+3\cdot 2^{-\frac{6}{5}\alpha+\frac{6}{5}}\right)^2\ln(2)}{25\left(2^{\frac{4}{5}\alpha-\frac{9}{5}}-2^{-\frac{6}{5}\alpha+\frac{6}{5}}+1\right)^2}+\frac{2\left(2^{\frac{4}{5}\alpha+\frac{1}{5}}\ln(2)-9\cdot 2^{-\frac{6}{5}\alpha+\frac{6}{5}}\ln(2)\right)}{25\left(2^{\frac{4}{5}\alpha-\frac{9}{5}}-2^{-\frac{6}{5}\alpha+\frac{6}{5}}+1\right)}.$$

One can verify that this is negative and concave for $\alpha\in[85/100,1)$. Thus, one can derive the lower bound

$$h(\alpha)>\frac{6751}{150}\alpha-\frac{148}{3}.$$

Therefore we get a lower bound

$$\frac{\partial^2 F}{\partial\alpha^2}>\frac{176}{(16\alpha-11)^2}-\frac{11}{\alpha(16\alpha-11)}+\frac{6751}{150}\alpha-\frac{148}{3}.$$

Standard calculus reveals that the minimum is strictly positive.

Again, this means $F(\alpha)$ is convex and it suffices to check the boundary. It is easily verified that for $c\in(0,\ln^{-1}(2))$,

$$\begin{aligned}\lim_{\alpha\to 0.85}F(\alpha)&\le-((1.15\ln(2)^2-0.41687\ln(2)+0.5586)c-1.15)/(c\ln(2))\\&<\frac{2(1-c\ln(2))}{c\ln(2)},\qquad\text{and}\\\lim_{\alpha\to 1}F(\alpha)&=\frac{1-c\ln(2)}{c\ln(2)}<\frac{2(1-c\ln(2))}{c\ln(2)}.\end{aligned}$$

Finally, the lemma follows from combination of the three cases. Indeed, this proves that there is an $\varepsilon>0$ such that for all $\alpha\in(0,1]$,

$$\frac{1}{k\Delta}G(\alpha,q_{01},q_{11})=F(\alpha)<(1-\varepsilon)\frac{2(1-c\ln 2)}{c\ln 2}$$

as desired. □

Proposition 5.7 now follows, since by Lemma 6.2 and Stirling's approximation,

$$\begin{aligned}&\exp\left(G(\alpha,q_{01},q_{11})\right)\\&\quad=\exp\left(o\left(k\Delta\right)\right)\binom{N}{\alpha k,\,(1-\alpha)k,\,(1-\alpha)k}\\&\qquad\cdot\frac{\left(1-2(1-q_{01}-q_{11})^\Gamma+q_{00}^\Gamma\right)^M}{\binom{N\Delta}{\alpha k\Delta,\,(1-\alpha)k\Delta,\,(1-\alpha)k\Delta,\,(N-2k+\alpha k)\Delta}q_{11}^{\alpha k\Delta}q_{10}^{2(k-\alpha k)\Delta}q_{00}^{N\Delta-2k\Delta+\alpha k\Delta}}\\&\quad\ge\mathbb{E}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[Z(\boldsymbol{G},\alpha)]\exp\left(o\left(k\Delta\right)\right),\end{aligned}$$

and then using Proposition 6.3 concludes the proof.

6.3 Proof of Lemma 5.6

We have two adjustments to take care of in order to transfer our results from $\mathbb{Q}^{\star}_{\Delta,\Gamma}$ to $\mathbb{Q}_{\Delta}$. First, the configuration model $\mathbb{Q}^{\star}_{\Delta,\Gamma}$ may feature multi-edges, while $\mathbb{Q}_{\Delta}$ does not. Second, under $\mathbb{Q}^{\star}_{\Delta,\Gamma}$ we assume the test degrees to be regular. These two issues are handled in Sections 6.3.1 and 6.3.2, respectively.

Our proof will pass from $\mathbb{Q}^{\star}_{\Delta,\Gamma}$ to $\mathbb{Q}_{\Delta}$ by way of a third null model $\mathbb{Q}^{\star}_{\Delta}$ which is defined exactly like $\mathbb{Q}_{\Delta}$ with the sole difference that now each individual chooses Δ tests *with replacement* (i.e., multi-edges are possible).

Formally, the proof of Lemma 5.6 follows immediately by combining Lemmas 6.5, 6.6, and 6.7 below.

6.3.1 Existence of Multi-edges

In this section we show how to compare important properties of $\mathbb{Q}_{\Delta}$ and $\mathbb{Q}^{\star}_{\Delta}$. Our first result concerns $\boldsymbol{Z}(G)$.

Lemma 6.5. *We have*

$$\mathbb{E}_{\mathbb{Q}_{\Delta}}\left[\boldsymbol{Z}(G)\right] \geq \mathbb{E}_{\mathbb{Q}^{\star}_{\Delta}}\left[\boldsymbol{Z}(G)\right].$$

Proof. Given a sample $G^{\star} \sim \mathbb{Q}^{\star}_{\Delta}$, we can produce a sample $G \sim \mathbb{Q}_{\Delta}$ by resampling the duplicate edges until no multi-edges remain. This process can only increase the number of solutions: for every $\tau \in \boldsymbol{S}(G^{\star})$, we also have $\tau \in \boldsymbol{S}(G)$. □

We also have the converse bound for $\boldsymbol{Z}(G, \alpha)$.

Lemma 6.6. *For any fixed* $0 < c < \ln^{-1}(2)$, $0 < \theta < 1$, *and* $0 < \delta \leq \alpha \leq 1$,

$$\mathop{\mathbb{E}}_{\mathbb{Q}_{\Delta}}\left[\boldsymbol{Z}(G, \alpha)\right] \leq \mathop{\mathbb{E}}_{\mathbb{Q}^{\star}_{\Delta}}\left[\boldsymbol{Z}(G, \alpha)\right] \exp(o(k\Delta)).$$

Proof. Fix an arbitrary pair $\sigma, \tau \in \{0,1\}^N$ with Hamming weight k and overlap αk. Using linearity of expectation,

$$\mathop{\mathbb{E}}_{\mathbb{Q}_{\Delta}}\left[\boldsymbol{Z}(G, \alpha)\right] = \binom{N}{(1-\alpha)k, \alpha k, \alpha k} \mathbb{Q}_{\Delta}(\sigma, \tau \in \mathcal{S}(G))$$

and

$$\mathop{\mathbb{E}}_{\mathbb{Q}^{\star}_{\Delta}}\left[\boldsymbol{Z}(G, \alpha)\right] = \binom{N}{(1-\alpha)k, \alpha k, \alpha k} \mathbb{Q}^{\star}_{\Delta}(\sigma, \tau \in \mathcal{S}(G)).$$

Therefore it suffices to show

$$\mathbb{Q}_{\Delta}(\sigma, \tau \in \mathcal{S}(G)) \leq \exp(o(k\Delta)) \mathbb{Q}^{\star}_{\Delta}(\sigma, \tau \in \mathcal{S}(G)). \tag{6.24}$$

Under $\boldsymbol{G} \sim \mathbb{Q}^{\star}_{\Delta}$, let $\mathcal{E}$ denote the event that there are no multi-edges incident to individuals that have label 1 under σ or τ (or both). Notice that

$$\mathbb{Q}_{\Delta}(\sigma, \tau \in \mathcal{S}(G)) = \mathbb{Q}^{\star}_{\Delta}(\sigma, \tau \in \mathcal{S}(G) \mid \mathcal{E})$$

because the event $\{\sigma, \tau \in \mathcal{S}(G)\}$ depends only the edges incident to individuals in the union of supports $\operatorname{supp}(\sigma) \cup \operatorname{supp}(\tau)$. One can directly bound the probability $\mathbb{Q}^{\star}_{\Delta}(\mathcal{E}_M) = k^{-O(1)} = \exp(o(k\Delta))$ as in the proof of Lemma 8.8, and so we conclude (6.24). □

6.3.2 The Regularisation Process

In Section 6.3.1 we showed how to transfer results from $\mathbb{Q}^\star_\Delta$ to $\mathbb{Q}_\Delta$. In this section we show how to transfer results from $\mathbb{Q}^\star_{\Delta,\Gamma}$ to $\mathbb{Q}^\star_\Delta$. Namely, our goal is to establish the following result which (combined with Lemmas 6.5 and 6.6) completes the proof of Lemma 5.6.

Lemma 6.7. *For any fixed $\alpha \in (0,1]$,*

$$\mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[\boldsymbol{Z}(G,\alpha)] = \mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta}}[\boldsymbol{Z}(G,\alpha)]\exp\left(o(k\Delta)\right).$$

In particular,

$$\mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[\boldsymbol{Z}(G)] = \mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta}}[\boldsymbol{Z}(G)]\exp\left(o(k\Delta)\right).$$

Before proving this lemma, we introduce some notation. For $j \in [M]$, we use $\boldsymbol{\Gamma}_j$ to denote the random quantity $|\partial a_j|$, i.e., the number of individuals in test j. For technical reasons we will need to condition on the following high-probability event which states that the test degrees are well concentrated.

Lemma 6.8. *With probability $1-o(1)$ over $G \sim \mathbb{Q}^*_\Delta$,*

$$\frac{N\Delta}{M} - \ln^2(N)\sqrt{\frac{N\Delta}{M}} \le \min_j \boldsymbol{\Gamma}_j \le \max_j \boldsymbol{\Gamma}_j \le \frac{N\Delta}{M} + \ln^2(N)\sqrt{\frac{N\Delta}{M}}. \tag{6.25}$$

Since $\boldsymbol{\Gamma}_j \sim \mathrm{Bin}(N\Delta, 1/M)$, the proof is a direct consequence of Bernstein's inequality and a union bound over tests. Let $\mathcal{N}$ denote the event that (6.25) holds. We next show that conditioning on $\mathcal{N}$ does not change the expectation of $\boldsymbol{Z}(G,\alpha)$ too much.

Lemma 6.9. *We have*

$$\mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta}}[\boldsymbol{Z}(G,\alpha) \mid \mathcal{N}] = (1+o(1)) \mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta}}[\boldsymbol{Z}(G,\alpha)].$$

Proof. Define a planted model $\mathbb{P}^\star_\alpha$ as follows. To sample $G \sim \mathbb{P}^\star_\alpha$, first draw two k-sparse binary vectors $\sigma, \tau \in \{0,1\}^N$ uniformly at random subject to having overlap $\langle \sigma, \tau \rangle = \alpha k$. Then draw G from $\mathbb{Q}^\star_\Delta$ conditioned on the event that both σ and τ are solutions. Note that $\mathbb{P}^\star_\alpha(G)$ is proportional to $\boldsymbol{Z}(G,\alpha)$, that is,

$$\mathbb{P}^\star_\alpha(G) = \frac{\mathbb{Q}^\star_\Delta(G)\boldsymbol{Z}(G,\alpha)}{\mathbb{E}_{\mathbb{Q}^\star_\Delta}[\boldsymbol{Z}(G,\alpha)]}.$$

This implies the identity

$$\frac{\mathbb{E}_{\mathbb{Q}^\star_\Delta}[\boldsymbol{Z}(G,\alpha) \mid \mathcal{N}]}{\mathbb{E}_{\mathbb{Q}^\star_\Delta}[\boldsymbol{Z}(G,\alpha)]} = \frac{\mathbb{P}^\star_\alpha(\mathcal{N})}{\mathbb{Q}^\star_\Delta(\mathcal{N})}.$$

The result follows because $\mathcal{N}$ is a high-probability event under both $\mathbb{Q}^\star_\Delta$ and $\mathbb{P}^\star_\alpha$. For $\mathbb{Q}^\star_\Delta$ this is Lemma 6.8, and the claim for $\mathbb{P}^\star_\alpha$ can be proved similarly by handling the contribution from "infected" individuals similarly to the proof of Lemma 8.4. □

Proof of Lemma 6.7. The second desired claim follows from the first by setting $\alpha = 1$, so we focus on establishing the first. Furthermore, using Lemma 6.9 it suffices to prove

$$\mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[\boldsymbol{Z}(G,\alpha)] = \mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta}}[\boldsymbol{Z}(G,\alpha) \mid \mathcal{N}] \exp\left(o(k\Delta)\right).$$

Fix an arbitrary pair of k-sparse binary vectors $\sigma, \tau \in \{0,1\}^N$ with overlap $\langle \sigma, \tau \rangle = \alpha k$. By linearity of expectation,

$$\mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta,\Gamma}}[\boldsymbol{Z}(G,\alpha)] = \binom{N}{k}\binom{k}{\alpha k}\binom{N-k}{(1-\alpha)k}\mathbb{Q}^\star_{\Delta,\Gamma}\{\sigma,\tau \in \boldsymbol{S}(G)\}$$

and

$$\mathop{\mathbb{E}}_{\mathbb{Q}^\star_{\Delta}}[\boldsymbol{Z}(G,\alpha) \mid \mathcal{N}] = \binom{N}{k}\binom{k}{\alpha k}\binom{N-k}{(1-\alpha)k}\mathbb{Q}^\star_{\Delta}\{\sigma,\tau \in \boldsymbol{S}(G) \mid \mathcal{N}\}.$$

Hence it suffices to show

$$\mathbb{Q}^\star_{\Delta,\Gamma}\{\sigma,\tau \in \mathcal{S}(G)\} = \mathbb{Q}^\star_{\Delta}\{\sigma,\tau \in \boldsymbol{S}(G) \mid \mathcal{N}\} \exp\left(o(k\Delta)\right). \tag{6.26}$$

To prove (6.26) we employ the auxiliary probability space used also in the proof of Lemma 6.2. We describe again here its definition and quick motivation. We fix an *arbitrary* (to be chosen appropriately later) choice of probability values $q_{c,d} > 0$, where $c, d \in \{0,1\}$, which are solely required to sum up to 1. Now notice that to prove (6.26) we are only interested for both $\mathbb{Q}^\star_{\Delta}$ and $\mathbb{Q}^\star_{\Delta,\Gamma}$ to model the status of the edges which connect an arbitrary test with some individual labelled 1 by σ or τ. Let us first construct the probability space for $\mathbb{Q}^\star_{\Delta,\Gamma}$. In this case, the edges can be modelled as the conditional product probability measure on the binary status of the total possible $M\Gamma$ edges (counting from the test side), say $(\omega_{ij})_{i=1\ldots M, j=1\ldots\Gamma} \in \{0,1\}^{M\Gamma}$, $(\omega'_{ij})_{i=1\ldots M, j=1\ldots\Gamma} \in \{0,1\}^{M\Gamma}$, conditioned on the event $\mathcal{R}$ which makes sure to satisfy the Hamming weight k and overlap αk constraint on the individual side of σ, τ, that is we condition on

$$\mathcal{R} = \left\{ \sum_{i,j} \boldsymbol{\omega}_{ij} = \sum_{i,j} \boldsymbol{\omega}'_{ij} = k\Delta \qquad \text{and} \qquad \sum_{i,j} \boldsymbol{\omega}_{ij}\boldsymbol{\omega}'_{ij} = \alpha k\Delta. \right\}$$

The product law simply asks $(\omega_{ij})_{i=1\ldots M, j=1\ldots\Gamma}, (\omega'_{ij})_{i=1\ldots M, j=1\ldots\Gamma}$ to be independent random variables such that q_{cd} is the probability that $\omega_{ij} = c, \omega'_{ij} = d$ for $c, d \in \{0,1\}$. The symmetries of the model suffice to conclude that for any choice of $q_{c,d} > 0$ the conditional law is indeed the law also induced by $\mathbb{Q}^\star_{\Delta,\Gamma}$ on the edge status of σ, τ. One can construct in a straightforward manner the corresponding construction for $\mathbb{Q}^\star_{\Delta}$ conditional on the (varying) test degrees $\boldsymbol{\Gamma}_1, \ldots, \boldsymbol{\Gamma}_M$. We define the corresponding conditioning event as $\tilde{R}$.

Now recall that we care to compare the event of $\sigma, \tau \in \boldsymbol{S}(G)$ between the two null models. For this reason in the auxiliary spaces, we denote by $\mathcal{S}$ the event that all used edges in the auxiliary space for $\mathbb{Q}^\star_{\Delta,\Gamma}$ "cover all the M tests," and similarly define the event $\tilde{\mathcal{S}}$ "cover all the M tests" for $\mathbb{Q}^\star_{\Delta}$. Given the above it holds,

$$\mathbb{Q}^\star_{\Delta,\Gamma}\{\sigma,\tau \in \boldsymbol{S}(G)\} = \Pr(\mathcal{S} \mid \mathcal{R})$$

and

$$\mathbb{Q}^{\star}_{\Delta}\left\{\sigma,\tau\in\boldsymbol{S}(G)\mid\mathcal{N}\right\}=\mathbb{E}_{\boldsymbol{\Gamma}_i}\Pr(\tilde{\mathcal{S}}\mid\tilde{\mathcal{R}},\mathcal{N},\boldsymbol{\Gamma}_1,\ldots,\boldsymbol{\Gamma}_M)=\Pr(\tilde{\mathcal{S}}\mid\tilde{\mathcal{R}},\mathcal{N}).$$

Hence we turn our focus on proving

$$\Pr(\mathcal{S}\mid\mathcal{R})=\Pr(\tilde{\mathcal{S}}\mid\tilde{\mathcal{R}},\mathcal{N})\exp\left(o(k\Delta)\right), \tag{6.27}$$

or equivalently by Baye's rule,

$$\frac{\Pr(\mathcal{S})\Pr(\mathcal{R}\mid\mathcal{S})}{\Pr(\mathcal{R})}=\frac{\Pr(\tilde{\mathcal{S}}\mid\mathcal{N})\Pr(\tilde{\mathcal{R}}\mid\tilde{\mathcal{S}},\mathcal{N})}{\Pr(\tilde{\mathcal{R}}\mid\mathcal{N})}\exp\left(o(k\Delta)\right). \tag{6.28}$$

For the purpose of intuition, notice that (6.27) and (6.28) can be interpreted as "degree concentration" conditions in terms of the $\boldsymbol{\Gamma}_i$'s.

Recall now that so far we have defined the auxiliary probability spaces for arbitrary $q_{cd}>0$. To prove (6.28) we choose the values of the q_{cd} appropriately, similar to the proof of Lemma 6.2. We first handle the case that $0<\alpha<1$. We define q and $q_{00},\ldots,q_{11}$ such that the equations (6.1), (6.5) – (6.6) are satisfied and prove that in this case

$$q_{10},q_{01},q_{11}=\Theta\left(\frac{k}{N}\right)$$

and therefore $q_{00}=1-2q_{01}-q_{11}=1-\Theta(kN^{-1})$. Indeed, the r.h.s. of (6.1) is $\Theta\left(\frac{k}{N}\right)$, because $M=\Theta(k\Delta)$ and $\Gamma=\Theta\left(\frac{N}{k}\right)$. Because α does not depend on N, equation (6.13) implies that $q_{10},q_{01},q_{11}=\Theta\left(\frac{k}{N}\right)$.

We find that

$$\Pr(\tilde{\mathcal{S}}\mid\mathcal{N},\boldsymbol{\Gamma}_1,...,\boldsymbol{\Gamma}_M)=\prod_{i=1}^{M}\left(1-2(1-q_{01}-q_{11})^{\boldsymbol{\Gamma}_i}+q_{00}^{\boldsymbol{\Gamma}_i}\right).$$

Because by assumption $q_{01},q_{11}=\Theta\left(\frac{k}{N}\right)$, the following follows from a simple Taylor expansion of the logarithm. Recall that $\mathcal{N}$ ensures that $\boldsymbol{\Gamma}_i\sim\Theta\left(\frac{N}{k}\right)$ and, given $\mathcal{N}$,

$$\max_i\boldsymbol{\Gamma}_i\leq\min_i\boldsymbol{\Gamma}_i+O\left(\ln(N)\sqrt{\frac{N}{k}}\right).$$

Thus, given $\mathcal{N}$ we we have

$$\begin{aligned}
&\sum_{i=1}^{M}\ln\left(\frac{1-2\left(1-q_{01}-q_{11}\right)^{\boldsymbol{\Gamma}_i}+q_{00}^{\boldsymbol{\Gamma}_i}}{1-2\left(1-q_{01}-q_{11}\right)^{\Gamma}+q_{00}^{\Gamma}}\right)\\
&\quad=O\left(M\left|\max_i\boldsymbol{\Gamma}_i-\min_i\boldsymbol{\Gamma}_i\right|\left(\ln\left(1-q_{01}-q_{11}\right)\pm\ln\left(1-2q_{01}-q_{11}\right)\right)\right)\\
&\quad=\tilde{O}\left(M\sqrt{\frac{N}{k}}\cdot\frac{k}{N}\right)=o(k\Delta).
\end{aligned}$$

Therefore, we find

$$\mathbb{E}_{\mathbf{\Gamma}_i} \Pr(\tilde{\mathcal{S}} \mid \mathcal{N}, \mathbf{\Gamma}_1, \ldots, \mathbf{\Gamma}_M) = \Pr\left(\tilde{\mathcal{S}} \mid \mathcal{N}\right) = \Pr\left(\mathcal{S}\right) \exp\left(o(k\Delta)\right). \tag{6.29}$$

A similar Taylor expansion directly shows that as in Lemma 6.2

$$\begin{aligned}\Pr\left[\mathcal{R}\right] &= \binom{N\Delta}{\alpha k\Delta,\ (1-\alpha)k\Delta,\ (1-\alpha)k\Delta} q_{11}^{\alpha k\Delta} q_{10}^{2(k-\alpha k)\Delta} q_{00}^{N\Delta-2k\Delta+\alpha k\Delta} \\ &= \exp\left(o\left(k\Delta\right)\right) \Pr\left(\tilde{\mathcal{R}} \mid \mathcal{N}\right).\end{aligned}$$

We are left to prove that the conditional probabilities compare as well, more precisely that we have

$$\mathbb{E}_{\mathbf{\Gamma}_i} \Pr(\tilde{\mathcal{R}} \mid \tilde{\mathcal{S}}, \mathcal{N}, \mathbf{\Gamma}_1, \ldots, \mathbf{\Gamma}_M) = \Pr\left(\tilde{\mathcal{R}} \mid \tilde{\mathcal{S}}, \mathcal{N}\right) = \Pr\left(\mathcal{R} \mid \mathcal{S}\right) \exp\left(o\left(k\Delta\right)\right). \tag{6.30}$$

We know as in Lemma 6.2 that $\Pr\left(\mathcal{R} \mid \mathcal{S}\right) = N^{-O(1)} = \exp(o(k\Delta))$. Using an appropriate modification of the local limit theorem technique explained in Section 6 of [COHKL+21] one can similarly deduce $\Pr\left(\tilde{\mathcal{R}} \mid \tilde{\mathcal{S}}, \mathcal{N}\right) = \exp(o(k\Delta))$, completing the proof in the case $\alpha \in (0, 1)$.

The case $\alpha = 1$ follows from an almost identical line of reasoning for the case $\alpha = 1$. In this case, we have $q_{01} = q_{10} = 0$ and $q_{11} = \Theta\left(kN^{-1}\right)$ as previously. The calculation of $\Pr\left(\mathcal{S}\right) = \exp\left(o(k\Delta)\right) \Pr\left(\mathcal{S} \mid \mathcal{N}\right)$ works as above by setting $q_{01} = 0$. Indeed, given $\mathcal{N}$ it suffices to prove

$$(1 - (1 - q_{11})^{\mathbb{E}[\mathbf{\Gamma}_1]})^M = \exp\left(o\left(k\Delta\right)\right) \prod_{i=1}^{M} (1 - (1 - q_{11})^{\mathbf{\Gamma}_i}).$$

This again follows from a Taylor expansion with $\mathbb{E}\left[\mathbf{\Gamma}_1\right] \sim 2 \ln 2 \frac{N}{k}$, $q_{11} = \Theta\left(\frac{k}{N}\right)$ and $M \sim \frac{k\Delta}{2 \ln 2}$ and verifies

$$\Pr\left(\tilde{\mathcal{S}} \mid \mathcal{N}\right) = \exp\left(o\left(k\Delta\right)\right) \Pr\left(\mathcal{S}\right).$$

Analogously, as in Lemma 6.1, we can also verify that

$$\Pr\left(\mathcal{R}\right) = \binom{M\Gamma}{k\Delta} {q_{11}}^{\Delta k} (1 - q_{11})^{M\Gamma - \Delta k} = \exp\left(o\left(k\Delta\right)\right) \Pr\left(\tilde{\mathcal{R}} \mid \mathcal{N}\right).$$

and that the local central limit theorem argument carries through again to give $\Pr\left(\mathcal{R} \mid \mathcal{S}\right) = N^{-O(1)} = \exp(o(k\Delta))$ and $\Pr\left(\tilde{\mathcal{R}} \mid \tilde{\mathcal{S}}, \mathcal{N}\right) = \exp(o(k\Delta))$. □

7 Background on Hypothesis Testing and Low-Degree Polynomials

Suppose we are interested in distinguishing between two probability distributions $\mathbb{P} = \mathbb{P}_n$ and $\mathbb{Q} = \mathbb{Q}_n$ over $\mathbb{R}^p$ (in our case, $\{0, 1\}^p$), where $p = p_n$ grows with the problem size n. Given a single sample X drawn from either $\mathbb{P}$ or $\mathbb{Q}$ (each chosen with probability $1/2$), the goal is to correctly determine whether X came from $\mathbb{P}$ or $\mathbb{Q}$. There are two different objectives of interest:

- **Strong detection**: test succeeds with probability $1 - o(1)$ as $n \to \infty$.
- **Weak detection**: test succeeds with probability $\frac{1}{2} + \varepsilon$ for some constant $\varepsilon > 0$ (not depending on n).

A natural sufficient condition to obtain strong (respectively, weak) detection via a polynomial-based test is strong (resp., weak) separation, as discussed in Section 2.2. We recall the definitions here for convenience. For a multivariate polynomial $f : \mathbb{R}^p \to \mathbb{R}$,

- **Strong separation**: $\sqrt{\max\{\mathrm{Var}_{\mathbb{P}}[f], \mathrm{Var}_{\mathbb{Q}}[f]\}} = o\left(|\mathbb{E}_{\mathbb{P}}[f] - \mathbb{E}_{\mathbb{Q}}[f]|\right)$.
- **Weak separation**: $\sqrt{\max\{\mathrm{Var}_{\mathbb{P}}[f], \mathrm{Var}_{\mathbb{Q}}[f]\}} = O\left(|\mathbb{E}_{\mathbb{P}}[f] - \mathbb{E}_{\mathbb{Q}}[f]|\right)$.

7.1 Chi-Squared Divergence

The *chi-squared divergence* $\chi^2(\mathbb{P} \,\|\, \mathbb{Q})$ is a standard quantity that can be defined in a number of equivalent ways. Let $L = \frac{d\mathbb{P}}{d\mathbb{Q}}$ denote the *likelihood ratio*. Since our distributions $\mathbb{P}, \mathbb{Q}$ are on the finite set $\{0,1\}^p$, the likelihood ratio is simply $L(X) = \frac{\mathbb{P}(X)}{\mathbb{Q}(X)} := \frac{\Pr_{X' \sim \mathbb{P}}(X'=X)}{\Pr_{X' \sim \mathbb{Q}}(X'=X)}$. To ensure that L is defined, we will always assume $\mathbb{P}$ is absolutely continuous with respect to $\mathbb{Q}$, which on the finite domain $\{0,1\}^p$ simply means the support of $\mathbb{P}$ is contained in the support of $\mathbb{Q}$ (we can define $L(X) = 0$ outside the support of $\mathbb{Q}$). We have

$$\begin{aligned}
\chi^2(\mathbb{P} \,\|\, \mathbb{Q}) &:= \mathop{\mathbb{E}}_{X \sim \mathbb{Q}} L(X)^2 - 1 \\
&= \sup_{f : \mathbb{R}^p \to \mathbb{R}} \frac{\left(\mathbb{E}_{X \sim \mathbb{P}} f(X)\right)^2}{\mathbb{E}_{X \sim \mathbb{Q}} f(X)^2} - 1 \\
&= \sup_{\substack{f : \mathbb{R}^p \to \mathbb{R} \\ \mathbb{E}_{X \sim \mathbb{Q}} f(X) = 0}} \frac{\left(\mathbb{E}_{X \sim \mathbb{P}} f(X)\right)^2}{\mathbb{E}_{X \sim \mathbb{Q}} f(X)^2}.
\end{aligned}$$

The equivalence between these definitions is standard, and follows as a special case of Lemma 7.2 below. Standard arguments use the chi-squared divergence to show information-theoretic impossibility of detection (see for example Lemma 2 of [MRZ15]):

Lemma 7.1.

- *If* $\chi^2(\mathbb{P} \,\|\, \mathbb{Q}) = O(1)$ *as* $n \to \infty$ *then strong detection is impossible.*
- *If* $\chi^2(\mathbb{P} \,\|\, \mathbb{Q}) = o(1)$ *as* $n \to \infty$ *then weak detection is impossible.*

One can use either $\chi^2(\mathbb{P} \,\|\, \mathbb{Q})$ or $\chi^2(\mathbb{Q} \,\|\, \mathbb{P})$ for this purpose, but it is typically more tractable to bound $\chi^2(\mathbb{P} \,\|\, \mathbb{Q})$ where $\mathbb{Q}$ is the "simpler" distribution.

7.2 Low-Degree Chi-Squared Divergence

The *degree-D chi-squared divergence* $\chi^2_{\le D}(\mathbb{P} \,\|\, \mathbb{Q})$ is an analogous quantity which measures whether or not $\mathbb{P}, \mathbb{Q}$ can be distinguished by a degree-D polynomial. Let $\mathbb{R}[X]_{\le D}$ denote the space of multivariate polynomials $\mathbb{R}^p \to \mathbb{R}$ of degree (at most) D. For functions $\mathbb{R}^p \to \mathbb{R}$, define the inner product $\langle f, g \rangle_{\mathbb{Q}} := \mathbb{E}_{X \sim \mathbb{Q}}[f(X)g(X)]$ and the associated norm $\|f\|_{\mathbb{Q}} = \sqrt{\langle f, f \rangle_{\mathbb{Q}}}$. Also let $f^{\le D}$ denote the orthogonal (with respect to $\langle \cdot, \cdot \rangle_{\mathbb{Q}}$) projection of f onto $\mathbb{R}[X]_{\le D}$. Recall that $L = \frac{d\mathbb{P}}{d\mathbb{Q}}$ denotes the likelihood ratio. We have the equivalent definitions

$$\chi^2_{\le D}(\mathbb{P} \,\|\, \mathbb{Q}) := \mathop{\mathbb{E}}_{X \sim \mathbb{Q}} L^{\le D}(X)^2 - 1 = \|L^{\le D}\|^2_{\mathbb{Q}} - 1 \tag{7.1}$$

$$= \sup_{f \in \mathbb{R}[X]_{\le D}} \frac{(\mathbb{E}_{X \sim \mathbb{P}} f(X))^2}{\mathbb{E}_{X \sim \mathbb{Q}} f(X)^2} - 1 \tag{7.2}$$

$$= \sup_{\substack{f \in \mathbb{R}[X]_{\le D} \\ \mathbb{E}_{X \sim \mathbb{Q}} f(X) = 0}} \frac{(\mathbb{E}_{X \sim \mathbb{P}} f(X))^2}{\mathbb{E}_{X \sim \mathbb{Q}} f(X)^2}. \tag{7.3}$$

These equivalences are standard (see e.g. [Hop18, KWB19]), and we include the proof for convenience.

Lemma 7.2. *Suppose $\mathbb{P}$ and $\mathbb{Q}$ are distributions over $\mathbb{R}^p$ with $\mathbb{P}$ absolutely continuous with respect to $\mathbb{Q}$. The three definitions for $\chi^2_{\le D}(\mathbb{P} \,\|\, \mathbb{Q})$ in (7.1)-(7.3) are equivalent.*

Proof. For (7.1)=(7.2),

$$\sup_{f \in \mathbb{R}[X]_{\le D}} \frac{(\mathbb{E}_{X \sim \mathbb{P}} f(X))^2}{\mathbb{E}_{X \sim \mathbb{Q}} f(X)^2} = \sup_{f \in \mathbb{R}[X]_{\le D}} \frac{(\mathbb{E}_{X \sim \mathbb{Q}} f(X)L(X))^2}{\mathbb{E}_{X \sim \mathbb{Q}} f(X)^2} = \sup_{f \in \mathbb{R}[X]_{\le D}} \frac{\langle f, L \rangle^2_{\mathbb{Q}}}{\|f\|^2_{\mathbb{Q}}}$$

which is optimized by $f = L^{\le D}$, so

$$= \frac{\langle L^{\le D}, L \rangle^2_{\mathbb{Q}}}{\|L^{\le D}\|^2_{\mathbb{Q}}} = \frac{\|L^{\le D}\|^4_{\mathbb{Q}}}{\|L^{\le D}\|^2_{\mathbb{Q}}} = \|L^{\le D}\|^2_{\mathbb{Q}}.$$

For (7.1)=(7.3), define the subspace $V = \{f \in \mathbb{R}[X]_{\le D} : \mathbb{E}_{X \sim \mathbb{Q}}[f] = 0\} = \{f \in \mathbb{R}[X]_{\le D} : \langle f, 1 \rangle_{\mathbb{Q}} = 0\}$ and let f^V denote orthogonal projection of f onto this subspace. Similarly to above,

$$\sup_{f \in V} \frac{(\mathbb{E}_{X \sim \mathbb{P}} f(X))^2}{\mathbb{E}_{X \sim \mathbb{Q}} f(X)^2} = \|L^V\|^2_{\mathbb{Q}}.$$

Now $L^V = (L - \langle L, 1 \rangle_{\mathbb{Q}})^{\le D} = (L - 1)^{\le D} = L^{\le D} - 1$ and so

$$\begin{aligned} \|L^V\|^2_{\mathbb{Q}} &= \|L^{\le D} - 1\|^2_{\mathbb{Q}} = \|L^{\le D}\|^2_{\mathbb{Q}} - 2\langle L^{\le D}, 1 \rangle_{\mathbb{Q}} + 1 \\ &= \|L^{\le D}\|^2_{\mathbb{Q}} - 2\langle L, 1 \rangle_{\mathbb{Q}} + 1 = \|L^{\le D}\|^2_{\mathbb{Q}} - 1, \end{aligned}$$

completing the proof. □

Note that on the finite domain $\{0,1\}^p$, the degree-D chi-squared divergence recovers the usual chi-squared divergence whenever $D \geq p$, since any function $\{0,1\}^p \to \mathbb{R}$ can be written as a degree-p polynomial. From (7.1) we can see that the quantity $\sqrt{\chi^2_{\leq D}(\mathbb{P} \,\|\, \mathbb{Q}) + 1}$ is equal to $\|L^{\leq D}\|_{\mathbb{Q}}$, which is commonly called the *norm of the low-degree likelihood ratio* (see [Hop18, KWB19]). Analogous to the standard chi-squared divergence, we have the following interpretation for $\chi^2_{\leq D}(\mathbb{P} \,\|\, \mathbb{Q})$.

- If $\chi^2_{\leq D}(\mathbb{P} \,\|\, \mathbb{Q}) = O(1)$ for some $D = \omega(\ln p)$, this suggests that strong detection has no polynomial-time algorithm and furthermore requires runtime $\exp(\tilde{\Omega}(D))$.
- If $\chi^2_{\leq D}(\mathbb{P} \,\|\, \mathbb{Q}) = o(1)$ for some $D = \omega(\ln p)$, this suggests that weak detection has no polynomial-time algorithm and furthermore requires runtime $\exp(\tilde{\Omega}(D))$.

To justify the above interpretations, recall the notions of strong/weak separation and low-degree hardness from Section 2.2. We will see (Lemma 7.3) that if $\chi^2_{\leq D}(\mathbb{P} \,\|\, \mathbb{Q}) = O(1)$ then no degree-D polynomial can strongly separate $\mathbb{P}$ and $\mathbb{Q}$, and similarly, if $\chi^2_{\leq D}(\mathbb{P} \,\|\, \mathbb{Q}) = o(1)$ then no degree-D polynomial can weakly separate $\mathbb{P}$ and $\mathbb{Q}$. For further discussion on some other sense(s) in which $\chi^2_{\leq D}(\mathbb{P} \,\|\, \mathbb{Q})$ can be used to rule out polynomial-based tests, we refer the reader to [KWB19], Section 4.1 (for strong detection) and [LWB20], Section 2.3 (for weak detection).

7.3 Conditional Chi-Squared Divergence

It is well known that in some instances, the chi-squared divergence is not sufficient to prove sharp impossibility results: there are cases where detection is impossible, yet $\chi^2(\mathbb{P} \,\|\, \mathbb{Q}) \to \infty$ due to a rare "bad" event under $\mathbb{P}$. Sharper results can sometimes be obtained by a conditional chi-squared calculation. This amounts to defining a modified planted distribution $\tilde{\mathbb{P}}$ by conditioning $\mathbb{P}$ on some high-probability event (that is, an event of probability $1 - o(1)$). Note that any algorithm for strong (respectively, weak) detection between $\mathbb{P}$ and $\mathbb{Q}$ also achieves strong (respectively, weak) detection between $\tilde{\mathbb{P}}$ and $\mathbb{Q}$. As a result, bounds on $\chi^2(\tilde{\mathbb{P}} \,\|\, \mathbb{Q})$ can be used to prove impossibility of detection between $\mathbb{P}$ and $\mathbb{Q}$. This technique is classical, and it turns out to have a low-degree analogue: bounds on $\chi^2_{\leq D}(\tilde{\mathbb{P}} \,\|\, \mathbb{Q})$ can be used to show failure of low-degree polynomials to strongly/weakly separate $\mathbb{P}$ and $\mathbb{Q}$, as we see below. (This result also appears in [BEH+22, Proposition 6.2] and we include the proof here for convenience.)

Lemma 7.3. *Suppose $\mathbb{P} = \mathbb{P}_n$ and $\mathbb{Q} = \mathbb{Q}_n$ are distributions over $\mathbb{R}^p$ for some $p = p_n$. Let $A = A_n$ be a high-probability event under $\mathbb{P}$, that is, $\mathbb{P}(A) = 1 - o(1)$. Define the conditional distribution $\tilde{\mathbb{P}} = \mathbb{P} \,|\, A$.*

- *If $\chi^2_{\leq D}(\tilde{\mathbb{P}} \,\|\, \mathbb{Q}) = O(1)$ as $n \to \infty$ for some $D = D_n$, then no degree-D polynomial strongly separates $\mathbb{P}$ and $\mathbb{Q}$ in the sense of (2.1).*
- *If $\chi^2_{\leq D}(\tilde{\mathbb{P}} \,\|\, \mathbb{Q}) = o(1)$ as $n \to \infty$ for some $D = D_n$, then no degree-D polynomial weakly separates $\mathbb{P}$ and $\mathbb{Q}$ in the sense of (2.2).*

Proof. We prove the contrapositive. Suppose $f = f_n$ strongly (respectively, weakly) separates $\mathbb{P}$ and $\mathbb{Q}$. By shifting and rescaling we can assume without loss of generality that $\mathbb{E}_{\mathbb{Q}}[f] = 0$ and

$\mathbb{E}_{\mathbb{P}}[f] = 1$, and that $\mathrm{Var}_{\mathbb{Q}}[f], \mathrm{Var}_{\mathbb{P}}[f]$ are both $o(1)$ (resp., $O(1)$). Note that $\mathbb{E}_{\mathbb{Q}}[f^2] = \mathrm{Var}_{\mathbb{Q}}[f]$. It suffices to show $\mathbb{E}_{\tilde{\mathbb{P}}}[f] \geq 1 - o(1)$ so that, using (7.3),

$$\chi^2_{\leq D}(\tilde{\mathbb{P}} \,\|\, \mathbb{Q}) \geq \frac{(\mathbb{E}_{\tilde{\mathbb{P}}}[f])^2}{\mathbb{E}_{\mathbb{Q}}[f^2]} \geq \frac{1 - o(1)}{\mathrm{Var}_{\mathbb{Q}}[f]}$$

which is $\omega(1)$ (resp., $\Omega(1)$), completing the proof.

It remains to prove $\mathbb{E}_{\tilde{\mathbb{P}}}[f] \geq 1 - o(1)$. Letting A^c denote the complement of the event A, we have

$$1 = \mathop{\mathbb{E}}_{\mathbb{P}}[f] = \mathbb{P}(A) \mathop{\mathbb{E}}_{\tilde{\mathbb{P}}}[f] + \mathbb{P}(A^c) \mathop{\mathbb{E}}_{\mathbb{P}}[f \mid A^c],$$

and so, solving for $\mathbb{E}_{\tilde{\mathbb{P}}}[f]$,

$$\mathop{\mathbb{E}}_{\tilde{\mathbb{P}}}[f] = \mathbb{P}(A)^{-1}(1 - \mathbb{P}(A^c) \mathop{\mathbb{E}}_{\mathbb{P}}[f \mid A^c]).$$

Since $\mathbb{P}(A) = 1 - o(1)$, it suffices to show $|\mathbb{P}(A^c)\, \mathbb{E}_{\mathbb{P}}[f \mid A^c]| = o(1)$. We can also repeat the above argument for the second moment:

$$\mathop{\mathbb{E}}_{\mathbb{P}}[f^2] = \mathbb{P}(A) \mathop{\mathbb{E}}_{\tilde{\mathbb{P}}}[f^2] + \mathbb{P}(A^c) \mathop{\mathbb{E}}_{\mathbb{P}}[f^2 \mid A^c],$$

and so

$$\mathbb{P}(A^c) \mathop{\mathbb{E}}_{\mathbb{P}}[f^2 \mid A^c] \leq \mathop{\mathbb{E}}_{\mathbb{P}}[f^2] = \mathop{\mathrm{Var}}_{\mathbb{P}}[f] + 1.$$

We can use the above to conclude

$$\begin{aligned}
\left|\mathbb{P}(A^c) \mathop{\mathbb{E}}_{\mathbb{P}}[f \mid A^c]\right| &\leq \mathbb{P}(A^c) \sqrt{\mathop{\mathbb{E}}_{\mathbb{P}}[f^2 \mid A^c]} \\
&\leq \mathbb{P}(A^c) \sqrt{\mathbb{P}(A^c)^{-1} (\mathop{\mathrm{Var}}_{\mathbb{P}}[f] + 1)} \\
&= \sqrt{\mathbb{P}(A^c)} \cdot \sqrt{\mathop{\mathrm{Var}}_{\mathbb{P}}[f] + 1} \\
&= o(1) \cdot O(1) = o(1),
\end{aligned}$$

completing the proof. □

7.4 Proof Technique for Low-Degree Lower Bounds: Low-Overlap Second Moment

We now give an overview of the proof strategy for our low-degree hardness results. We will bound the low-degree chi-squared divergence using a "low-overlap chi-squared calculation." (This is not to be confused with the *conditional* chi-squared from the previous section, although we will sometimes use both together—a "low-overlap conditional chi-squared calculation." But for now, suppose we are simply working with $\mathbb{P}$ instead of $\tilde{\mathbb{P}}$.) This strategy was employed implicitly by [BBK+21, BKW20, KWB19] and is investigated in more detail by [BEH+22].

Recall that for the group testing models we consider, the planted distribution $\mathbb{P}$ takes the following form: first a set of k infected individuals is chosen uniformly at random, which we encode using a k-sparse indicator vector $u \in \{0,1\}^N$; then the observation X is drawn from an appropriate

distribution $\mathbb{P}_u$. We can therefore write $L(X) = \mathbb{E}_{u\sim\mathcal{U}} L_u(X)$ with $L_u = d\mathbb{P}_u/d\mathbb{Q}$, where $\mathcal{U}$ denotes the uniform measure on k-sparse binary vectors. This means, using linearity of the degree-D projection operator,

$$\begin{aligned}\chi^2_{\le D}(\mathbb{P}\,\|\,\mathbb{Q}) + 1 = \left\|L^{\le D}\right\|^2_{\mathbb{Q}} &= \left\|\left(\mathop{\mathbb{E}}_{u\sim\mathcal{U}} L_u\right)^{\le D}\right\|^2_{\mathbb{Q}} = \left\|\mathop{\mathbb{E}}_{u\sim\mathcal{U}}\left(L_u^{\le D}\right)\right\|^2_{\mathbb{Q}} \\ &= \left\langle \mathop{\mathbb{E}}_{u\sim\mathcal{U}} L_u^{\le D}, \mathop{\mathbb{E}}_{u'\sim\mathcal{U}} L_{u'}^{\le D}\right\rangle_{\mathbb{Q}} = \mathop{\mathbb{E}}_{u,u'\sim\mathcal{U}} \langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}}\end{aligned}$$

where u and u' are drawn independently from $\mathcal{U}$. For some threshold $\delta > 0$ to be chosen later (which may scale with n), we will break this expression down into two parts and handle them separately:

$$\chi^2_{\le D}(\mathbb{P}\,\|\,\mathbb{Q}) + 1 = \mathcal{R}_{\le\delta} + \mathcal{R}_{>\delta}$$

where

$$\mathcal{R}_{\le\delta} := \mathop{\mathbb{E}}_{u,u'\sim\mathcal{U}} \mathbb{1}_{\langle u,u'\rangle\le\delta} \langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}}$$

and

$$\mathcal{R}_{>\delta} := \mathop{\mathbb{E}}_{u,u'\sim\mathcal{U}} \mathbb{1}_{\langle u,u'\rangle>\delta} \langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}}.$$

We now sketch the arguments for bounding these two terms. We will show $\mathcal{R}_{>\delta} = o(1)$ by leveraging the fact that $\langle u, u'\rangle > \delta$ is a very low-probability event, combined with a crude upper bound on $\langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}}$. For $\mathcal{R}_{\le\delta}$, we will first use a symmetry argument from [BEH+22, Proposition 3.6] (we include the details in Lemmas 8.12 and 9.6) to show $\langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}} \le \langle L_u, L_{u'}\rangle_{\mathbb{Q}}$ for all u, u', and so

$$\mathcal{R}_{\le\delta} \le \mathcal{T}_{\le\delta} := \mathop{\mathbb{E}}_{u,u'\sim\mathcal{U}} \mathbb{1}_{\langle u,u'\rangle\le\delta} \langle L_u, L_{u'}\rangle_{\mathbb{Q}}.$$

Thus it suffices to bound the "low-overlap second moment" $\mathcal{T}_{\le\delta}$. Since this quantity does not involve low-degree projection, it will be tractable to compute directly.

We will sometimes need to bound the *conditional* low-degree chi-squared divergence, in which case we follow the above proof sketch with a modified planted distribution $\tilde{\mathbb{P}}$ in place of $\mathbb{P}$.

We remark that the "standard" approach to bounding the low-degree chi-squared divergence involves direct moment computations with a basis of $\mathbb{Q}$-orthogonal polynomials (see e.g. [Hop18], Section 2.3 or [KWB19], Section 2.3). For the group testing models we consider here, this approach seems prohibitively complicated: for the Bernoulli design we will need a modified planted distribution $\tilde{\mathbb{P}}$, under which it seems difficult to directly compute expectations of orthogonal polynomials; for the constant-column design, the orthogonal polynomials themselves are quite complicated and arduous to work with directly. By following the more indirect proof sketch outlined above, we are able to drastically simplify these calculations: for the Bernoulli design, the low-overlap second moment $\mathcal{T}_{\le\delta}$ "plays well" with the conditional distribution $\tilde{\mathbb{P}}$; for the constant-column design, we manage to largely avoid working with the specific details of the orthogonal polynomials (aside from some very basic properties used when bounding $\mathcal{R}_{>\delta}$).

8 Detection in the Constant-Column Design

8.1 Detection Algorithm: Proof of Theorem 3.2(a)

Recall that our goal is to derive conditions under which there exists a low-degree algorithm that achieves strong separation (as defined in (2.1)) for the following two distributions:

- Null model $\mathbb{Q}$: N individuals each participate in exactly Δ distinct tests, chosen uniformly at random (from a total number of M tests).
- Planted Model $\mathbb{P}$: a set of k infected individuals out of N is chosen uniformly at random. Then a graph is drawn as in the null model conditioned on having at least one infected individual in every test.

Proposition 8.1. *Fix an arbitrary constant $\varepsilon > 0$. If $k^3 \geq N^{2+\varepsilon}$ then there is a degree-2 polynomial that strongly separates $\mathbb{P}$ and $\mathbb{Q}$.*

This implies Theorem 3.2(a) because the condition $c > c_{\mathrm{LD}}^{\mathrm{CC}}$ is equivalent to $k^3 \geq N^{2+\varepsilon}$. The polynomial achieving strong separation is T defined in (8.1). The value of T is computable in polynomial time, so by Chebyshev's inequality, this also gives a polynomial-time algorithm for strong detection by thresholding T.

The rest of this section is devoted to proving Proposition 8.1. Given an (N, M)-bipartite graph $X \in \{0,1\}^{NM}$ drawn from either $\mathbb{P}$ or $\mathbb{Q}$, let $\Gamma_1, \dots, \Gamma_M$ denote the degree sequence of the tests, i.e., Γ_j is the number of individuals in test j. The polynomial we use to distinguish will be $T : \{0,1\}^{NM} \to \mathbb{R}$ defined by

$$T(X) = \sum_{j=1}^{M} \left(\Gamma_j - \frac{N\Delta}{M} \right)^2 . \tag{8.1}$$

Note that each Γ_j is a degree-1 polynomial in X, and so T is a degree-2 polynomial in X.

Remark 8.2. *Since the total number of edges in the graph is exactly $N\Delta = \sum_j \Gamma_j$, we can expand the square in (8.1) to deduce*

$$T(X) = \sum_{j=1}^{M} \Gamma_j^2 - \frac{N^2\Delta^2}{M},$$

which means the simpler polynomial $\sum_j \Gamma_j^2$ also achieves strong separation in the same regime that T does. However, the centered version (8.1) will be more convenient for our analysis.

In the planted model, decompose $\Gamma_j = Z_j + W_j$ where W_j is the contribution from infected edges and Z_j is the contribution from non-infected edges. There are two key claims we need to prove:

Lemma 8.3. *In the null model, $|T - \mathbb{E}[T]| \leq \tilde{O}(N/\sqrt{k})$ with overwhelming probability $1 - n^{-\omega(1)}$.*

Lemma 8.4. *In the planted model,*

$$\left| \left(\sum_j W_j^2 \right) - (1 + \ln 2 + o(1))k\Delta \right| \leq \tilde{O}(\sqrt{k})$$

with overwhemling probability $1 - n^{-\omega(1)}$.

8.1.1 Proof of Proposition 8.1

We first show how to complete the proof of Proposition 8.1 assuming Lemmas 8.3 and 8.4.

Lemma 8.5.

$$\operatorname*{Var}_{\mathbb{Q}}[T] = \tilde{O}(N^2/k).$$

Proof. Since $T \le n^{O(1)}$ almost surely, this is immediate from Lemma 8.3. □

Lemma 8.6.

$$\left|\operatorname*{\mathbb{E}}_{\mathbb{P}}[T] - \operatorname*{\mathbb{E}}_{\mathbb{Q}}[T]\right| = \tilde{\Omega}(k).$$

Proof. Under $\mathbb{Q}$ we have $\Gamma_j \sim \mathrm{Bin}(N, \frac{\Delta}{M})$ for each j (but these are not independent), so we can compute

$$\operatorname*{\mathbb{E}}_{\mathbb{Q}}[T] = M \cdot \mathrm{Var}\left[\mathrm{Bin}\left(N, \frac{\Delta}{M}\right)\right] = N\Delta\left(1 - \frac{\Delta}{M}\right). \tag{8.2}$$

Under $\mathbb{P}$, let $\overline{Z}_j = Z_j - (N-k)\frac{\Delta}{M}$ and $\overline{W}_j = W_j - k\frac{\Delta}{M}$, and write

$$T = \sum_j (\overline{Z}_j + \overline{W}_j)^2 = \sum_j \overline{Z}_j^2 + \sum_j \overline{W}_j^2 + 2\sum_j \overline{Z}_j \overline{W}_j. \tag{8.3}$$

Similarly to (8.2),

$$\mathbb{E}\left[\sum_j \overline{Z}_i^2\right] = (N-k)\Delta\left(1 - \frac{\Delta}{M}\right). \tag{8.4}$$

Also, $\mathbb{E}[\overline{Z}_j \overline{W}_j] = 0$ due to the independence between the Z's and W's along with the centering $\mathbb{E}[\overline{Z}_j] = \mathbb{E}[\overline{W}_j] = 0$. The centering for W follows because the total number of infected edges is exactly $k\Delta = \sum_j W_j$. Finally, using this same fact again,

$$\sum_j \overline{W}_j^2 = \sum_j \left(W_j^2 - 2k\frac{\Delta}{M}W_j + k^2\frac{\Delta^2}{M^2}\right) = \sum_j W_j^2 - \frac{k^2\Delta^2}{M}.$$

Combining the above, we conclude

$$\operatorname*{\mathbb{E}}_{\mathbb{P}}[T] - \operatorname*{\mathbb{E}}_{\mathbb{Q}}[T] = \mathbb{E}\left[\sum_j W_j^2\right] - k\Delta - k(k-1)\frac{\Delta^2}{M} = \mathbb{E}\left[\sum_j W_j^2\right] - (1 + 2\ln 2 + o(1))k\Delta.$$

Finally, since $\sum_j W_j^2 \le n^{O(1)}$ almost surely, Lemma 8.4 implies

$$\mathbb{E}\left[\sum_j W_j^2\right] = (1 + \ln 2 + o(1))k\Delta \pm \tilde{O}(\sqrt{k}), \tag{8.5}$$

and so

$$\operatorname*{\mathbb{E}}_{\mathbb{P}}[T] - \operatorname*{\mathbb{E}}_{\mathbb{Q}}[T] = -(\ln 2 + o(1))k\Delta \pm \tilde{O}(\sqrt{k}) = -\tilde{\Theta}(k),$$

completing the proof. □

Lemma 8.7.

$$\operatorname*{Var}_{\mathbb{P}}[T] = \tilde{O}(N^2/k).$$

Proof. Recall from (8.3) the decomposition

$$T = \sum_j \overline{Z}_j^2 + \sum_j \overline{W}_j^2 + 2\sum_j \overline{Z}_j \overline{W}_j.$$

We claim that all pairwise covariances between the three terms in the right-hand side above are zero. For the first two terms,

$$\operatorname{Cov}\left(\sum_j \overline{Z}_j^2\,,\ \sum_j \overline{W}_j^2\right) = 0$$

follows immediately because the Z's are independent from the W's. We can also compute

$$\begin{aligned}
\operatorname{Cov}\left(\sum_j \overline{Z}_j^2\,,\ \sum_j \overline{Z}_j \overline{W}_j\right) &= \sum_{ij} \mathbb{E}[\overline{Z}_i^2 \overline{Z}_j \overline{W}_j] - \mathbb{E}\left[\sum_j \overline{Z}_j^2\right] \mathbb{E}\left[\sum_j \overline{Z}_j \overline{W}_j\right] \\
&= \sum_{ij} \mathbb{E}[\overline{Z}_i^2 \overline{Z}_j]\,\mathbb{E}[\overline{W}_j] - \mathbb{E}\left[\sum_j \overline{Z}_j^2\right] \left(\sum_j \mathbb{E}[\overline{Z}_j]\,\mathbb{E}[\overline{W}_j]\right) \\
&= 0,
\end{aligned}$$

where we have used independence between the Z's and W's along with the centering $\mathbb{E}[\overline{Z}_j] = \mathbb{E}[\overline{W}_j] = 0$. The third covariance can similarly be computed to be zero. As a result,

$$\operatorname*{Var}_{\mathbb{P}}[T] = \operatorname{Var}\left[\sum_j \overline{Z}_j^2\right] + \operatorname{Var}\left[\sum_j \overline{W}_j^2\right] + \operatorname{Var}\left[\sum_j \overline{Z}_j \overline{W}_j\right].$$

The first two terms are $\tilde{O}(N^2/k)$ and $\tilde{O}(k)$ respectively, using Lemmas 8.3 and 8.4 respectively. We will compute the third term. Since $\sum_i \overline{Z}_i = 0$ almost surely, we have, using symmetry,

$$0 = \mathbb{E}\left[\left(\sum_j \overline{Z}_j\right)^2\right] = M\,\mathbb{E}[\overline{Z}_1^2] + M(M-1)\,\mathbb{E}[\overline{Z}_1 \overline{Z}_2].$$

Therefore $\mathbb{E}[\overline{Z}_1 \overline{Z}_2] = -\frac{1}{M-1}\,\mathbb{E}[\overline{Z}_1^2]$ and similarly, $\mathbb{E}[\overline{W}_1 \overline{W}_2] = -\frac{1}{M-1}\,\mathbb{E}[\overline{W}_1^2]$. We can use this to

compute

$$\begin{aligned}
\operatorname{Var}\left[\sum_j \overline{Z}_j \overline{W}_j\right] &= \sum_{ij} \mathbb{E}[\overline{Z}_i \overline{Z}_j \overline{W}_i \overline{W}_j] \\
&= \sum_{ij} \mathbb{E}[\overline{Z}_i \overline{Z}_j]\, \mathbb{E}[\overline{W}_i \overline{W}_j] \\
&= \sum_{i} \mathbb{E}[\overline{Z}_i^2]\, \mathbb{E}[\overline{W}_i^2] + \sum_{i \neq j} \mathbb{E}[\overline{Z}_i \overline{Z}_j]\, \mathbb{E}[\overline{W}_i \overline{W}_j] \\
&= M\, \mathbb{E}[\overline{Z}_1^2]\, \mathbb{E}[\overline{W}_1^2] + M(M-1) \cdot \frac{-1}{M-1}\, \mathbb{E}[\overline{Z}_1^2] \cdot \frac{-1}{M-1}\, \mathbb{E}[\overline{W}_1^2] \\
&= \frac{M^2}{M-1}\, \mathbb{E}[\overline{Z}_1^2]\, \mathbb{E}[\overline{W}_1^2] \\
&= \frac{1}{M-1}\, \mathbb{E}\left[\sum_j \overline{Z}_j^2\right] \mathbb{E}\left[\sum_j \overline{W}_j^2\right] = \tilde{O}\left(\frac{1}{k} \cdot N \cdot k\right) = \tilde{O}(N),
\end{aligned}$$

where we have used (8.4) and (8.5) in the final line. Since $k \leq N \leq N^2/k$, we conclude $\operatorname{Var}_{\mathbb{P}}[T] = \tilde{O}(N^2/k + k + N) = \tilde{O}(N^2/k)$. □

Proof of Proposition 8.1. This follows immediately from the definition of strong separation (2.1) by combining Lemmas 8.5, 8.6, and 8.7. □

8.1.2 Proof of Lemma 8.3

Proof of Lemma 8.3. Under $\mathbb{Q}$ we have $\Gamma_j \sim \operatorname{Bin}(N, \frac{\Delta}{M})$ for each j (although these are not independent), which has mean $\frac{N\Delta}{M} \geq n^{\Omega(1)}$ and variance $\leq \frac{N\Delta}{M}$. Bernstein's inequality gives $|\Gamma_j - \frac{N\Delta}{M}| \leq \sqrt{\frac{N\Delta}{M}} \ln n$ with probability $n^{-\omega(1)}$. Let $\Gamma_\pm := \frac{N\Delta}{M} \pm \sqrt{\frac{N\Delta}{M}} \ln n$. Define Γ'_j to be the restriction of Γ_j to the interval $[\Gamma_-, \Gamma_+]$, that is,

$$\Gamma'_j := \begin{cases} \Gamma_- & \text{if } \Gamma_j < \Gamma_- \\ \Gamma_j & \text{if } \Gamma_- \leq \Gamma_j \leq \Gamma_+ \\ \Gamma_+ & \text{if } \Gamma_j > \Gamma_+ \end{cases}$$

and let

$$T' := \sum_{j=1}^{M} \left(\Gamma'_j - \frac{N\Delta}{M}\right)^2.$$

The Bernstein bound above implies $T' = T$ with probability $1 - n^{-\omega(1)}$ and (since $T, T' \leq n^{O(1)}$) $\mathbb{E}[T'] = \mathbb{E}[T] \pm n^{-\omega(1)}$. It therefore suffices to prove the lemma with T' in place of T.

We will apply McDiarmid's inequality to T'. Let $X_i \subseteq [M]$ denote individual i's choice of Δ distinct tests. Note that $\{X_i\}$ are independent and that T' is a deterministic function of $\{X_i\}$; we write $T' = T'(X_1, \ldots, X_N)$. To apply McDiarmid's inequality, we need to bound the maximum possible change in T' induced by changing a single X_i. If a single X_i changes, this changes at

most $2\Delta = \tilde{O}(1)$ different Γ'_j values, each of which changes by at most 1. When Γ'_j changes to $\Gamma'_j + \delta$ for $\delta \in \{\pm 1\}$, the induced change in T' is

$$\left|\left(\Gamma'_j + \delta - \frac{N\Delta}{M}\right)^2 - \left(\Gamma'_j - \frac{N\Delta}{M}\right)^2\right| = \left|2\delta\left(\Gamma'_j - \frac{N\Delta}{M}\right) + 1\right| \leq 2\sqrt{\frac{N\Delta}{M}}\ln n + 1 = \tilde{O}(\sqrt{N/k}).$$

McDiarmid's inequality now yields

$$|T' - \mathbb{E}[T']| \leq \tilde{O}(N/\sqrt{k}) \qquad \text{with probability } 1 - n^{-\omega(1)},$$

completing the proof. □

8.1.3 Proof of Lemma 8.4

Proof of Lemma 8.4. We first give an overview of the proof, which involves a series of comparisons to simpler models. Since the infected and non-infected individuals behave independently, we only need to consider the infected individuals in this proof. We will define quantities R_j that are similar to W_j except with multi-edges allowed. The R_j's can be generated by a balls-into-bins experiment conditioned on having at least one ball (infected edge) in each bin (test). We then approximate the load per bin as a family of independent random variables R'_j with distribution $\mathrm{Poi}_{\geq 1}(\lambda)$ (Poisson conditioned on value at least 1), for a certain choice of λ. Standard concentration arguments imply the desired result for the R'_j's with overwhelming probability $1 - n^{-\omega(1)}$. We next show that with non-trivial probability $n^{-O(1)}$, the sum of the R'_j's is exactly $k\Delta$, in which case the R'_j's have the same joint distribution as the R_j's. This lets us conclude the desired result for the R_j's with overwhelming probability. Finally, we show that with non-trivial probability $n^{-O(1)}$, the balls-into-bins experiment did not feature any multi-edges, allowing us to conclude the desired result for the original W_j's. In the following, we will fill in this sketch with details.

Suppose $k\Delta$ balls are thrown into M bins independently and uniformly at random, conditioned on having at least one ball in every bin. Let R_j denote the random number of balls in bin j. Also let $R'_1, \ldots, R'_M$ be a collection of independent $\mathrm{Poi}_{\geq 1}(\lambda)$ random variables with $\lambda = (1 + o(1))\ln 2$ chosen such that $\mathbb{E}[R'_j] = \frac{k\Delta}{M} = (1 + o(1))2\ln 2$. Our first step is to prove the desired result for the $\{R'_j\}$. One can compute $\mathbb{E}[(R'_j)^2] = (2\ln 2)(1 + \ln 2) + o(1) = (1 + \ln 2 + o(1))\frac{k\Delta}{M}$. Standard sub-exponential tail bounds on the Poisson distribution (see [Can16]) imply $R'_j \leq \ln^2 n$ with probability $1 - n^{-\omega(1)}$ and $\mathbb{E}[(R'_j)^2 \mid R'_j \leq \ln^2 n] = \mathbb{E}[(R'_j)^2] \pm n^{-\omega(1)}$. Apply Hoeffding's inequality conditioned on the event $\{R'_j \leq \ln^2 n \text{ for all } j\}$ to conclude

$$\left|\left(\sum_j (R'_j)^2\right) - (1 + \ln 2 + o(1))k\Delta\right| \leq \tilde{O}(\sqrt{k}) \qquad \text{with probability } 1 - n^{-\omega(1)}.$$

Our next step is to transfer this claim to $\{R_j\}$ and then finally to $\{W_j\}$. Define the event $\mathcal{R} = \left\{\sum_{j=1}^M R'_j = k\Delta\right\}$. A folklore fact (e.g., implicit in [Dur19, Chapter 3.6]) is that the bin loads of the balls-into-bins experiment has the same distribution as i.i.d. Poisson random variables (of any variance) conditioned on the total number of balls being correct; this gives the equality of distributions

$$(R_1, \ldots, R_M) \overset{d}{=} (R'_1, \ldots, R'_M) \quad \text{given } \mathcal{R}.$$

Also, by the local limit theorem for sums of independent random variables, since $k\Delta$ is the expectation of $\sum_j R'_j$, we have $\Pr(\mathcal{R}) = n^{-O(1)}$. This means the probability of any event can only increase by a factor of $n^{O(1)}$ when passing from $\{R'_j\}$ to $\{R_j\}$, and in particular,

$$\left|\left(\sum_j R_j^2\right) - (1 + \ln 2 + o(1))k\Delta\right| \le \tilde{O}(\sqrt{k}) \qquad \text{with probability } 1 - n^{-\omega(1)}.$$

Finally, we use a similar argument to pass from $\{R_j\}$ to $\{W_j\}$. In Lemma 8.8 below, we show that with probability $n^{-O(1)}$, the balls-into-bins experiment generating $\{R_j\}$ features no multi-edges (i.e., the Δ balls from each infected individual fall into Δ distinct bins). Conditioned on having no multi-edges, $\{R_j\}$ has the same distribution as $\{W_j\}$, so similarly to above we conclude

$$\left|\left(\sum_j W_j^2\right) - (1 + \ln 2 + o(1))k\Delta\right| \le \tilde{O}(\sqrt{k}) \qquad \text{with probability } 1 - n^{-\omega(1)}.$$

as desired. □

Lemma 8.8. *Suppose k infected individuals each choose Δ tests out of M uniformly at random with replacement (so that multi-edges may occur), conditioned on having at least one infected individual in every test. With probability $n^{-O(1)}$, no multi-edges occur.*

Proof. Suppose each individual chooses Δ tests with replacement. Let A be the event that all M tests contain at least one infected individual, and let B be the event that no multi-edges occur. Our goal is to show $\Pr(B \mid A) = n^{-O(1)}$. It is clear that $\Pr(A \mid B) \ge \Pr(A \mid B^c)$. Using Bayes' rule,

$$\begin{aligned}\Pr(B \mid A) &= \frac{\Pr(A \mid B)\Pr(B)}{\Pr(A)} = \frac{\Pr(A \mid B)\Pr(B)}{\Pr(A \mid B)\Pr(B) + \Pr(A \mid B^c)\Pr(B^c)} \\ &\ge \frac{\Pr(B)}{\Pr(B) + \Pr(B^c)} = \Pr(B).\end{aligned}$$

Thus it suffices to show $\Pr(B) = n^{-O(1)}$, which is easy to establish directly due to independence across individuals. For any one individual, the expected number of "edge collisions" is $\binom{\Delta}{2}\frac{1}{M} \le \frac{\Delta^2}{M}$, so by Markov's inequality, the probability that this individual has no multi-edges is $\ge 1 - \frac{\Delta^2}{M}$. Now

$$\Pr(B) \ge \left(1 - \frac{\Delta^2}{M}\right)^k = \left(1 - \Theta\left(\frac{\ln n}{k}\right)\right)^k = \exp(-\Theta(\ln n)) = n^{-\Theta(1)},$$

completing the proof. □

8.2 Low-Degree Lower Bound: Proof of Theorem 3.2(b)

8.2.1 Orthogonal Polynomials

A key ingredient for the analysis will be an orthonormal (with respect to $\langle \cdot, \cdot \rangle_{\mathbb{Q}}$ defined in Section 7.2) basis for the polynomials $\{0,1\}^{NM} \to \mathbb{R}$. We first discuss orthogonal polynomials on a slice of the hypercube (which corresponds to the edges incident to one individual), and then show how to combine these to build an orthonormal basis for $\mathbb{Q}$.

Orthogonal Polynomials on a Slice of the Hypercube Consider the uniform distribution on the "slice of the hypercube" $\binom{[M]}{\Delta} := \{x \in \{0,1\}^M : \sum_i x_i = \Delta\}$, where $\Delta \le M/2$. The associated inner product between functions $\binom{[M]}{\Delta} \to \mathbb{R}$ is $\langle f, g\rangle := \mathbb{E}_{x \sim \mathrm{Unif}\binom{[M]}{\Delta}}[f(x)g(x)]$ and the associated norm is $\|f\| := \sqrt{\langle f, f\rangle}$. An orthonormal basis of polynomials with respect to this inner product is given in [Sri11, Fil16]. For ease of readability, we will not give the (somewhat complicated) full definition of the basis here. Instead, we will state only the properties of this basis that we actually need for the proof. See Appendix B for further details on how to extract these properties from [Fil16].

The basis elements are called $(\hat{\chi}_B)_{B \in \mathcal{B}_M}$. These are multivariate polynomials $\mathbb{R}^M \to \mathbb{R}$ that are orthonormal with respect to the above inner product $\langle \cdot, \cdot \rangle$ on the slice. The indices B belong to some set $\mathcal{B}_M$, the details of which will not be important for us. The indices have a notion of "size" $|B| \in \mathbb{N} := \{0, 1, 2, \ldots\}$, which coincides with the degree of the polynomial $\hat{\chi}_B$.

Fact 8.9. *For any integer $D \ge 0$, the set $\{\hat{\chi}_B : B \in \mathcal{B}_M, |B| \le \min(D, \Delta)\}$ is a complete orthonormal basis for the degree-D polynomials on $\binom{[M]}{\Delta}$. That is, for any polynomial $\mathbb{R}^M \to \mathbb{R}$ of degree (at most) D, there is a unique $\mathbb{R}$-linear combination of these basis elements that is equivalent[3] to f on $\binom{[M]}{\Delta}$.*

In particular, *any* function on the slice can be written as a polynomial of degree at most Δ.

Luckily, we will not need to use many specific details about the functions $\hat{\chi}_B$. We only need the following crude upper bound on their maximum value.

Fact 8.10. *For any $x \in \binom{[M]}{\Delta}$ and any $B \in \mathcal{B}_M$ with $|B| \le \Delta$, we have $|\hat{\chi}_B(x)| \le M^{2|B|}$.*

Orthogonal Polynomials for the Null Distribution The null distribution $\mathbb{Q}$ consists of N independent copies of the uniform distribution on $\binom{[M]}{\Delta}$, one for each individual. We can therefore use the following standard construction to build an orthonormal basis of polynomials for $\mathbb{Q}$. We denote the basis by $\{H_S\}_{S \in \mathcal{S}_{M,\Delta}}$ where

$$\mathcal{S}_{M,\Delta} = \{S = (B_1, \ldots, B_N) : B_i \in \mathcal{B}_M, |B_i| \le \Delta\},$$

defined by $H_S(X) = \prod_{i \in [N]} \hat{\chi}_{B_i}(X_i)$ where X_i is the collection of edge-indicator variables for edges incident to individual i. For $S = (B_1, \ldots, B_N)$, we define $|S| = \sum_{i \in [N]} |B_i|$, which is the degree of the polynomial H_S. As a consequence of Fact 8.9, $\{H_S : S \in \mathcal{S}_{M,\Delta}, |S| \le D\}$ is a complete orthonormal (with respect to $\langle \cdot, \cdot \rangle_{\mathbb{Q}}$) basis for the degree-$D$ polynomials $\{0,1\}^{NM} \to \mathbb{R}$.

We will need an upper bound on the number of basis elements of a given degree. Since $\{H_S\}$ are linearly independent, the number of indices $S \in \mathcal{S}_{M,\Delta}$ with $|S| \le D$ is at most the dimension (as a vector space over $\mathbb{R}$) of the degree-D polynomials $\{0,1\}^{NM} \to \mathbb{R}$. This dimension is at most the number of multilinear monomials of degree $\le D$, i.e., the number of subsets of $[NM]$ of cardinality $\le D$. This immediately gives the following.

Fact 8.11. *For any integer $D \ge 0$,*

$$|\{S \in \mathcal{S}_{M,\Delta} : |S| \le D\}| \le (1 + NM)^D.$$

[3]Here, "equivalent" means the two functions output the same value when given any input from $\binom{[M]}{\Delta}$. This is not the same as being equal as formal polynomials, e.g., x_1 is equivalent to x_1^2, and $\sum_i x_i$ is equivalent to the constant Δ.

8.2.2 Low-Degree Hardness

We follow the proof outline in Section 7.4, defining $\mathcal{U}$, $\mathbb{P}_u$, and $L_u = d\mathbb{P}_u/d\mathbb{Q}$ accordingly. With some abuse of notation, we will use u to refer to both the set of infected individuals and its indicator vector $u \in \{0,1\}^N$.

Lemma 8.12. *For any u, u', we have $\langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}} \le \langle L_u, L_{u'}\rangle_{\mathbb{Q}}$.*

Proof. We use a symmetry argument inspired by [BEH+22, Proposition 3.6]. Expanding in the orthonormal basis $\{H_S\}$ from Section 8.2.1, we have

$$\langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}} = \sum_{|S|\le D} \langle L_u, H_S\rangle_{\mathbb{Q}}\langle L_{u'}, H_S\rangle_{\mathbb{Q}} = \sum_{|S|\le D} \mathop{\mathbb{E}}_{X\sim\mathbb{P}_u}[H_S(X)] \mathop{\mathbb{E}}_{X\sim\mathbb{P}_{u'}}[H_S(X)]. \tag{8.6}$$

Let $V(S) = \{i \in [N] : \exists a \in [M], (i,a) \in S\}$, the set of all individuals "involved" in the basis function S. Note that if $V(S) \not\subseteq u$ then there exists some $i \in V(S)$ such that under $X \sim \mathbb{P}_u$ we have $X_i \sim \mathrm{Unif}\binom{[M]}{\Delta}$ independently from the rest of X, and thus $\mathbb{E}_{X\sim\mathbb{P}_u}[H_S(X)] = 0$. Similarly, if $V(S) \not\subseteq u'$ then $\mathbb{E}_{X\sim\mathbb{P}_{u'}}[H_S(X)] = 0$. On the other hand, if $V(S) \subseteq u \cap u'$ then (by symmetry) $\mathbb{P}_u$ and $\mathbb{P}_{u'}$ have the same marginal distribution when restricted to the variables $\{(i,a) : i \in u \cap u'\}$ and so $\mathbb{E}_{X\sim\mathbb{P}_u}[H_S(X)] = \mathbb{E}_{X\sim\mathbb{P}_{u'}}[H_S(X)]$. As a result, we have $\mathbb{E}_{X\sim\mathbb{P}_u}[H_S(X)]\,\mathbb{E}_{X\sim\mathbb{P}_{u'}}[H_S(X)] \ge 0$ for all S, i.e., every term on the right-hand side of (8.6) is nonnegative. This means $\langle L_u^{\le 0}, L_{u'}^{\le 0}\rangle_{\mathbb{Q}} \le \langle L_u^{\le 1}, L_{u'}^{\le 1}\rangle_{\mathbb{Q}} \le \langle L_u^{\le 2}, L_{u'}^{\le 2}\rangle_{\mathbb{Q}} \le \cdots \le \langle L_u^{\le \infty}, L_{u'}^{\le \infty}\rangle_{\mathbb{Q}} = \langle L_u, L_{u'}\rangle_{\mathbb{Q}}$. □

Following Section 7.4, recall the decomposition

$$\chi^2_{\le D}(\mathbb{P}\,\|\,\mathbb{Q}) + 1 = \mathcal{R}_{\le\delta}(D) + \mathcal{R}_{>\delta}(D) \tag{8.7}$$

(where we have made the dependence on D explicit) and choose

$$\delta = \max\left\{\frac{k^2}{N}, 1\right\} \cdot n^{2\gamma} \tag{8.8}$$

for a small constant $\gamma > 0$ to be chosen later. In light of Lemma 8.12, we have

$$\mathcal{R}_{\le\delta}(D) := \mathop{\mathbb{E}}_{u,u'\sim\mathcal{U}} \mathbb{1}_{\langle u,u'\rangle\le\delta} \langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}} \le \mathop{\mathbb{E}}_{u,u'\sim\mathcal{U}} \mathbb{1}_{\langle u,u'\rangle\le\delta} \langle L_u, L_{u'}\rangle_{\mathbb{Q}} =: \mathcal{T}_{\le\delta}. \tag{8.9}$$

It therefore remains to bound $\mathcal{R}_{>\delta}(D)$ and $\mathcal{T}_{\le\delta}$, which we will do in Lemmas 8.14 and 8.17 respectively.

Towards bounding $\mathcal{R}_{>\delta}(D)$, we need the following crude upper bound on $\langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}}$, which makes use of some basic properties of the orthogonal polynomials discussed in Section 8.2.1.

Lemma 8.13. *For any u, u', we have $\langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}} \le (NM+1)^D M^{4D}$.*

Proof. Consider the expansion (8.6). The number of terms in the sum on the right-hand side is at most $(NM+1)^D$ by Fact 8.11. Using Fact 8.10 and the definition of H_S (see Section 8.2.1), we have for any $|S| \le D$ and any $X \in \{0,1\}^{N\times M}$ that $|H_S(X)| \le M^{2D}$. Plugging these bounds back into (8.6) yields the claim. □

Lemma 8.14. *For any fixed $\theta \in (0,1)$, $c \in (0, (\ln 2)^{-2})$, and $\gamma > 0$, if δ is chosen according to (8.8) and $D = D_n$ satisfies $D \le n^\gamma$ then $\mathcal{R}_{>\delta}(D) = o(1)$.*

Proof. Fix u and consider the randomness over u'. In order to have $\langle u, u' \rangle > \delta$, there must exist a subset of size exactly $\lceil \delta \rceil$ contained in both u and u'. For any fixed subset of u of this size, the probability (over u') that it is also contained in u' is $\binom{N-\lceil\delta\rceil}{k-\lceil\delta\rceil}/\binom{N}{k}$. Taking a union bound over these subsets and using the choice of δ (8.8) along with the binomial bound $\binom{n}{k} \le \left(\frac{en}{k}\right)^k$ for $1 \le k \le n$,

$$\begin{aligned}
\Pr_{u,u' \sim \mathcal{U}}(\langle u, u' \rangle > \delta) &\le \binom{k}{\lceil\delta\rceil} \frac{\binom{N-\lceil\delta\rceil}{k-\lceil\delta\rceil}}{\binom{N}{k}} \le \binom{k}{\lceil\delta\rceil} \left(\frac{k}{N - \lceil\delta\rceil + 1}\right)^{\lceil\delta\rceil} \\
&\le \left(\frac{ek}{\lceil\delta\rceil}\right)^{\lceil\delta\rceil} \left(\frac{k}{N-k}\right)^{\lceil\delta\rceil} = \left(\frac{ek}{\lceil\delta\rceil} \cdot \frac{k}{N-k}\right)^{\lceil\delta\rceil} \\
&\le \left(\frac{2e}{n^{2\gamma}}\right)^{\lceil\delta\rceil} \le \left(\frac{2e}{n^{2\gamma}}\right)^{n^{2\gamma}} \le n^{-\gamma n^{2\gamma}},
\end{aligned} \tag{8.10}$$

provided $c < (\ln 2)^{-2}$ (so that $k = o(N)$). Combining this with Lemma 8.13,

$$\begin{aligned}
\mathcal{R}_{>\delta}(D) := \mathop{\mathbb{E}}_{u,u' \sim \mathcal{U}} \mathbb{1}_{\langle u,u' \rangle > \delta} \langle L_u^{\le D}, L_{u'}^{\le D} \rangle_{\mathbb{Q}} &\le \Pr_{u,u' \sim \mathcal{U}}(\langle u, u' \rangle > \delta) \cdot (NM+1)^D M^{4D} \\
&= n^{-\Omega(n^{2\gamma})} \cdot n^{O(D)},
\end{aligned} \tag{8.11}$$

which is $o(1)$ provided $D \le n^\gamma$. □

8.2.3 Low-Overlap Second Moment

This section is devoted to bounding $\mathcal{T}_{\le\delta}$ as defined in (8.9). Letting $E(u, X)$ denote the event that every test contains at least one individual from u, we can write

$$L_u(X) = \frac{d\mathbb{P}_u}{d\mathbb{Q}}(X) = \mathbb{Q}(E(u,X))^{-1} \mathbb{1}_{E(u,X)}$$

and

$$\langle L_u, L_{u'} \rangle_{\mathbb{Q}} = \mathbb{Q}(E(u,X))^{-2} \Pr_{X \sim \mathbb{Q}}(E(u,X) \cap E(u',X)) = \frac{\Pr_{X\sim\mathbb{Q}}(E(u',X) \mid E(u,X))}{\Pr_{X\sim\mathbb{Q}}(E(u,X))}. \tag{8.12}$$

Let $\mathcal{N}(u) \subseteq [M]$ denote the neighborhood of u, that is, the set of tests that contain at least one individual from u. Let $B(u, u', X)$ denote the event that the neighborhood of $u \cap u'$ has maximal size, that is, $|\mathcal{N}(u \cap u')| = \Delta \cdot |u \cap u'|$.

Lemma 8.15. *For any fixed u, u',*

$$\frac{\Pr_{X\sim\mathbb{Q}}(E(u',X) \mid E(u,X))}{\Pr_{X\sim\mathbb{Q}}(E(u,X))} \le \frac{1}{\Pr_{X\sim\mathbb{Q}}(B(u,u',X))}.$$

Proof. First, observe that the events $E(u,X)$ and $E(u',X)$ are conditionally independent given $|\mathcal{N}(u\cap u')|$. Furthermore, since $E(u',X)$ is clearly a monotone event with respect to $|\mathcal{N}(u\cap u')|$, we have for every $x\in\{0,1,\ldots,\Delta|u\cap u'|\}$,

$$\begin{aligned}\Pr_{X\sim\mathbb{Q}}(E(u',X)\,|\,|\mathcal{N}(u\cap u')|=x) &\le \Pr_{X\sim\mathbb{Q}}(E(u',X)\,|\,|\mathcal{N}(u\cap u')|=\Delta|u\cap u'|)\\ &= \Pr_{X\sim\mathbb{Q}}(E(u',X)\,|\,B(u,u',X)).\end{aligned}$$

Hence, combining with the aforementioned conditional independence we get

$$\Pr_{X\sim\mathbb{Q}}(E(u',X)\,|\,|\mathcal{N}(u\cap u')|=x,E(u,X))\le \Pr_{X\sim\mathbb{Q}}(E(u',X)\,|\,B(u,u',X)). \tag{8.13}$$

Using now (8.13) and the law of total probability we have

$$\begin{aligned}&\Pr_{X\sim\mathbb{Q}}(E(u',X)\,|\,E(u,X))\\ &=\sum_{x=0}^{\Delta|u\cap u'|}\Pr_{X\sim\mathbb{Q}}(|\mathcal{N}(u\cap u')|=x\,|\,E(u,X))\Pr_{X\sim\mathbb{Q}}(E(u',X)\,|\,|\mathcal{N}(u\cap u')|=x,E(u,X))\\ &\le \Pr_{X\sim\mathbb{Q}}(E(u',X)\,|\,B(u,u',X)).\end{aligned}\tag{8.14}$$

Given (8.14) and symmetry we conclude

$$\begin{aligned}\frac{\Pr_{X\sim\mathbb{Q}}(E(u',X)\,|\,E(u,X))}{\Pr_{X\sim\mathbb{Q}}(E(u,X))} &\le \frac{\Pr_{X\sim\mathbb{Q}}(E(u',X)\,|\,B(u,u',X))}{\Pr_{X\sim\mathbb{Q}}(E(u,X))}\\ &= \frac{\Pr_{X\sim\mathbb{Q}}(E(u',X)\,|\,B(u,u',X))}{\Pr_{X\sim\mathbb{Q}}(E(u,X)\,|\,B(u,u',X))\Pr_{X\sim\mathbb{Q}}(B(u,u',X))}\\ &= \frac{1}{\Pr_{X\sim\mathbb{Q}}(B(u,u',X))},\end{aligned}$$

completing the proof. □

Lemma 8.16. *For any fixed u,u' with $\langle u,u'\rangle=\ell$,*

$$\Pr_{X\sim\mathbb{Q}}(B(u,u',X))\ge 1-\ell^2M^{-1}\Delta^2.$$

Proof. We will compute $\mathbb{E}[Z]$ where Z is defined to be the number of "collisions", i.e., the number of tuples (i,j,a) where $i,j\in u\cap u'$ (with $i<j$) and $a\in[M]$ such that test a contains both individuals i and j. The number of tuples (i,j,a) is $\binom{\ell}{2}M$ and the probability that any fixed tuple is a collision is $(\Delta/M)^2$. Therefore $\mathbb{E}[Z]=\binom{\ell}{2}M^{-1}\Delta^2$. Since $B(u,u',X)$ is the event that $Z=0$, we have by Markov's inequality, $\Pr(B)=1-\Pr(Z\ge 1)\ge 1-\mathbb{E}[Z]\ge 1-\ell^2M^{-1}\Delta^2$. □

Lemma 8.17. *For any fixed $\theta\in(0,1)$ and $c>0$ satisfying $c<c_{\mathrm{LD}}^{\mathrm{CC}}$, there exists $\gamma=\gamma(\theta,c)$ such that if δ is chosen according to (8.8) then $\mathcal{T}_{\le\delta}=1+o(1)$.*

Proof. Combining (8.12) with Lemmas 8.15 and 8.16, we have

$$\langle L_u,L_{u'}\rangle_{\mathbb{Q}}\le(1-\langle u,u'\rangle^2M^{-1}\Delta^2)^{-1}$$

and so

$$\mathcal{T}_{\le\delta} := \mathop{\mathbb{E}}_{u,u'\sim\mathcal{U}} \mathbb{1}_{\langle u,u'\rangle\le\delta} \langle L_u, L_{u'}\rangle_{\mathbb{Q}} \le (1-\delta^2 M^{-1}\Delta^2)^{-1}.$$

Recalling $M^{-1}\Delta^2 = \tilde{\Theta}(k^{-1})$, we have $\mathcal{T}_{\le\delta} = 1 + o(1)$ provided that $\delta \ll \sqrt{k}$ (where $\ll$ hides factors of $\ln n$). Recalling the choice of δ (8.8), this reduces to the sufficient conditions $\frac{k^2}{N} n^{2\gamma} \ll \sqrt{k}$ and $n^{2\gamma} \ll \sqrt{k}$. Choosing γ sufficiently small and recalling the scaling for N, these reduce to $\frac{3}{2}\theta + (1-\theta)c(\ln 2)^2 < 1$, which is equivalent to $c < c_{\mathrm{LD}}^{\mathrm{CC}}$. □

Proof of Theorem 3.2(b). Provided $c < c_{\mathrm{LD}}^{\mathrm{CC}}$ (which also implies $c < (\ln 2)^{-2}$), we can combine (8.7), (8.9), Lemma 8.14, and Lemma 8.17 to conclude $\chi^2_{\le D}(\mathbb{P} \,\|\, \mathbb{Q}) = o(1)$ for any $D \le n^\gamma = n^{\Omega(1)}$. By Lemma 7.3, this completes the proof of Theorem 3.2(b). □

9 Detection in the Bernoulli Design

For convenience we recall the definition

$$c_{\mathrm{LD}}^{\mathrm{B}} = \begin{cases} -\frac{1}{\ln^2 2} W_0(-\exp(-\frac{\theta}{1-\theta}\ln 2 - 1)) & \text{if } 0 < \theta < \frac{1}{2}(1 - \frac{1}{4\ln 2 - 1}), \\ \frac{1}{\ln 2} \cdot \frac{1-2\theta}{1-\theta} & \text{if } \frac{1}{2}(1 - \frac{1}{4\ln 2-1}) \le \theta < \frac{1}{2}, \\ 0 & \text{if } \frac{1}{2} \le \theta < 1, \end{cases}$$

where $W_0(x)$ denotes the unique $y \ge -1$ satisfying $ye^y = x$. Throughout this section, the following reformulation will be helpful: for $\theta \in (0,1)$ and $c > 0$, the condition $c > c_{\mathrm{LD}}^{\mathrm{B}}$ is equivalent to $\tau(c) < \frac{\theta}{1-\theta}$, where the function τ is given by

$$\tau(c) = \begin{cases} 1 - c\ln 2 & \text{if } 0 < c \le \frac{1}{2(\ln 2)^2}, \\ c\ln 2 - \frac{1}{\ln 2}[1 + \ln(c(\ln 2)^2)] & \text{if } \frac{1}{2(\ln 2)^2} < c < \frac{1}{(\ln 2)^2}, \\ 0 & \text{if } c \ge \frac{1}{(\ln 2)^2}. \end{cases} \tag{9.1}$$

9.1 Upper Bounds: Proof of Theorem 3.3(a) and Theorem 3.4(a)

First, for Theorem 3.4(a), it is known that if $c > 1/\ln 2$ then approximate recovery is possible (see e.g. [IZ21, Lemma 2.1]). Hence, by Proposition C.1 strong detection is also possible.

In this section we give a polynomial-time algorithm for strong detection whenever $\tau(c) < \frac{\theta}{1-\theta}$ (recall the reformulation in (9.1)). We also show how to turn this algorithm into an $O(\ln n)$-degree polynomial that achieves strong separation (see Section 9.1.4). This will complete the proof of both Theorem 3.4(a) and Theorem 3.3(a).

Define the test statistic T to be the number of individuals of (graph-theoretic) degree at least $d = 2tqM$ for a constant $t > 1$ to be chosen later. That is,

$$T = \sum_{i=1}^{N} \mathbb{1}_{d_i \ge d}$$

where d_i is the degree of individual i (i.e., the number of tests that i participates in).

9.1.1 Non-Infected

First consider the contribution T_- to T from non-infected individuals. (Under $\mathbb{Q}$, we consider all individuals to be "non-infected.") Let $N' = |V_-|$ be the number of non-infected individuals, which is equal to N under $\mathbb{Q}$ and $N-k$ under $\mathbb{P}$. The degree of each $i \in V_-$ is $d_i \sim \mathrm{Bin}(M,q)$ and these are independent. Define

$$p_- = \Pr(\mathrm{Bin}(M,q) \geq d)$$

so that $T_- \sim \mathrm{Bin}(N', p_-)$. This means $\mathbb{E}[T_-] = N'p_-$ and $\mathrm{Var}(T_-) = N'p_-(1-p_-) \leq N'p_-$. We can bound p_- using the Binomial tail bound (Proposition A.2):

$$p_- \leq \exp(-MD(2tq \,\|\, q))$$

where, using Lemma A.4,

$$D(2tq \,\|\, q) \geq q(2t\ln 2t - 2t + 1) - O(q^2),$$

where $O(\cdot)$ hides a constant depending only on t. This means

$$\begin{aligned} p_- &\leq \exp\left[-\left(\frac{c}{2} + o(1)\right) k\ln(n/k) \cdot q(2t\ln 2t - 2t + 1 - o(1))\right] \\ &\leq n^{-(1-\theta)\frac{c}{2}(\ln 2)(2t\ln 2t - 2t + 1) + o(1)}. \end{aligned} \tag{9.2}$$

9.1.2 Infected

Now consider the contribution T_+ to T from infected individuals (under $\mathbb{P}$). Under $\mathbb{P}$ there are $k = |V_+|$ infected individuals. Each $i \in V_+$ has degree $d_i \sim \mathrm{Bin}(M, 2q)$ (see (9.3)), but these are not independent. Define

$$p_+ = \Pr(\mathrm{Bin}(M, 2q) \geq d).$$

Lemma 9.1. *We have*

$$p_+ = n^{-(1-\theta)c(\ln 2)(t\ln t - t + 1) + o(1)}.$$

Proof. We first give a lower bound using the Binomial tail lower bound (Proposition A.3 and Lemma A.4):

$$\begin{aligned} p_+ &\geq \frac{1}{\sqrt{8d(1-d/M)}} \exp\left(-MD\left(\frac{d}{M} \,\middle\|\, 2q\right)\right) \\ &\geq \frac{1}{\sqrt{16tqM}} \exp(-MD(2tq \,\|\, 2q)) \\ &\geq \frac{1}{\sqrt{16t}} \left(\left(\frac{c}{2}\ln 2 + o(1)\right)\ln(n/k)\right)^{-1/2} \exp[-M(2tq\ln t + 2q - 2tq + O(q^2))] \\ &\geq n^{-o(1)} \exp[-(c\ln 2 + o(1))(t\ln t - t + 1 + o(1))\ln(n/k)] \\ &= n^{-(1-\theta)c(\ln 2)(t\ln t - t + 1) - o(1)} \end{aligned}$$

as desired. The matching upper bound is proved similarly, using the Binomial tail upper bound (Proposition A.2). □

This gives us control of the mean of T_+, since $\mathbb{E}[T_+] = kp_+$. Next we will bound the variance of T_+ which is more difficult because the d_i are not independent. However, we will leverage negative correlations between the d_i to effectively reduce to the independent case. Fix two distinct infected individuals i, j and a test a. Recall that X_{ia} is the indicator for edge (i, a). We will compute the joint distribution of X_{ia} and X_{ja}. Letting E_a be the event that a is connected to at least one of the k infected individuals,

$$\begin{aligned} q^2 = \mathop{\mathbb{E}}_{\mathbb{Q}}[X_{ia}X_{ja}] &= \mathbb{Q}(E_a) \mathop{\mathbb{E}}_{\mathbb{Q}}[X_{ia}X_{ja}|E_a] + \mathbb{Q}(\overline{E_a}) \mathop{\mathbb{E}}_{\mathbb{Q}}[X_{ia}X_{ja}|\overline{E_a}] \\ &= \frac{1}{2} \cdot \mathop{\mathbb{E}}_{\mathbb{Q}}[X_{ia}X_{ja}|E_a] + \frac{1}{2} \cdot 0 \end{aligned}$$

and so

$$\mathbb{P}(X_{ia} = X_{ja} = 1) = \mathop{\mathbb{E}}_{\mathbb{P}}[X_{ia}X_{ja}] = \mathop{\mathbb{E}}_{\mathbb{Q}}[X_{ia}X_{ja}|E_a] = 2q^2.$$

Similarly, we can compute

$$\mathbb{P}(X_{ia} = X_{ja} = 0) = 1 - 4q + 2q^2$$

and

$$\mathbb{P}(X_{ia} = 1 \wedge X_{ja} = 0) = \mathbb{P}(X_{ia} = 0 \wedge X_{ja} = 1) = 2q(1 - q),$$

and so we know the joint distribution of X_{ia} and X_{ja} under $\mathbb{P}$. Due to independence across tests, we also know the joint distribution of $\{X_{ia}\}_{a\in[M]}$ and $\{X_{ja}\}_{a\in[M]}$. In particular, we have the conditional probabilities

$$\mathbb{P}(X_{ja} = 1 \,|\, X_{ia} = 1) = \frac{2q^2}{2q} = q$$

and

$$\mathbb{P}(X_{ja} = 1 \,|\, X_{ia} = 0) = \frac{2q(1 - q)}{1 - 2q},$$

as well as the conditional distribution

$$d_j \,|\, \{d_i = w\} \sim \mathrm{Bin}(w, q) + \mathrm{Bin}\left(M - w, \frac{2q(1 - q)}{1 - 2q}\right) =: \mathcal{D}_w$$

where the two binomials are independent. Since $\frac{2q(1-q)}{1-2q} > q$ (recall $q = \frac{\nu}{k} \to 0$), the distribution $\mathcal{D}_w$ stochastically dominates $\mathcal{D}_{w+1}$ for all $0 \le w < M$. As a result,

$$\mathbb{P}(d_j \ge d \,|\, d_i \ge d) \le \mathbb{P}(d_j \ge d),$$

and so

$$\mathbb{P}(d_i \ge d \wedge d_j \ge d) = \mathbb{P}(d_i \ge d)\mathbb{P}(d_j \ge d \,|\, d_i \ge d) \le \mathbb{P}(d_i \ge d)\mathbb{P}(d_j \ge d) = p_+^2.$$

We can now compute

$$\begin{aligned}
\mathrm{Var}(T_+) &= \mathbb{E}[T_+^2] - \mathbb{E}[T_+]^2 \\
&= \mathbb{E}\left[\left(\sum_{i \in V_+} \mathbb{1}_{d_i \geq d}\right)^2\right] - (kp_+)^2 \\
&= \mathbb{E}\left[\sum_i \mathbb{1}_{d_i \geq d} + \sum_{i \neq j} \mathbb{1}_{d_i \geq d}\mathbb{1}_{d_j \geq d}\right] - (kp_+)^2 \\
&\leq kp_+ + k(k-1)p_+^2 - (kp_+)^2 \\
&= kp_+(1-p_+) \\
&\leq kp_+.
\end{aligned}$$

9.1.3 Putting it Together

Let's recap what we have so far. Under $\mathbb{Q}$, we have $T = T_-$, which has mean and variance

$$\mathop{\mathbb{E}}_{\mathbb{Q}}[T] = Np_- \qquad \text{and} \qquad \mathop{\mathrm{Var}}_{\mathbb{Q}}(T) \leq Np_-.$$

Under $\mathbb{P}$, we have $T = T_+ + T_-$ (with T_+ and T_- independent), which has mean and variance

$$\mathop{\mathbb{E}}_{\mathbb{P}}[T] = (N-k)p_- + kp_+ \qquad \text{and} \qquad \mathop{\mathrm{Var}}_{\mathbb{P}}(T) \leq (N-k)p_- + kp_+.$$

In order to distinguish $\mathbb{P}$ and $\mathbb{Q}$ with high probability by thresholding T, it suffices (by Chebyshev's inequality) to have

$$\sqrt{\mathop{\mathrm{Var}}_{\mathbb{Q}}(T)} + \sqrt{\mathop{\mathrm{Var}}_{\mathbb{P}}(T)} = o\left(\mathop{\mathbb{E}}_{\mathbb{P}}[T] - \mathop{\mathbb{E}}_{\mathbb{Q}}[T]\right),$$

which yields the sufficient condition

$$\sqrt{Np_-} + \sqrt{kp_+} = o(k(p_+ - p_-)).$$

Thus, it sufficies to have all of the following three conditions:

(i) $p_- = o(p_+)$,

(ii) $\sqrt{Np_-} = o(kp_+)$,

(iii) $\sqrt{kp_+} = o(kp_+)$.

Recall from above (see (9.2) and Lemma 9.1) the asymptotics

$$k = n^{\theta + o(1)}, \qquad N = n^{1-(1-\theta)\frac{c}{2}\ln 2 + o(1)}, \qquad p_- \leq n^{-(1-\theta)\frac{c}{2}(\ln 2)(2t\ln 2t - 2t + 1) + o(1)},$$

$$p_+ = n^{-(1-\theta)c(\ln 2)(t\ln t - t + 1) + o(1)}.$$

These can be used to rewrite the three conditions as the following sufficient conditions:

(i') $t > 1$ (which, recall, we also assumed earlier),

(ii') $1 + c(\ln 2)(t \ln \frac{t}{2} - t + 1) < \frac{\theta}{1-\theta}$,

(iii') $c(\ln 2)(t \ln t - t + 1) < \frac{\theta}{1-\theta}$.

First consider the case $0 \le c \le \frac{1}{2(\ln 2)^2}$. In this case, choose $t = 2$ (which minimizes the left-hand side of (ii')). This causes (iii') to become subsumed by (ii'). Also, (ii') simplifies to $1 - c \ln 2 < \frac{\theta}{1-\theta}$, which matches the desired condition $\tau(c) < \frac{\theta}{1-\theta}$.

Next consider the case $\frac{1}{2(\ln 2)^2} < c < \frac{1}{(\ln 2)^2}$. In this case, choose $t = \frac{1}{c(\ln 2)^2}$, which satisfies (i') due to the assumption on c. This causes (ii') and (iii') to become equivalent, both reducing to the desired condition $c \ln 2 - \frac{1}{\ln 2}[1 + \ln(c(\ln 2)^2)] < \frac{\theta}{1-\theta}$.

Finally, consider the case $c \ge \frac{1}{(\ln 2)^2}$. For any $\theta \in (0, 1)$, it suffices to take $t = 1 + \varepsilon$ for sufficiently small $\varepsilon > 0$ for all the conditions to be satisfied.

9.1.4 Polynomial Approximation

Above, we have shown that the test statistic $T = T(X)$ strongly separates $\mathbb{P}$ and $\mathbb{Q}$, but T is not a polynomial. We will now show that when $\tau(c) < \frac{\theta}{1-\theta}$ there is a degree-$O(\ln n)$ polynomial that strongly separates $\mathbb{P}$ and $\mathbb{Q}$, and we will do this using a polynomial approximation for T.

Recall $T = \sum_{i=1}^{N} \mathbb{1}_{d_i \ge d}$ where d_i is the degree of individual i in the graph. We define the following polynomial approximation for the indicator $\mathbb{1}_{x \ge d}$: for $a := \lceil d \rceil$ and some integer $b > a$ (to be chosen later),

$$I_b(x) = \sum_{a \le j < b} \prod_{\substack{0 \le \ell < b \\ \ell \ne j}} \frac{x - \ell}{j - \ell}.$$

Note that I_b is a polynomial in x of degree $b - 1$, which we will choose to be $O(\ln n)$. By construction, $I_b(x) = \mathbb{1}_{x \ge d}$ for all $x \in \{0, 1, 2, \ldots, b - 1\}$. Therefore

$$I_b(d_i) = \mathbb{1}_{d_i \ge d} + \mathbb{1}_{d_i \ge b} \cdot (I_b(d_i) - 1).$$

The key calculation we need is a bound on the second moment of the error term

$$E_{i,b} := \mathbb{1}_{d_i \ge b} \cdot (I_b(d_i) - 1).$$

Recall $d_i \sim \mathrm{Bin}(M, \bar{q})$ where $\bar{q}$ is either q or $2q$ (depending on whether individual i is infected).

Lemma 9.2. *Suppose $d_i \sim \mathrm{Bin}(M, \bar{q})$ for $\bar{q} \in \{q, 2q\}$. For any constant $C > 0$ there exists a constant $B = B(C, \theta, c) > 0$ such that when choosing b to be the first odd integer greater than $B \ln n$,*

$$\mathbb{E}[E_{i,b}^2] \le n^{-C}.$$

Proof. We first note that it suffices (up to a change in the constant B) to show the result for

$$\tilde{E}_{i,b} := \mathbb{1}_{d_i \ge b} I_b(d_i)$$

in place of $E_{i,b}$. This is because

$$E_{i,b}^2 \le 2(\tilde{E}_{i,b}^2 + \mathbb{1}_{d_i \ge b})$$

and

$$\mathbb{E}[\mathbb{1}_{d_i \geq b}] = \Pr(d_i \geq b),$$

which can be made smaller than n^{-2C} by choosing B large enough (similarly to the calculation in Section 9.1.1).

Now for any $x \geq b$ we have the bound

$$|I_b(x)| \leq (b-a)\frac{x^{b-1}}{\left[\left(\frac{b-1}{2}\right)!\right]^2}$$

where we have used the fact that $\prod_{0 \leq \ell < b,\, \ell \neq j} |j - \ell|$ is minimized when j lies at the center of the range $\{0, 1, \ldots, b-1\}$. We will also make use of the bounds $\binom{n}{k} \leq \left(\frac{ne}{k}\right)^k$ (for all $1 \leq k \leq n$) and $n! \geq \left(\frac{n}{e}\right)^n$ (for all $n \geq 1$). We have

$$\begin{aligned}
\mathbb{E}[\tilde{E}_{i,b}^2] &= \sum_{x=b}^{\infty} \Pr(d_i = x) I_b(x)^2 \\
&\leq \sum_{x=b}^{\infty} \binom{M}{x} \bar{q}^x (1-\bar{q})^{M-x} \cdot (b-a)^2 \frac{x^{2(b-1)}}{\left[\left(\frac{b-1}{2}\right)!\right]^4} \\
&\leq \sum_{x=b}^{\infty} \left(\frac{Me}{x}\right)^x \bar{q}^x (1-\bar{q})^{M-x} \cdot (b-a)^2 \frac{x^{2(b-1)}}{\left[\left(\frac{b-1}{2e}\right)^{(b-1)/2}\right]^4} \\
&= \sum_{x=b}^{\infty} (b-a)^2 (1-\bar{q})^M \left(\frac{Me}{x}\right)^x \left(\frac{\bar{q}}{1-\bar{q}}\right)^x \left(\frac{2ex}{b-1}\right)^{2(b-1)} \\
&\leq \sum_{x=b}^{\infty} b^2 \left(\frac{Me}{x}\right)^x (3q)^x \left(\frac{2ex}{b-1}\right)^{2(b-1)} \\
&= \sum_{x=b}^{\infty} b^2 \left(\frac{3eMq}{x}\right)^x \left(\frac{2ex}{b-1}\right)^{2(b-1)} =: \sum_{x=b}^{\infty} r_x.
\end{aligned}$$

To complete the proof, we will show that the first term is $r_b \leq \frac{1}{2} n^{-C}$ and the ratio of successive terms is $\frac{r_{x+1}}{r_x} \leq \frac{1}{2}$ for all $x \geq b$. For the first step,

$$\begin{aligned}
r_b &= b^2 \left(\frac{3eMq}{b}\right)^b \left(\frac{2eb}{b-1}\right)^{2(b-1)} \\
&= b^2 \left(\frac{b-1}{2eb}\right)^2 \left(\frac{12e^3 Mqb^2}{b(b-1)^2}\right)^b \\
&\leq b^2 \left(\frac{12e^3 Mqb^2}{b(b-1)^2}\right)^b \\
&= b^2 \left(12e^3(c\nu/2 + o(1))(1-\theta) \cdot \frac{b \ln n}{(b-1)^2}\right)^b.
\end{aligned}$$

Recalling $B \ln n \le b \le B \ln n + 2$ and choosing B sufficiently large, the above is

$$\le b^2 (1/e)^b \le (B \ln n + 2)^2 e^{-B \ln n} \le \frac{1}{2} n^{-C}$$

as desired. For the second step, for $x \ge b$,

$$\begin{aligned}
\frac{r_{x+1}}{r_x} &= 3eMq \cdot \frac{x^x}{(x+1)^{x+1}} \left(\frac{x+1}{x}\right)^{2(b-1)} \\
&= \frac{3eMq}{x+1} \left(\frac{x+1}{x}\right)^{2(b-1)-x} \\
&\le \frac{3eMq}{x+1} \left(1 + \frac{1}{x}\right)^{b-2} \\
&\le \frac{3eMq}{x+1} \left(1 + \frac{1}{b}\right)^{b} \\
&\le \frac{3eMq}{x+1} \cdot e \\
&= \frac{3e^2 (c\nu/2 + o(1))(1-\theta) \ln n}{x+1} \\
&\le \frac{3e^2 (c\nu/2 + o(1))(1-\theta) \ln n}{B \ln n}
\end{aligned}$$

which can be made $\le \frac{1}{2}$ by choosing B sufficiently large. □

Using Lemma 9.2 we can now show that under either $\mathbb{P}$ or $\mathbb{Q}$, the first two moments of $I_b(d_i)$ and $\mathbb{1}_{d_i \ge d}$ nearly match:

$$\left| \mathop{\mathbb{E}}_{\mathbb{Q}}[I_b(d_i)] - \mathop{\mathbb{E}}_{\mathbb{Q}}[\mathbb{1}_{d_i \ge d}] \right| = \left| \mathop{\mathbb{E}}_{\mathbb{Q}} E_{i,b} \right| \le \sqrt{\mathop{\mathbb{E}}_{\mathbb{Q}} E_{i,b}^2} \le n^{-C/2},$$

$$\begin{aligned}
\left| \mathop{\mathbb{E}}_{\mathbb{Q}}[I_b(d_i) I_b(d_j)] - \mathop{\mathbb{E}}_{\mathbb{Q}}[\mathbb{1}_{d_i \ge d} \mathbb{1}_{d_j \ge d}] \right| &= \left| \mathop{\mathbb{E}}_{\mathbb{Q}}[\mathbb{1}_{d_j \ge d} E_{i,b} + \mathbb{1}_{d_i \ge d} E_{j,b} + E_{i,b} E_{j,b}] \right| \\
&\le \sqrt{\mathop{\mathbb{E}}_{\mathbb{Q}} E_{i,b}^2} + \sqrt{\mathop{\mathbb{E}}_{\mathbb{Q}} E_{j,b}^2} + \sqrt{\mathop{\mathbb{E}}_{\mathbb{Q}} E_{i,b}^2 \cdot \mathop{\mathbb{E}}_{\mathbb{Q}} E_{j,b}^2} \\
&\le 3 n^{-C/2},
\end{aligned}$$

and similarly for $\mathbb{P}$.

Define the polynomial

$$\tilde{T}(X) = \sum_{i=1}^{N} I_b(d_i),$$

which has degree $b - 1 = O(\ln n)$. Using the bounds above, the first two moments of $\tilde{T}$ and T nearly match:

$$\left| \mathop{\mathbb{E}}_{\mathbb{Q}}[\tilde{T}] - \mathop{\mathbb{E}}_{\mathbb{Q}}[T] \right| = \left| \sum_{i=1}^{N} \mathop{\mathbb{E}}_{\mathbb{Q}}[I_b(d_i) - \mathbb{1}_{d_i \ge d}] \right| \le N \cdot n^{-C/2} = n^{O(1) - C/2},$$

$$\left|\mathop{\mathbb{E}}_{\mathbb{Q}}[\tilde{T}^2] - \mathop{\mathbb{E}}_{\mathbb{Q}}[T^2]\right| = \left|\sum_{1\le i,j\le N} \mathop{\mathbb{E}}_{\mathbb{Q}}[I_b(d_i)I_b(d_j) - \mathbb{1}_{d_i\ge d}\mathbb{1}_{d_j\ge d}]\right|$$
$$\le N^2 \cdot 3n^{-C/2} = n^{O(1)-C/2},$$

$$\begin{aligned}
\left|\mathop{\mathrm{Var}}_{\mathbb{Q}}[\tilde{T}] - \mathop{\mathrm{Var}}_{\mathbb{Q}}[T]\right| &= \left|\mathop{\mathbb{E}}_{\mathbb{Q}}[\tilde{T}^2] - \mathop{\mathbb{E}}_{\mathbb{Q}}[T^2] - \mathop{\mathbb{E}}_{\mathbb{Q}}[\tilde{T}]^2 + \mathop{\mathbb{E}}_{\mathbb{Q}}[T]^2\right| \\
&\le \left|\mathop{\mathbb{E}}_{\mathbb{Q}}[\tilde{T}^2] - \mathop{\mathbb{E}}_{\mathbb{Q}}[T^2]\right| + \left|\mathop{\mathbb{E}}_{\mathbb{Q}}[\tilde{T} - T]\mathop{\mathbb{E}}_{\mathbb{Q}}[\tilde{T} + T]\right| \\
&\le 3N^2n^{-C/2} + Nn^{-C/2}\left|\mathop{\mathbb{E}}_{\mathbb{Q}}[\tilde{T} + T]\right| \\
&\le 3N^2n^{-C/2} + Nn^{-C/2}\left(2\mathop{\mathbb{E}}_{\mathbb{Q}}[T] + Nn^{-C/2}\right) \\
&\le 3N^2n^{-C/2} + Nn^{-C/2}\left(2N + Nn^{-C/2}\right) \\
&= n^{O(1)-C/2}
\end{aligned}$$

and similarly for $\mathbb{P}$ (where the $O(1)$ terms do not depend on C).

Suppose $\tau(c) < \frac{\theta}{1-\theta}$. We have shown previously (see Section 9.1.3) that T strongly separates $\mathbb{P}$ and $\mathbb{Q}$ with separation $\mathbb{E}_{\mathbb{P}}[T] - \mathbb{E}_{\mathbb{Q}}[T] = (1 - o(1))kp_+ \ge n^{-O(1)}$. (In fact, the separation is larger than 1, but the simpler bound $n^{-O(1)}$ will suffice.) By taking C sufficiently large, the mean and variance of $\tilde{T}$ match those of T (under either $\mathbb{P}$ or $\mathbb{Q}$) up to an error that is negligible compared to the separation $\mathbb{E}_{\mathbb{P}}[T] - \mathbb{E}_{\mathbb{Q}}[T]$. Therefore $\tilde{T}$ strongly separates $\mathbb{P}$ and $\mathbb{Q}$.

9.2 Lower Bounds: Proof of Theorem 3.3(b) and Theorem 3.4(b)

The proofs in this section are based on bounding the chi-squared divergence and its conditional/low-degree variants as described in Section 7.

9.2.1 Conditional Planted Distribution

We will condition $\mathbb{P}$ on the following "good" event A. Let A be the event that all infected individuals have degree at most d, for a particular d which will be chosen so that $\mathbb{P}(A) = 1 - o(1)$. Below, we will show that it is sufficient to take $d = 2tqM$ for any constant $t > 1$ satisfying (9.5). Let $\tilde{\mathbb{P}}$ be the conditional distribution $\mathbb{P} \,|\, A$.

Suppose individual i is infected and let a be a test. Letting X_{ia} be the indicator for edge (i, a) and letting E_a be the event that a is connected to at least one infected individual,

$$q = \mathop{\mathbb{E}}_{\mathbb{Q}}[X_{ia}] = \mathbb{Q}(E_a)\mathop{\mathbb{E}}_{\mathbb{Q}}[X_{ia}|E_a] + \mathbb{Q}(\overline{E_a})\mathop{\mathbb{E}}_{\mathbb{Q}}[X_{ia}|\overline{E_a}] = \frac{1}{2}\cdot\mathop{\mathbb{E}}_{\mathbb{Q}}[X_{ia}|E_a] + \frac{1}{2}\cdot 0$$

and so

$$\mathop{\mathbb{E}}_{\mathbb{P}}[X_{ia}] = \mathop{\mathbb{E}}_{\mathbb{Q}}[X_{ia}|E_a] = 2q. \tag{9.3}$$

So under $\mathbb{P}$, the degree d_i of individual i is distributed as $d_i \sim \mathrm{Bin}(M, 2q)$ (but these are not independent across i).

Using the Binomial tail bound (Proposition A.2), for any constant $t > 1$,

$$\Pr\left(d_i \geq 2tqM\right) \leq \exp\left(-MD\left(2tq \,\|\, 2q\right)\right)$$

where, using Lemma A.4,

$$D(2tq \,\|\, 2q) \geq 2q(t \ln t - t + 1) - O(q^2),$$

where $O(\cdot)$ hides a constant depending only on t. This means, letting V_+ denote the set of infected individuals,

$$\begin{aligned}\Pr\left(\exists i \in V_+, d_i \geq 2tqM\right) &\leq k \exp\left[-2qM(t \ln t - t + 1 - O(q))\right] \\ &= n^{\theta + o(1)} n^{-(1-\theta)c(\ln 2)(t \ln t - t + 1) + o(1)} \\ &= n^{\theta - (1-\theta)c(\ln 2)(t \ln t - t + 1) + o(1)}. \end{aligned} \tag{9.4}$$

To ensure that A is a high-probability event under $\mathbb{P}$, we need to choose d so that (9.4) is $o(1)$, that is, $d = 2tqM$ where $t > 1$ is a constant satisfying

$$c(\ln 2)(t \ln t - t + 1) > \frac{\theta}{1-\theta}. \tag{9.5}$$

9.2.2 Conditional Chi-Squared

With some abuse of notation, we will use u to refer to both the set of infected individuals and its indicator vector $u \in \{0,1\}^N$. Let $A = A(u, X)$ be the "good" event defined in Section 9.2.1 above (namely, the individuals in u all have degree at most d), and let $\tilde{\mathbb{P}}$ denote the conditional distribution $\mathbb{P} \,|\, A$. For a test a, let $E_a = E_a(u, X)$ be the event that a contains at least one infected individual. Let $E = \cap_a E_a$. Define $\mathcal{U}$, $\tilde{\mathbb{P}}_u$, and $L_u = d\tilde{\mathbb{P}}_u / d\mathbb{Q}$ as in Section 7.4. Compute

$$\begin{aligned} L_u(X) = \frac{d\tilde{\mathbb{P}}}{d\mathbb{P}}(X) \cdot \frac{d\mathbb{P}}{d\mathbb{Q}}(X) &= \mathbb{P}(A)^{-1} \mathbb{1}_{A(u,X)} \cdot \mathbb{Q}(E(u,X))^{-1} \mathbb{1}_{E(u,X)} \\ &= \mathbb{P}(A)^{-1} \, 2^M \mathbb{1}_{E(u,X)} \mathbb{1}_{A(u,X)} \end{aligned}$$

and

$$\langle L_u, L_{u'} \rangle_{\mathbb{Q}} = \mathbb{P}(A)^{-2} \, 2^{2M} \Pr_{X \sim \mathbb{Q}} \left(E(u,X) \cap E(u',X) \cap A(u,X) \cap A(u',X)\right). \tag{9.6}$$

Letting $\ell = \langle u, u' \rangle$,

$$\chi^2(\tilde{\mathbb{P}} \,\|\, \mathbb{Q}) + 1 = \mathop{\mathbb{E}}_{u,u' \sim \mathcal{U}} \langle L_u, L_{u'} \rangle_{\mathbb{Q}} = \sum_{\ell=0}^{k} \Pr(\ell) \langle L_u, L_{u'} \rangle_{\mathbb{Q}}, \tag{9.7}$$

where $\Pr(\ell)$ is shorthand for

$$\Pr_{u,u' \sim \mathcal{U}} (\langle u, u' \rangle = \ell) = \frac{\binom{k}{\ell}\binom{N-k}{k-\ell}}{\binom{N}{k}}. \tag{9.8}$$

Note that the term $\langle L_u, L_{u'}\rangle_{\mathbb{Q}}$ in (9.7) depends on u, u' only through $\ell = \langle u, u'\rangle$ and is thus well-defined as a function of ℓ alone.

We will now work on bounding various parts of the formula (9.7). First recall $\mathbb{P}(A) = 1 - o(1)$. To handle $\Pr(\ell)$ we have

$$\frac{\binom{N-k}{k-\ell}}{\binom{N}{k}} \le \frac{\binom{N}{k-\ell}}{\binom{N}{k}} = \frac{k!(N-k)!}{(k-\ell)!(N-k+\ell)!} \le \left(\frac{k}{N-k}\right)^{\ell} = n^{-\ell[(1-\theta)(1-\frac{c}{2}\ln 2)+o(1)]} \tag{9.9}$$

provided $c < \frac{2}{\ln 2}$ (so that $k = o(N)$). Also, for $\ell \ge 1$ we have the standard bound

$$\binom{k}{\ell} \le \left(\frac{ek}{\ell}\right)^{\ell}. \tag{9.10}$$

Next we will bound the final term $\Pr_{X\sim\mathbb{Q}}(\cdots)$ in (9.6). Let $\tilde{E}_a(u, u', X)$ be the event that test a contains at least one individual from $u \cap u'$. Note that $\tilde{E}_a(u, u', X) \subseteq E_a(u, X) \cap E_a(u', X)$. Recalling $(1-q)^k = 1/2$, we have

$$\Pr_{X\sim\mathbb{Q}}(\tilde{E}_a(u, u', X)) = 1 - (1-q)^{\ell} = 1 - 2^{-\ell/k}$$

and

$$\begin{aligned}\Pr_{X\sim\mathbb{Q}}(E_a(u, X) \cap E_a(u', X)) &= (1 - 2^{-\ell/k}) + 2^{-\ell/k}(1 - 2^{-(k-\ell)/k})^2 \\ &= 1 - 2\cdot 2^{-\ell/k-(1-\ell/k)} + 2^{-\ell/k-2(1-\ell/k)} \\ &= 2^{\ell/k-2}.\end{aligned}$$

Note that $A(u, X) \cap A(u', X)$ implies that the sum of all degrees in $u \cap u'$ is at most ℓd, which means $\tilde{E}_a(u, u', X)$ holds for at most ℓd tests a. Thus,

$$\Pr_{X\sim\mathbb{Q}}(E(u, X) \cap E(u', X) \cap A(u, X) \cap A(u', X)) \le (2^{\ell/k-2})^M \Pr(\mathrm{Bin}(M, r) \le \ell d) \tag{9.11}$$

where r is the conditional probability

$$r := \Pr_{X\sim\mathbb{Q}}(\tilde{E}_a(u, u', X) \mid E_a(u, X) \cap E_a(u', X)) = \frac{1 - 2^{-\ell/k}}{2^{\ell/k-2}} = 4\cdot 2^{-\ell/k}(1 - 2^{-\ell/k}).$$

We will treat the contributions to (9.7) from small ℓ and large ℓ separately.

Small ℓ. First consider the terms in (9.7) where $\ell \le \varepsilon k$ for a small constant $\varepsilon > 0$ to be chosen later. We need to bound the expression $\Pr(\mathrm{Bin}(M, r) \le \ell d)$ from (9.11). To this end, we have[4]

$$2^{-\ell/k} = \exp\left(-\frac{\ell}{k}\ln 2\right) = 1 - \frac{\ell}{k}\ln 2 + O((\ell/k)^2),$$

[4]Here and in the remainder of this section, we use $O(\cdot)$ with the understanding that its argument is small. Formally, $O(\cdot)$ hides an absolute constant factor provided that its argument is smaller than some absolute constant, and may also hide $1 + o(1)$ factors (in the usual sense).

$$r = 4(1 - O(\ell/k)) \left(\frac{\ell}{k} \ln 2 - O((\ell/k)^2) \right) = (1 - O(\varepsilon)) \cdot 4 \ln 2 \cdot \frac{\ell}{k} = (1 - O(\varepsilon)) \cdot 4\ell q,$$

$$\ell d = 2t\ell q M,$$

and

$$\mathbb{E}[\mathrm{Bin}(M, r)] = rM = (1 - O(\varepsilon)) \cdot 4\ell q M.$$

Note that if $t \geq 2$ then $\{\mathrm{Bin}(M, r) \leq \ell d\}$ is not a rare event and so we will simply upper-bound its probability by 1; in this case, we do not gain anything from using the conditional planted distribution $\tilde{\mathbb{P}}$ instead of $\mathbb{P}$. On the other hand, if $t < 2$ then we can apply the Binomial tail bound (Proposition A.2): writing $r = 4t'\ell q$ where $t' = 1 - O(\varepsilon)$, and taking ε small enough so that $t < 2t'$,

$$\Pr(\mathrm{Bin}(M, r) \leq \ell d) \leq \exp\left(-MD\left(\frac{\ell d}{M} \middle\| r \right) \right) = \exp\left(-MD(2t\ell q \,\|\, 4t'\ell q) \right)$$

where (using Lemma A.4)

$$D(2t\ell q \,\|\, 4t'\ell q) \geq 2\ell q(t \ln \frac{t}{2t'} + 2t' - t) - O((\ell q)^2).$$

This means

$$\begin{aligned} \Pr(\mathrm{Bin}(M, r) \leq \ell d) &\leq \exp\left(-2M\ell q(t \ln \frac{t}{2t'} + 2t' - t) + M\ell q \cdot O(\varepsilon) \right) \\ &= \exp\left(-2M\ell q \left(t \ln \frac{t}{2} + 2 - t - O(\varepsilon) \right) \right) \\ &= n^{-\ell\left[(1-\theta)c(\ln 2)\left(t \ln \frac{t}{2} + 2 - t\right) - O(\varepsilon)\right]}. \end{aligned} \tag{9.12}$$

We can now put everything together to bound the chi-squared divergence: using (9.11) and $\mathbb{P}(A) = 1 - o(1)$, the contribution to (9.7) from $\ell \leq \varepsilon k$ is at most

$$\begin{aligned} \mathcal{T}_{\leq \varepsilon k}(t) :&= \mathop{\mathbb{E}}_{u,u' \sim \mathcal{U}} \mathbb{1}_{\langle u, u' \rangle \leq \varepsilon k} \langle L_u, L_{u'} \rangle \\ &= \mathbb{P}(A)^{-2} \, 2^{2M} \sum_{0 \leq \ell \leq \varepsilon k} \Pr(\ell) \, (2^{\ell/k - 2})^M \Pr(\mathrm{Bin}(M, r) \leq \ell d) \\ &= \mathbb{P}(A)^{-2} \sum_{0 \leq \ell \leq \varepsilon k} \Pr(\ell) \, (2^{\ell/k})^M \Pr(\mathrm{Bin}(M, r) \leq \ell d) \\ &\leq \mathbb{P}(A)^{-2} \left[1 + \sum_{1 \leq \ell \leq \varepsilon k} \Pr(\ell) \, (2^{\ell/k})^M \Pr(\mathrm{Bin}(M, r) \leq \ell d) \right]. \end{aligned} \tag{9.13}$$

Note that we have made the dependence of $\mathcal{T}_{\leq \varepsilon k}(t)$ on t explicit; recall that t is a constant appearing in the definition of $\tilde{\mathbb{P}}$. Using

$$2^{M/k} = 2^{(c/2 + o(1)) \ln(n/k)} = \left(\frac{n}{k} \right)^{\frac{c}{2} \ln 2 + o(1)} = n^{(1-\theta)\frac{c}{2} \ln 2 + o(1)}$$

along with (9.8),(9.9),(9.10),(9.12)(9.13), we have

$$\mathcal{T}_{\le \varepsilon k}(t) \le \mathbb{P}(A)^{-2}\left[1+\sum_{1\le \ell\le \varepsilon k}\left(\frac{ek}{\ell}\right)^{\ell} n^{-\ell[(1-\theta)(1-\frac{c}{2}\ln 2)+o(1)]}\right. \tag{9.14}$$

$$\left. n^{\ell[(1-\theta)\frac{c}{2}\ln 2+o(1)]} n^{-\ell\left[(1-\theta)c(\ln 2)\left(t\ln\frac{t}{2}+2-t\right)-O(\varepsilon)\right]}\right]$$

$$= \mathbb{P}(A)^{-2}\left[1+\sum_{1\le \ell\le \varepsilon k}\left(\frac{e}{\ell} n^{\theta-(1-\theta)[1+c(\ln 2)(t\ln\frac{t}{2}-t+1)]+O(\varepsilon)}\right)^{\ell}\right]. \tag{9.15}$$

This is $1+o(1)$ for sufficiently small ε provided that the following three conditions hold:

(i) $t>1$ and $c(\ln 2)(t\ln t - t+1) > \frac{\theta}{1-\theta}$ so that $\mathbb{P}(A) = 1-o(1)$; see (9.5),

(ii) $t<2$ so that the bound (9.12) is valid,

(iii) $\theta-(1-\theta)[1+c(\ln 2)(t\ln\frac{t}{2}-t+1)]<0$ so that (9.15) is $1+o(1)$.

Provided $\frac{1}{2(\ln 2)^2} < c < \frac{1}{(\ln 2)^2}$ and $c\ln 2 - \frac{1}{\ln 2}[1+\ln(c(\ln 2)^2)] > \frac{\theta}{1-\theta}$, the choice $t = \frac{1}{c(\ln 2)^2}$ satisfies (i),(ii),(iii) above. This means we have proved the following.

Lemma 9.3. *For any fixed $\theta\in(0,1)$ and $c\in\left(\frac{1}{2(\ln 2)^2},\frac{1}{(\ln 2)^2}\right)$ satisfying*

$$c\ln 2 - \frac{1}{\ln 2}[1+\ln(c(\ln 2)^2)] > \frac{\theta}{1-\theta},$$

there exist constants $\varepsilon>0$ and $t>1$ such that $\mathbb{P}(A)=1-o(1)$ and $\mathcal{T}_{\le\varepsilon k}(t)=1+o(1)$.

Alternatively, we can drop the requirement (ii) $t<2$ and replace (9.12) with the trivial bound $\Pr(\mathrm{Bin}(M,r)\le \ell d)\le 1$ (which reverts to the non-conditional chi-squared). In this case the result is, similarly to (9.15),

$$\mathcal{T}_{\le \varepsilon k}(t) \le \mathbb{P}(A)^{-2}\left[1+\sum_{1\le \ell\le \varepsilon k}\left(\frac{ek}{\ell}\right)^{\ell} n^{-\ell[(1-\theta)(1-\frac{c}{2}\ln 2)+o(1)]} n^{\ell[(1-\theta)\frac{c}{2}\ln 2+o(1)]}\right]$$

$$= \mathbb{P}(A)^{-2}\left[1+\sum_{1\le \ell\le \varepsilon k}\left(\frac{e}{\ell} n^{\theta-(1-\theta)(1-c\ln 2)+o(1)}\right)^{\ell}\right]. \tag{9.16}$$

This is $1+o(1)$ for any $\varepsilon\in(0,1]$ (we have not required ε to be small in this case) provided that the following two conditions hold:

(i) $t>1$ and $c(\ln 2)(t\ln t - t+1) > \frac{\theta}{1-\theta}$ so that $\mathbb{P}(A) = 1-o(1)$; see (9.5),

(ii) $\theta-(1-\theta)(1-c\ln 2)<0$ so that (9.16) is $1+o(1)$.

We can satisfy (i) by choosing $t=\infty$ (i.e., $\tilde{\mathbb{P}}=\mathbb{P}$), so we are left with the condition (ii), which simplifies to $1-c\ln 2 > \frac{\theta}{1-\theta}$. This means we have proved the following.

Lemma 9.4. *For any fixed $\theta\in(0,1)$ and $c>0$ satisfying*

$$1-c\ln 2 > \frac{\theta}{1-\theta},$$

and for any $\varepsilon\in(0,1]$, we have $\mathcal{T}_{\le\varepsilon k}(\infty)=1+o(1)$.

Large ℓ. Now consider the contribution to (9.7) from $\varepsilon k \le \ell \le k$ for any fixed constant $\varepsilon > 0$. Use the trivial bound instead of (9.12); the conditioning will not be important here. Similarly to (9.16), the contribution is at most

$$\begin{aligned}\mathcal{T}_{>\varepsilon k}(t) :&= \mathop{\mathbb{E}}_{u,u'\sim\mathcal{U}} \mathbb{1}_{\langle u,u'\rangle>\varepsilon k} \langle L_u, L_{u'}\rangle \\ &= \mathbb{P}(A)^{-2} \sum_{\varepsilon k<\ell\le k} \left(\frac{ek}{\ell}\right)^{\ell} n^{-\ell[(1-\theta)(1-\frac{c}{2}\ln 2)+o(1)]} n^{\ell[(1-\theta)\frac{c}{2}\ln 2+o(1)]} \\ &\le (1+o(1)) \sum_{\varepsilon k<\ell\le k} \left(\frac{e}{\varepsilon} n^{-(1-\theta)(1-c\ln 2)+o(1)}\right)^{\ell},\end{aligned}$$

which is $o(1)$ provided $c < \frac{1}{\ln 2}$. This means we have proved the following.

Lemma 9.5. *For any constants $\theta \in (0,1)$, $c \in \left(0, \frac{1}{\ln 2}\right)$, $\varepsilon > 0$, and $t > 1$, we have $\mathcal{T}_{>\varepsilon k}(t) = o(1)$.*

9.2.3 Impossibility of Detection: Proof of Theorem 3.4(b)

Proof of Theorem 3.4(b). Recalling Lemma 7.1 and the reformulation in (9.1), our goal is to show $\chi^2(\tilde{\mathbb{P}} \,\|\, \mathbb{Q}) = o(1)$ provided $c < 1/\ln 2$ and $\tau(c) > \frac{\theta}{1-\theta}$. Recall $\chi^2(\tilde{\mathbb{P}} \,\|\, \mathbb{Q}) + 1 = \mathcal{T}_{\le\varepsilon k}(t) + \mathcal{T}_{>\varepsilon k}(t)$. For $\frac{1}{2(\ln 2)^2} < c < \frac{1}{\ln 2} < \frac{1}{(\ln 2)^2}$, the result follows from Lemmas 9.3 and 9.5. For $0 < c \le \frac{1}{2(\ln 2)^2}$, the result follows from Lemma 9.4 with $\varepsilon = 1$. □

9.2.4 Low-Degree Hardness of Detection: Proof of Theorem 3.3(b)

Proof of Theorem 3.3(b). Recalling Lemma 7.3 and the reformulation in (9.1), our goal is to show $\chi^2_{\le D}(\tilde{\mathbb{P}} \,\|\, \mathbb{Q}) = o(1)$ provided $\tau(c) > \frac{\theta}{1-\theta}$. Note that from (9.1), the assumption $\tau(c) > \frac{\theta}{1-\theta}$ implies $c < 1/(\ln 2)^2$, so we can assume this throughout this section. We will follow the proof outline explained in Section 7.4. We need an orthonormal basis of polynomials for $\mathbb{Q}$. Such a basis is given by $\{h_S\}_{S\subseteq[N]\times[M]}$ where $h_S(X) = [q(1-q)]^{-|S|/2} \prod_{(i,a)\in S}(X_{ia} - q)$. These are orthonormal with respect to the inner product $\langle\cdot,\cdot\rangle_{\mathbb{Q}}$. Furthermore, $\{h_S\}_{|S|\le D}$ is a basis for the subspace consisting of polynomials of degree (at most) D.

Following Section 7.4, define $\mathcal{U}$, $\tilde{\mathbb{P}}_u$, and $L_u = d\tilde{\mathbb{P}}_u/d\mathbb{Q}$, and recall the decomposition

$$\chi^2_{\le D}(\tilde{\mathbb{P}} \,\|\, \mathbb{Q}) + 1 = \mathcal{R}_{\le\varepsilon k}(t, D) + \mathcal{R}_{>\varepsilon k}(t, D),$$

where we have made explicit the dependence on t (the constant appearing in the definition of $\tilde{\mathbb{P}}$) and D. The following key fact is proved later in this section.

Lemma 9.6. *For any u, u', we have $\langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}} \le \langle L_u, L_{u'}\rangle_{\mathbb{Q}}$.*

In light of Lemma 9.6, we have

$$\begin{aligned}\mathcal{R}_{\le\varepsilon k}(t, D) &:= \mathop{\mathbb{E}}_{u,u'\sim\mathcal{U}} \mathbb{1}_{\langle u,u'\rangle\le\varepsilon k} \langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}} \\ &\le \mathop{\mathbb{E}}_{u,u'\sim\mathcal{U}} \mathbb{1}_{\langle u,u'\rangle\le\varepsilon k} \langle L_u, L_{u'}\rangle_{\mathbb{Q}} =: \mathcal{T}_{\le\varepsilon k}(t),\end{aligned}$$

and we have already shown $\mathcal{T}_{\le\varepsilon k}(t) = 1 + o(1)$ (Lemmas 9.3, 9.4) under the assumption $\tau(c) > \frac{\theta}{1-\theta}$. The other term $\mathcal{R}_{>\varepsilon k}(t, D)$ can be controlled by the following lemma, proved later in this section. (Recall we are assuming $c < \frac{1}{(\ln 2)^2} < \frac{2}{\ln 2}$ in this section.)

Lemma 9.7. *For any constants $\theta \in (0,1)$, $c \in \left(0, \frac{2}{\ln 2}\right)$, $\varepsilon > 0$, and $t > 1$, and for any $D = D_n$ satisfying $D = o(k)$, we have $\mathcal{R}_{>\varepsilon k}(t, D) = o(1)$.*

This completes the proof of the theorem, modulo the two lemmas that remain to be proved below. □

Proof of Lemma 9.6. We use a symmetry argument from [BEH+22, Proposition 3.6]. Expanding in the orthonormal basis $\{h_S\}$, we have

$$\langle L_u^{\le D}, L_{u'}^{\le D}\rangle_{\mathbb{Q}} = \sum_{|S|\le D} \langle L_u, h_S\rangle_{\mathbb{Q}} \langle L_{u'}, h_S\rangle_{\mathbb{Q}} = \sum_{|S|\le D} \mathop{\mathbb{E}}_{X\sim\tilde{\mathbb{P}}_u}[h_S(X)] \mathop{\mathbb{E}}_{X\sim\tilde{\mathbb{P}}_{u'}}[h_S(X)]. \tag{9.17}$$

Let $V(S) = \{i \in [N] : \exists a \in [M], (i,a) \in S\}$, the set of all individuals "involved" in the basis function S. Note that if $V(S) \not\subseteq u$ then there exists some $(i,a) \in S$ such that under $X \sim \tilde{\mathbb{P}}_u$ we have $X_{ia} \sim \text{Bernoulli}(q)$ independently from the rest of X, and thus $\mathbb{E}_{X\sim\tilde{\mathbb{P}}_u}[h_S(X)] = 0$. (Here it is important that conditioning on the event A only affects infected individuals.) Similarly, if $V(S) \not\subseteq u'$ then $\mathbb{E}_{X\sim\tilde{\mathbb{P}}_{u'}}[h_S(X)] = 0$. On the other hand, if $V(S) \subseteq u \cap u'$ then (by symmetry) $\tilde{\mathbb{P}}_u$ and $\tilde{\mathbb{P}}_{u'}$ have the same marginal distribution when restricted to the variables $\{(i,a) : i \in u \cap u'\}$ and so $\mathbb{E}_{X\sim\tilde{\mathbb{P}}_u}[h_S(X)] = \mathbb{E}_{X\sim\tilde{\mathbb{P}}_{u'}}[h_S(X)]$. As a result, we have $\mathbb{E}_{X\sim\tilde{\mathbb{P}}_u}[h_S(X)]\,\mathbb{E}_{X\sim\tilde{\mathbb{P}}_{u'}}[h_S(X)] \ge 0$ for all S, i.e., every term on the right-hand side of (9.17) is nonnegative. This means $\langle L_u^{\le 0}, L_{u'}^{\le 0}\rangle_{\mathbb{Q}} \le \langle L_u^{\le 1}, L_{u'}^{\le 1}\rangle_{\mathbb{Q}} \le \langle L_u^{\le 2}, L_{u'}^{\le 2}\rangle_{\mathbb{Q}} \le \cdots \le \langle L_u^{\le \infty}, L_{u'}^{\le \infty}\rangle_{\mathbb{Q}} = \langle L_u, L_{u'}\rangle_{\mathbb{Q}}$. □

Proof of Lemma 9.7. For any S and X, we have the bound $|h_S(X)| \le \left(\frac{1-q}{q}\right)^{|S|/2} \le q^{-|S|/2}$ (assuming $q \le 1/2$, which holds for sufficiently large n). Expanding $\mathcal{R}_{>\varepsilon k}(t, D)$ using (9.17), and using the fact that the number of subsets $S \subseteq [N] \times [M]$ of size $|S| \le D$ is at most $(NM+1)^D$,

$$\begin{aligned}
\mathcal{R}_{>\varepsilon k}(t, D) &= \mathop{\mathbb{E}}_{u,u'} \mathbb{1}_{\langle u,u'\rangle > \varepsilon k} \sum_{|S|\le D} \mathop{\mathbb{E}}_{X\sim\tilde{\mathbb{P}}_u}[h_S(X)] \mathop{\mathbb{E}}_{X\sim\tilde{\mathbb{P}}_{u'}}[h_S(X)] \\
&\le \mathop{\mathbb{E}}_{u,u'} \mathbb{1}_{\langle u,u'\rangle > \varepsilon k} \sum_{|S|\le D} q^{-|S|} \\
&\le \Pr_{u,u'}(\langle u,u'\rangle > \varepsilon k)\,(NM+1)^D q^{-D}.
\end{aligned}$$

Similarly to (8.10),

$$\begin{aligned}
\Pr_{u,u'}(\langle u,u'\rangle > \varepsilon k) \le \binom{k}{\lceil \varepsilon k\rceil} \frac{\binom{N-\lceil\varepsilon k\rceil}{k-\lceil\varepsilon k\rceil}}{\binom{N}{k}} &\le \binom{k}{\lceil\varepsilon k\rceil} \left(\frac{k}{N-\lceil\varepsilon k\rceil+1}\right)^{\lceil\varepsilon k\rceil} \\
&\le \left(\frac{ek}{\lceil\varepsilon k\rceil}\right)^{\lceil\varepsilon k\rceil} \left(\frac{k}{N-k}\right)^{\lceil\varepsilon k\rceil} = n^{-\Omega(k)}
\end{aligned}$$

provided $c < \frac{2}{\ln 2}$ (so that $k = o(N)$). Also,

$$(NM+1)^D q^{-D} = n^{O(D)}$$

and so

$$\mathcal{R}_{>\varepsilon k}(t, D) \le n^{-\Omega(k)} n^{O(D)}$$

which is $o(1)$ provided $D = o(k)$. □

A Tool Box

The following lemmas will be useful to us.

Lemma A.1 (Stirling approximation [Mar65])**.** *We have for $n \to \infty$ that*

$$n! = (1 + O(1/n))\sqrt{2\pi n}\, n^n \exp(-n).$$

We will use the following standard Binomial tail bound.

Proposition A.2 ([AG89])**.** *Let $n \in \mathbb{N}$ and $p \in (0,1)$. For $a \in (0,1)$, define*

$$D(a \,\|\, p) := a \ln \frac{a}{p} + (1-a) \ln \frac{1-a}{1-p}. \tag{A.1}$$

- *For all $0 < k < pn$,*

$$\Pr(\mathrm{Bin}(n,p) \le k) \le \exp\left(-nD\left(\frac{k}{n} \,\middle\|\, p\right)\right).$$

- *For all $pn < k < n$,*

$$\Pr(\mathrm{Bin}(n,p) \ge k) \le \exp\left(-nD\left(\frac{k}{n} \,\middle\|\, p\right)\right).$$

There is also a nearly-matching *lower bound* on the tail probability.

Proposition A.3 ([Ash90])**.** *Let $n \in \mathbb{N}$ and $p \in (0,1)$. Define $D(a \,\|\, p)$ as in (A.1).*

- *For all $0 < k < pn$,*

$$\Pr(\mathrm{Bin}(n,p) \le k) \ge \frac{1}{\sqrt{8k(1-k/n)}} \exp\left(-nD\left(\frac{k}{n} \,\middle\|\, p\right)\right).$$

- *For all $pn < k < n$,*

$$\Pr(\mathrm{Bin}(n,p) \ge k) \ge \frac{1}{\sqrt{8k(1-k/n)}} \exp\left(-nD\left(\frac{k}{n} \,\middle\|\, p\right)\right).$$

The following bounds on $D(a \,\|\, p)$ will be convenient.

Lemma A.4. *Suppose $a, p \in (0, \delta]$ for some $\delta \in (0, 1/2]$. Then*

$$a \ln \frac{a}{p} + p - a - 3\delta^2 \le D(a \,\|\, p) \le a \ln \frac{a}{p} + p - a + 3\delta^2.$$

Proof. For the first inequality, bound the second term in the definition (A.1) as follows:

$$\begin{aligned}(1-a)\ln\frac{1-a}{1-p} &\geq (1-a)\ln[(1-a)(1+p)]\\ &= (1-a)\ln(1+p-a-ap).\end{aligned}$$

Note that $1-\delta \leq (1-a)(1+p) \leq 1+\delta$ and so $-\delta \leq p-a-ap \leq \delta$. Taylor-expand the logarithm:

$$\begin{aligned}&= (1-a)\sum_{k=1}^{\infty}\frac{(-1)^{k+1}}{k}(p-a-ap)^k\\ &\geq (1-a)\left(p-a-ap-\frac{1}{2}\sum_{k=2}^{\infty}\delta^k\right)\\ &\geq (1-a)\left(p-a-2\delta^2\right)\\ &= p-a-2\delta^2-ap+a^2+2a\delta^2\\ &\geq p-a-3\delta^2\end{aligned}$$

as desired.

Now, for the second inequality,

$$\ln\frac{1-a}{1-p} = \ln(1-a)+\ln(1+p+p^2+p^3+\cdots) \leq \ln(1-a)+\ln(1+p+2p^2) \leq p-a+2p^2$$

where we have used $p \leq 1/2$ and $\ln(1+x) \leq x$. This means

$$(1-a)\ln\frac{1-a}{1-p} \leq p-a+2p^2-ap+a^2-2ap^2 \leq p-a+2p^2+a^2 \leq p-a+3\delta^2$$

as desired. □

B Orthogonal Polynomials

In this section we give more details about the orthogonal polynomials on a slice of the hypercube. In particular, we explain how to deduce the claims in Section 8.2.1 from the results of [Fil16] (definition/theorem numbers for [Fil16] pertain to arXiv v2).

Throughout this section, the inner product and norm for functions are with respect to the uniform distribution on the slice $\binom{[M]}{\Delta}$, as defined in Section 8.2.1. The basis elements are $\hat{\chi}_B := \chi_B/\|\chi_B\|$ where χ_B is defined in [Fil16, Definition 3.2]. The indices B are elements of a particular set $\mathcal{B}_M$; each $B \in \mathcal{B}_M$ is a strictly increasing sequence of elements from $[M]$, whose length we denote $|B|$. The set $\mathcal{B}_M$ does not contain all such sequences, only those that are "top sets" [Fil16, Definition 2.3] but the details of this will not be important for us. The functions χ_B (and therefore also $\hat{\chi}_B$) are orthogonal; see Theorems 3.1 and 4.1 of [Fil16].

For convenience, we recap the definition of χ_B from [Fil16]. For sequences $A = a_1, \ldots, a_d$ and $B = b_1, \ldots, b_d$ where $a_1, \ldots, a_d, b_1, \ldots, b_d$ are $2d$ distinct numbers from $[M]$, define

$$\chi_{A,B} = \prod_{i=1}^{d}(x_{a_i} - x_{b_i})$$

as in [Fil16, Definition 2.2]. Now following [Fil16, Definition 3.2], define

$$\chi_B = \sum_{A<B} \chi_{A,B}$$

where the sum over $A < B$ is over sequences $A = a_1, \ldots, a_d$ of length $d = |B|$, whose elements are distinct and disjoint from those of B, with $a_i < b_i$ entrywise.

Proof of Fact 8.9. The basis elements $\hat{\chi}_B = \chi_B / \|\chi_B\|$ have norm 1 by construction. By [Fil16, Theorem 4.1], the set $\{\chi_B : B \in \mathcal{B}_M, |B| \le \Delta\}$ is a complete orthogonal basis (as a vector space over $\mathbb{R}$) for all functions $\binom{[M]}{\Delta} \to \mathbb{R}$. This means for any degree-D polynomial $f : \mathbb{R}^M \to \mathbb{R}$, there is a unique collection of coefficients $\alpha_B \in \mathbb{R}$ such that the linear combination

$$\sum_{\substack{B \in \mathcal{B}_M \\ |B| \le \Delta}} \alpha_B \hat{\chi}_B$$

is equivalent to f on $\binom{[M]}{\Delta}$. It remains to show that this expansion only uses basis functions with $|B| \le D$, that is, we aim to show $\alpha_B = 0$ for all $|B| > D$. Since $\alpha_B = \langle f, \hat{\chi}_B \rangle$, this follows from Lemma B.1 below. □

Lemma B.1. *If $f : \mathbb{R}^M \to \mathbb{R}$ is a degree-D polynomial and $|B| > D$ then $\langle f, \chi_B \rangle = 0$.*

Proof. By linearity, it suffices to prove $\langle f, \chi_{A,B} \rangle = 0$ for an arbitrary $A < B$ in the case where f is a single degree-D *monomial*. Since f involves only D different variables and $|B| > D$, there must be an index j such that both x_{a_j} and x_{b_j} do not appear in f. Now write

$$\langle f, \chi_{A,B} \rangle = \mathop{\mathbb{E}}_{x \sim \mathrm{Unif}\binom{[M]}{\Delta}} \left(f(x) \prod_{i \ne j} (x_{a_i} - x_{b_i}) \right) (x_{a_j} - x_{b_j}),$$

which is equal to zero by symmetry, since for any fixed values for $\{x_i : i \ne j\}$, the events $\{x_{a_j} = 0, x_{b_j} = 1\}$ and $\{x_{a_j} = 1, x_{b_j} = 0\}$ are equally likely. □

We now prove Fact 8.10, which recall is the claim $|\hat{\chi}_B(x)| \le M^{2|B|}$ for all $x \in \binom{[M]}{\Delta}$ and all $B \in \mathcal{B}_M$ with $|B| \le \Delta$.

Proof of Fact 8.10. Since $\hat{\chi}_B = \chi_B / \|\chi_B\|$, the claim follows immediately from Lemmas B.2 and B.3 below. □

Lemma B.2. *For any $x \in \binom{[M]}{\Delta}$ and any $B \in \mathcal{B}_M$ with $|B| \le \Delta$, we have $|\chi_B(x)| \le M^{|B|}$.*

Proof. There are at most $M^{|B|}$ length-$|B|$ sequences of elements from $[M]$. Therefore, χ_B is the sum of at most $M^{|B|}$ terms $\chi_{A,B}$, and each $\chi_{A,B}$ can only take values in $\{-1, 0, 1\}$. □

Lemma B.3. *For any $B \in \mathcal{B}_M$ with $|B| \le \Delta$, we have $\|\chi_B\| \ge M^{-|B|}$.*

Proof. Let $d = |B|$. Theorem 4.1 of [Fil16] states that

$$\|\chi_B\|^2 = c_B 2^d \frac{\Delta^{\underline{d}}(M-\Delta)^{\underline{d}}}{M^{\underline{2d}}}$$

where $n^{\underline{k}} := n(n-1)\cdots(n-k+1)$ and (see [Fil16], Theorem 3.2)

$$c_B := \prod_{i=1}^{d} \binom{b_i - 2(i-1)}{2}. \tag{B.1}$$

We know that $c_B > 0$ because $\|\chi_B\|^2 > 0$ for all $B \in \mathcal{B}_M$ with $|B| \leq \Delta$ (see the proof of Theorem 4.1 in [Fil16]), and from (B.1) it is clear that c_B is an integer. This means $c_B \geq 1$. We now have

$$\|\chi_B\|^2 \geq \frac{1}{M^{\underline{2d}}} \geq M^{-2d}$$

as desired. □

C Reducing Detection to Approximate Recovery

In this section we show that any algorithm for approximate recovery can be made into an algorithm for strong detection, in both the Bernoulli (Proposition C.1) and constant-column (Proposition C.2) designs. We first focus on the Bernoulli design after the pre-processing step of COMP as discussed in Section 2.1.

Proposition C.1. *Assume the Bernoulli design for group testing with $c > 1/\ln 2$ and any $\theta \in (0,1)$. If an algorithm A defined on $N \times M$ bipartite graphs with worst-case termination time $T(A)$ achieves approximate recovery, then there is an algorithm B that achieves strong detection with worst-case termination time at most $T(A) + \mathrm{poly}(N, M)$.*

Recall that $c > 1/\ln 2$ is the condition for information-theoretic possibility of approximate recovery.

Proof. We choose $\delta > 0$ such that $cD(\delta \,\|\, 2^{-(1+\delta)})/(1+\delta) > 1$, where D is defined according to (A.1). Notice that such a $\delta > 0$ exists since $c > 1/\ln 2$.

The algorithm B acts as follows: it first runs A on the group testing instance and then checks if the output of A is a set of size at most $(1+\delta)k$ that explains all but δM of the (positive) tests. If YES, output that the distribution is planted. If NO, output that the distribution is the null. The termination time is immediate. We proceed with the analysis.

Success on the null model In this case, we will show the stronger result that with probability $1 - o(1)$, there is not a set of size at most $(1+\delta)k$ individuals which explains all but δM of the tests.

First notice that for a size-ℓ set of individuals, the number of tests they don't explain is distributed as $\mathrm{Bin}(M, (1-\nu/k)^\ell = 2^{-\ell/k})$. Hence, by a direct union bound the probability that there is a set of individuals of size $(1+\delta)k$ which satisfies all but δM of the tests is at most

$$\begin{aligned}
&\sum_{0\le\ell\le(1+\delta)k}\binom{N}{\ell}\Pr[\mathrm{Bin}(M,2^{-\ell/k})\le\delta M]\\
&\qquad\le k\binom{N}{(1+\delta)k}\Pr[\mathrm{Bin}(M,2^{-1-\delta})\le\delta M]\\
&\qquad\le k\exp[(1+\delta)k\ln(N/k)-D(\delta\,\|\,2^{-1-\delta})M]\\
&\qquad= k\exp[(1+\delta-cD(\delta\,\|\,2^{-1-\delta}))k\ln(N/k)]\\
&\qquad= o(1).
\end{aligned}$$

Success on the planted model Choose an arbitrary fixed $\delta' \in (0, \frac{\delta}{2\ln 2})$. Note the success of A in approximate recovery immediately implies that with probability $1-o(1)$, the size of A's output is at most $(1+\delta')k$ individuals and among these there are at least $(1-\delta')k$ infected individuals.

Given the above, we have the following: the probability that A's output explains fewer than $(1-\delta)M$ tests is, up to a $o(1)$ additive factor, at most the probability that there exists a subset of at most $\delta' k$ infected individuals with at least one participant in at least δM tests. This by a union bound and Proposition A.2 (since $\delta'\nu < \delta$ for large values of N) is at most

$$\begin{aligned}
\binom{k}{\delta' k}\Pr[\mathrm{Bin}(\delta' Mk,\nu/k)\ge\delta M] &\le \exp(-\delta' MkD(1/k\,\|\,\nu/k)+O(k))\\
&= \exp(-\Omega(M)+O(k))\\
&= o(1).
\end{aligned}$$

This completes the proof. □

We now prove the analogous result for the constant-column design.

Proposition C.2. *Assume the constant-column design for group testing with $c > 1/\ln 2$ and any $\theta \in (0,1)$. If an algorithm A defined on $N \times M$ bipartite graphs with worst-case termination time $T(A)$ achieves approximate recovery, then there is an algorithm B that achieves strong detection with worst-case termination time at most $T(A) + \mathrm{poly}(N, M)$.*

Proof. This proof follows along the lines of the Bernoulli case but it becomes a little bit easier. Intuitively, this is clear: the probability that a set of ℓ individuals is connected to all tests is comparable in the two designs but in the Bernoulli design the individual degrees fluctuate significantly.

Let $\eta > \frac{1}{2c\ln^2 2}$. The decision algorithm B reads as follows:

- Check the outcome of algorithm A.
  - If the outcome is a set of at most $(1+\eta)k$ individuals that are connected to at least $(1-\eta)M$ tests, return *planted*.
  - Otherwise, return *null*.
- This checking works in polynomial time.

Success on the planted model Let $0 < \delta < \frac{\eta}{2\ln 2}$. The algorithm A returns by assumption a set of at most $(1+\delta)k$ individuals, out of which at least $(1-\delta)k$ are truly infected, with probability $1 - o(1)$. As the model is a planted model, we know that there are at most δk additional infected individuals that can be used to explain the tests. Those δk individuals can be connected to at most

$$\delta k \Delta = \frac{\delta M}{2\ln 2} < \eta M$$

tests by construction. Therefore, the output of B is correct with probability $1 - o(1)$.

Success on the null model It suffices to prove that in a random almost regular graph with N individual nodes, M test-nodes and individual degree Δ, there is with high probability no set of at most $(1+\eta)k$ individuals that is connected to at least $(1-\eta)M$ tests.

We employ the balls-into-bins experiment. (We ignore the issue of multi-edges here, as this can be handled similarly to Section 6.3.1.) If $\ell\Delta$ balls are thrown onto $M = \frac{k\Delta}{2\ln 2}$ boxes, the expected number of empty boxes $\boldsymbol{A}_\ell$ is

$$\mathbb{E}\left[\boldsymbol{A}_\ell\right] = \ell\Delta\left(1 - \frac{1}{\ell\Delta}\right)^{\frac{k\Delta}{2\ln 2}}.$$

Let $p_\ell = \left(1 - \frac{1}{\ell\Delta}\right)^{\frac{k\Delta}{2\ln 2}}$. It is a well known fact that the indicator functions for the different boxes being empty are negatively associated Bernoulli random variables [DR96]. Therefore, the Chernoff bound implies

$$\Pr\left(\boldsymbol{A}_\ell \le p_\ell\ell\Delta - t\ell\Delta\right) \le \exp\left(-\ell D_{\mathrm{KL}}(p_\ell - t \,\|\, p_\ell)\right).$$

Therefore, the probability that a set of individuals of size at most $(1+\eta)k$ exists that explains all but ηM tests is upper bounded by

$$\sum_{\ell=0}^{(1+\eta)k} \binom{N}{\ell} \Pr\left(\boldsymbol{A}_\ell \le \eta M\right) \le (1+\eta)k \binom{N}{(1+\eta)k} \Pr\left(\boldsymbol{A}_{(1+\eta)k} \le \eta M\right).$$

The calculus is now identical to the Bernoulli case. □

D Comparison with [TAS20]

The detection boundary in Bernoulli group testing was studied by [TAS20], in a model similar to ours but with a slight difference. In the present work, we study detection in the Bernoulli design in the "post-COMP" setting discussed in Section 2. We repeat here the setting for convenience.

"Post-COMP" Bernoulli design (testing) Let n, $k = k_n$, $N = N_n$ and $M = M_n$ scale as $k = n^{\theta+o(1)}$, $N = n^{1-(1-\theta)\frac{c}{2}\ln 2+o(1)}$ and $M = (c/2 + o(1))k\ln(n/k)$. Consider the following distributions over (N, M)-bipartite graphs (encoding adjacency between N individuals and M tests).

- Under the null distribution $\mathbb{Q}$, each of the N individuals participates in each of the M tests with probability $q = \nu/k$ with $\nu > 0$ such that $(1-\nu/k)^k = 1/2$ (defined also in Section 2) independently.
- Under the planted distribution $\mathbb{P}$, a set of k infected individuals out of N is chosen uniformly at random. Then a graph is drawn from $\mathbb{Q}$ conditioned on having at least one infected individual in every test.

As described in Theorem 3.4, we have established in this work *the exact detection boundary* for the above setting. Previously, [TAS20] provided upper and lower bounds for the detection boundary in the "pre-COMP" Bernoulli design, defined as follows.

"Pre-COMP" Bernoulli design (testing) Let n, $k = k_n$, $m = m_n$ scale as $k = n^{\theta+o(1)}$ and $m = (c+o(1))k\ln(n/k)$. Consider the following distributions over $(G, \hat{\sigma})$ pairs, where G is an (n, m)-bipartite graph (encoding adjacency between n individuals and m tests) and $\hat{\sigma} \in \{0,1\}^m$ encodes positive/negative test results.

- Under the null distribution $\mathbb{Q}$, each of the n individuals participates in each of the m tests with probability q (defined above) independently. The test results are chosen independently to be positive or negative with probability $1/2$.
- Under the planted distribution $\mathbb{P}$, a set of k infected individuals out of n is chosen uniformly at random. Then a graph is drawn from $\mathbb{Q}$. Finally, each test result is labelled positive if at least one infected individual participated in it. Otherwise, it is labelled negative.

In this section we provide a short proof that our Theorem 3.4 can be used to establish the detection boundary of the pre-COMP Bernoulli design as well. We prove the following result, in particular improving both the upper and lower bounds of [TAS20].

Theorem D.1. *Consider the pre-COMP Bernoulli design with parameters $\theta \in (0,1)$ and $c > 0$. Recall $c_{\inf} := 1/\ln 2$ and $c^{\mathrm{B}}_{\mathrm{LD}}$ as defined in (3.2).*

(a) (Possible) If $c > \min\{c_{\inf}, c^{\mathrm{B}}_{\mathrm{LD}}\}$ then strong detection is possible.

(b) (Impossible) If $c < \min\{c_{\inf}, c^{\mathrm{B}}_{\mathrm{LD}}\}$ then weak detection is impossible.

D.1 Proof of Theorem D.1

For the proof of Theorem D.1 we need a lemma which almost follows immediately from standard results.

Lemma D.2. *Assume the pre-COMP planted distribution $\mathbb{P}$ for the Bernoulli design. For all $\theta \in (0,1)$ and $c \in (0, 1/\ln 2)$ it holds that the number of post-COMP remaining individuals N and post-COMP remaining tests M are distributed as $M \sim \mathrm{Bin}(m, 1/2)$ and $N|M \sim k + \mathrm{Bin}(n-k, 2^{-(m-M)/k})$. In particular, it holds with probability $1-o(1)$ that*

$$M \in [m/2 - \sqrt{m\ln n},\ m/2 + \sqrt{m\ln n}]$$

and

$$N \in [n^{1-(1-\theta)\frac{c}{2}\ln 2 - \frac{1}{\sqrt{\ln n}}},\ n^{1-(1-\theta)\frac{c}{2}\ln 2 + \frac{1}{\sqrt{\ln n}}}].$$

Proof. The distribution of M follows directly. Now, given M, each non-infected individual is removed by COMP with probability $(1-\nu/k)^{m-M} = 2^{-(m-M)/k}$. The high-probability event follows directly from a multiplicative Chernoff bound and the fact $c < 1/\ln 2 < 2/\ln 2$. □

We start with the fairly intuitive direction, proving that any successful algorithm for strong detection in the post-COMP model also achieves strong detection in the pre-COMP model. In particular, given Theorem 3.4, we conclude that if $c > \min\{c_{\inf}, c_{\mathrm{LD}}^{\mathrm{B}}\}$ then strong detection is possible in the pre-COMP Bernoulli design.

Proposition D.3. *Fix parameters $\theta \in (0,1)$ and $c \in (0, 1/\ln 2)$. If strong detection is information-theoretically possible in the post-COMP Bernoulli design then it is also information-theoretically possible in the pre-COMP Bernoulli design.*

Proof. Consider any algorithm A achieving strong detection in the post-COMP Bernoulli design. Then we claim the following algorithm B achieves strong detection in the pre-COMP Bernoulli design: First run COMP on the received input. If the remaining number of tests M and the remaining number of individuals N do not both satisfy

$$M \in [m/2 - \sqrt{m \ln n},\ m/2 + \sqrt{m \ln n}]$$

and

$$N \in [n^{1-(1-\theta)\frac{c}{2}\ln 2 - \frac{1}{\sqrt{\ln n}}},\ n^{1-(1-\theta)\frac{c}{2}\ln 2 + \frac{1}{\sqrt{\ln n}}}]$$

then output that the distribution is $\mathbb{Q}$. Otherwise, run A on the post-COMP instance and return the output of A.

The analysis is as follows.

Planted model Assume that the algorithm receives input from the planted model. In that case, based on Lemma D.2, after running COMP the parameters M, N satisfy the desired constraints, with probability $1 - o(1)$. Hence, with probability $1 - o(1)$, the algorithm does not terminate in the second step. In the third step, the algorithm then receives an instance of the planted distribution based on the post-COMP Bernoulli design, where in particular the assumptions on M, N are satisfied. Hence, it outputs that the distribution is $\mathbb{P}$ with probability $1 - o(1)$, by assumption on the performance of A.

Null model Assume that the algorithm receives input from the null model. In that case, either the algorithm outputs that the distribution is $\mathbb{Q}$ in the second step (which is correct), or after COMP is applied to the group testing instance the output has $M = (c/2 + o(1))k\ln(n/k)$ remaining tests and $N = n^{1-(1-\theta)\frac{c}{2}\ln 2 + o(1)}$ remaining individuals. In that case, the output of the second step is an instance of the null distribution based on the post-COMP Bernoulli design satisfying the desired assumptions on N, M. Hence, it outputs that the distribution is $\mathbb{Q}$ with probability $1 - o(1)$, by assumption on the performance of A in the post-COMP model. The proof is complete. □

Finally, we also prove the following, perhaps less immediate, direction. In particular, given Theorem 3.4, this implies that if $c < \min\{c_{\inf}, c_{\mathrm{LD}}^{\mathrm{B}}\}$ then strong detection is impossible in the pre-COMP Bernoulli design.

Proposition D.4. *Fix parameters $\theta \in (0,1)$ and $c > 0$ with $c < \min\{c_{\inf}, c^{\mathrm{B}}_{\mathrm{LD}}\}$. If weak detection is impossible in the post-COMP Bernoulli design then it is also impossible in the pre-COMP Bernoulli design.*

Proof. Let us first decompose any pre-COMP Bernoulli group testing graph instance (produced by either the planted or null distribution), seen as a bipartite graph between n individuals and m tests into two edge-disjoint parts: the graph G_1 between the N post-COMP individuals and the M positive tests, and the graph G_2 between the $n - N$ (healthy) individuals that COMP deleted, and the m (both positive and negative) tests.

We first show that under our assumptions, the distribution over (N, M) produced by the planted (pre-COMP) model and the distribution over (N, M) produced by the null (pre-COMP) model have vanishing total variation distance. It is straightforward to see that in both models the distribution of M is $\mathrm{Bin}(m, 1/2)$. Hence, using Lemma D.2 it suffices to couple for $X := m - M \sim \mathrm{Bin}(m, 1/2)$, the distribution $N_P \sim k + \mathrm{Bin}(n-k, r = e^{-(\ln 2)X/k})|M$ (coming from the planted) and the distribution $N_Q \sim \mathrm{Bin}(n, r = e^{-(\ln 2)X/k})|M$ (coming from the null). By Pinsker's inequality it suffices to prove that the KL divergence vanishes. We have by elementary inequalities,

$$\begin{aligned}
D_{\mathrm{KL}}(N_P|M \,\|\, N_Q|M) &= \mathop{\mathbb{E}}_{s\sim N_P|M} \ln \frac{\Pr(N_P = s)}{\Pr(N_Q = s)} \\
&= \mathop{\mathbb{E}}_{s\sim N_P|M} \ln \frac{\binom{n-k}{s-k} r^{s-k}(1-r)^{n-s}}{\binom{n}{s} r^s (1-r)^{n-s}} \\
&= \mathop{\mathbb{E}}_{s\sim N_P|M} \ln \frac{s!(n-k)!}{(s-k)!n!} r^{-k} \\
&\le \mathop{\mathbb{E}}_{s\sim N_P|M} \ln \frac{s^k}{(n-k)^k r^k} \\
&= k \mathop{\mathbb{E}}_{s\sim N_P|M} \ln \frac{s}{(n-k)r} \\
&\le k \mathop{\mathbb{E}}_{s\sim N_P|M} \frac{s-(n-k)r}{(n-k)r} \\
&= k \mathop{\mathbb{E}}_{X\sim \mathrm{Bin}(m,1/2)} \frac{k+nr-(n-k)r}{(n-k)r} \\
&\le \frac{2k^2}{n} \mathop{\mathbb{E}}_{X\sim \mathrm{Bin}(m,1/2)} e^{(\ln 2)X/k}.
\end{aligned}$$

Now, using the MGF of a Binomial distribution,

$$\begin{aligned}
D_{\mathrm{KL}}(N_P|M \,\|\, N_Q|M) &\le \frac{2k^2}{n}((e^{\ln 2/k}+1)/2)^m \\
&= \frac{2k^2}{n}(1 + \ln 2/(2k) + O(1/k^2))^m \\
&= \frac{2k^2}{n} e^{m \ln 2/(2k) + O(m/k^2)} \\
&= n^{2\theta - 1 + c(\ln 2)(1-\theta)/2 + o(1)}.
\end{aligned}$$

We will next show that the assumption $c < \min\{c_{\inf}, c_{\mathrm{LD}}^{\mathrm{B}}\}$ implies $2\theta - 1 + c(\ln 2)(1-\theta)/2 < 0$, which means $D_{\mathrm{KL}}(N_P|M \,\|\, N_Q|M) = o(1)$ and so we can couple (M, N) under the planted and the null models with probability $1 - o(1)$.

Under our assumption $c < c_{\mathrm{LD}}^{\mathrm{B}}$ we have that equivalently for the function

$$\tau(c) = \begin{cases} 1 - c\ln 2 & \text{if } 0 < c \le \frac{1}{2(\ln 2)^2}, \\ c\ln 2 - \frac{1}{\ln 2}[1 + \ln(c(\ln 2)^2)] & \text{if } \frac{1}{2(\ln 2)^2} < c < \frac{1}{(\ln 2)^2}, \end{cases}$$

that it holds $\tau(c) > \frac{\theta}{1-\theta}$. But for all $1/\ln 2 > c > 0$, we have

$$\tau(c) < 1 - c\ln 2/2.$$

Indeed if $c < \frac{1}{2\ln^2 2}$ that is clear. Now it also holds $c\ln 2 - \frac{1}{\ln 2}[1 + \ln(c(\ln 2)^2)] < 1 - \frac{c\ln 2}{2}$ when $\frac{1}{2(\ln 2)^2} < c < \frac{1}{(\ln 2)^2}$. This follows as

$$F(c) := c\ln 2 - \frac{1}{\ln 2}[1 + \ln(c(\ln 2)^2)] - (1 - c\ln 2/2), \qquad \frac{1}{2(\ln 2)^2} < c < \frac{1}{(\ln 2)^2},$$

is a convex function on c which is negative in the endpoints: $F(\frac{1}{2(\ln 2)^2}) = -\frac{1}{4\ln 2} < 0$ and also $F(\frac{1}{(\ln 2)^2}) = \frac{1}{2\ln 2} - 1 < 0$.

Hence, we have indeed established $\frac{\theta}{1-\theta} < 1 - \frac{c\ln 2}{2}$ and therefore $2\theta - 1 + c\ln 2(1-\theta)/2 < 0$. In particular, $D_{\mathrm{KL}}(N_P|M \,\|\, N_Q|M) = o(1)$ and indeed we can couple (M, N) under the planted and the null model with probability $1 - o(1)$.

Now that we have coupled the planted and null distributions for (N, M), we will use this to couple the entire pre-COMP planted distribution with the pre-COMP null distribution with probability $1 - o(1)$, implying impossibility of pre-COMP weak detection.

Recall from Lemma D.2 that (N, M) satisfy

$$M \in [m/2 - \sqrt{m\ln n}, m/2 + \sqrt{m\ln n}]$$

and

$$N \in [n^{1-(1-\theta)\frac{c}{2}\ln 2 - \frac{1}{\sqrt{\ln n}}}, n^{1-(1-\theta)\frac{c}{2}\ln 2 + \frac{1}{\sqrt{\ln n}}}]$$

with probability $1 - o(1)$. Conditioned on such an (N, M) pair, and conditioned on the identity of the N post-COMP individuals and M positive tests, it remains to couple the graphs G_1 and G_2. These graphs are conditionally independent so we can consider them separately. The assumption that post-COMP weak detection is impossible implies that the planted and null distributions over G_1 can be coupled with probability $1 - o(1)$. Also, the planted and null distributions over G_2 are identical, namely every individual among the $n - N$ deleted by COMP is independently connected to every test with probability q, conditioned on being connected to at least one negative test. This completes the proof. □

Acknowledgments

We thank Fotis Iliopoulos for helpful discussions during the first stages of this project.

References

[ABJ14] M. Aldridge, L. Baldassini, and O. Johnson. Group testing algorithms: bounds and simulations. *IEEE Transactions on Information Theory*, 60:3671–3687, 2014.

[ACO08] D. Achlioptas and A. Coja-Oghlan. Algorithmic barriers from phase transitions. *Proceedings of 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS'08)*, page 793–802, 2008.

[AG89] R. Arratia and L. Gordon. Tutorial on large deviations for the binomial distribution. *Bulletin of mathematical biology*, 51(1):125–131, 1989.

[AJS16] M. Aldridge, O. Johnson, and J. Scarlett. Improved group testing rates with constant column weight designs. *Proceedings of 2016 IEEE International Symposium on Information Theory (ISIT'16)*, pages 1381–1385, 2016.

[AJS19] M. Aldridge, O. Johnson, and J. Scarlett. Group testing: an information theory perspective. *Foundations and Trends in Communications and Information Theory*, 15:196–392, 2019.

[Ald19] M. Aldridge. Individual testing is optimal for nonadaptive group testing in the linear regime. *IEEE Transactions on Information Theory*, 65:2058–2061, 2019.

[Ash90] R. Ash. Information theory, 1990.

[BB20] M. Brennan and G. Bresler. Reducibility and statistical-computational gaps from secret leakage. In *Proceedings of 33rd Conference on Learning Theory (COLT'20)*, pages 648–847, 2020.

[BBH+21] M. Brennan, G. Bresler, S. Hopkins, J. Li, and T. Schramm. Statistical query algorithms and low-degree tests are almost equivalent. In *Proceedings of 34th Conference on Learning Theory (COLT'21)*, 2021.

[BBK+21] A. Bandeira, J. Banks, D. Kunisky, C. Moore, and A. Wein. Spectral planting and the hardness of refuting cuts, colorability, and communities in random graphs. In *Proceedings of 34th Conference on Learning Theory (COLT'21)*, pages 410–473, 2021.

[BEH+22] Afonso S Bandeira, Ahmed El Alaoui, Samuel B Hopkins, Tselil Schramm, Alexander S Wein, and Ilias Zadik. The Franz-Parisi criterion and computational trade-offs in high dimensional statistics. *arXiv preprint arXiv:2205.09727*, 2022.

[BHK+19] B. Barak, S. Hopkins, J. Kelner, P. Kothari, A. Moitra, and A. Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM Journal on Computing*, 48(2):687–735, 2019.

[BKW20] A. Bandeira, D. Kunisky, and A. Wein. Computational hardness of certifying bounds on constrained PCA problems. In *11th Innovations in Theoretical Computer Science Conference (ITCS'20)*, 2020.

[BMR20] J. Barbier, N. Macris, and C. Rush. All-or-nothing statistical and computational phase transitions in sparse spiked matrix estimation. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.

[BR13] Q. Berthet and P. Rigollet. Computational lower bounds for sparse PCA. *arXiv preprint arXiv:1304.0828*, 2013.

[Can16] Clément Canonne. A short note on Poisson tail bounds, 2016. Available online at http://www.cs.columbia.edu/~ccanonne/files/misc/2017-poissonconcentration.pdf. Accessed May 24, 2022.

[CCJS11] C. Chan, P. Che, S. Jaggi, and V. Saligrama. Non-adaptive probabilistic group testing with noisy measurements: near-optimal bounds with efficient algorithms. *Proceedings of 49th Annual Allerton Conference on Communication, Control, and Computing*, 1:1832–1839, 2011.

[CJSA14] C. Chan, S. Jaggi, V. Saligrama, and S. Agnihotri. Non-adaptive group testing: Explicit bounds and novel algorithms. *IEEE Transactions on Information Theory*, 60(5):3019–3035, 2014.

[COGHKL20a] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick. Information-theoretic and algorithmic thresholds for group testing. *IEEE Transactions on Information Theory*, 66(12):7911–7928, 2020.

[COGHKL20b] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick. Optimal group testing. *Proceedings of the 33rd Conference on Learning Theory (COLT'20)*, page 1–38, 2020.

[COHKL+21] Amin Coja-Oghlan, Max Hahn-Klimroth, Philipp Loick, Noela Müller, Konstantinos Panagiotou, and Matija Pasch. Inference and Mutual Information on Random Factor Graphs. *38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021)*, 187:24:1–24:15, 2021.

[DKWB19] Y. Ding, D. Kunisky, A. Wein, and A. Bandeira. Subexponential-time algorithms for sparse PCA. *arXiv preprint arXiv:1907.11635*, 2019.

[Dor43] R. Dorfman. The detection of defective members of large populations. *Annals of Mathematical Statistics*, 14:436–440, 1943.

[DR96] Devdatt P. Dubhashi and Desh Ranjan. Balls and bins: A study in negative dependence. *BRICS Report Series*, 3(25), Jan. 1996.

[Dur19] Rick Durrett. *Probability - Theory and Examples*. Cambridge University Press, Cambridge, 2019.

[EVM15] A. Emad, K. Varshney, and D. Malioutov. A semiquantitative group testing approach for learning interpretable clinical prediction rules. *Signal Processing with Adaptive Sparse Structured Representations (SPARS'15)*, 2015.

[Fil16] Y. Filmus. Orthogonal basis for functions over a slice of the boolean hypercube. *Electronic Journal of Combinatorics*, 23(P1.23), 2016.

[GJLR21] Oliver Gebhard, Oliver Johnson, Philipp Loick, and Maurice Rolvien. Improved bounds for noisy group testing with constant tests per item. *IEEE Transactions on Information Theory*, 2021.

[GSV05] D. Guo, S. Shamai, and S. Verdú. Mutual information and minimum mean-square error in gaussian channels. *IEEE Transactions on Information Theory*, 51(4):1261–1282, 2005.

[GZ17] D. Gamarnik and I. Zadik. High dimensional linear regression with binary coefficients: Mean squared error and a phase transition. *Proceedings of 30th Conference on Learning Theory (COLT'17)*, 2017.

[HKP+17] S. Hopkins, P. Kothari, A. Potechin, P. Raghavendra, T. Schramm, and D. Steurer. The power of sum-of-squares for detecting hidden structures. In *IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS'17)*, pages 720–731. IEEE, 2017.

[Hop18] S. Hopkins. *Statistical Inference and the Sum of Squares Method*. PhD thesis, Cornell University, 2018.

[HS17] S. Hopkins and D. Steurer. Efficient bayesian estimation from few samples: community detection and related problems. In *IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS'17)*, pages 379–390. IEEE, 2017.

[IZ21] F. Iliopoulos and I. Zadik. Group testing and local search: is there a computational-statistical gap? *Proceedings of the 34th Annual Conference on Learning Theory (COLT'21)*, 134:2499–2551, 2021.

[JLR11] S. Janson, T. Luczak, and A. Rucinski. *Random Graphs*. John Wiley and Sons, 2011.

[KMDZ06] H. Kwang-Ming and D. Ding-Zhu. Pooling designs and nonadaptive group testing: important tools for DNA sequencing. *World Scientific*, 2006.

[KWB19] D. Kunisky, A. Wein, and A. Bandeira. Notes on computational hardness of hypothesis testing: Predictions using the low-degree likelihood ratio. *arXiv preprint arXiv:1907.11636*, 2019.

[LBM20] C. Luneau, J. Barbier, and N. Macris. Information theoretic limits of learning a sparse rule. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*, 2020.

[LWB20] M. Löffler, A. Wein, and A. Bandeira. Computationally efficient sparse clustering. *arXiv preprint arXiv:2005.10817*, 2020.

[Mar65] A. J. Maria. A remark on Stirling's formula. *The American Mathematical Monthly*, 72(10):1096, 1965.

[MDM13] R. Mourad, Z. Dawy, and F. Morcos. Designing pooling systems for noisy high-throughput protein-protein interaction experiments using boolean compressed sensing. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 10:1478–1490, 2013.

[MNB+21] Leon Mutesa, Pacifique Ndishimye, Yvan Butera, Jacob Souopgui, Annette Uwineza, Robert Rutayisire, Ella Larissa Ndoricimpaye, Emile Musoni, Nadine Rujeni, Thierry Nyatanyi, et al. A pooled testing strategy for identifying SARS-CoV-2 at low prevalence. *Nature*, 589(7841):276–280, 2021.

[MRZ15] Andrea Montanari, Daniel Reichman, and Ofer Zeitouni. On the limitation of spectral methods: From the gaussian hidden clique problem to rank-one perturbations of gaussian tensors. *Advances in Neural Information Processing Systems*, 28, 2015.

[MTB12] C. McMahan, J. Tebbs, and C. Bilder. Informative Dorfman screening. *Journal of the International Biometric Society*, 68:287–296, 2012.

[ND00] H. Ngo and D. Du. A survey on combinatorial group testing algorithms with applications to DNA library screening. *Discrete Mathematical Problems with Medical Applications*, 7:171–182, 2000.

[NWZ20] J. Niles-Weed and I. Zadik. The all-or-nothing phenomenon in sparse tensor PCA. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*, 2020.

[NWZ21] J. Niles-Weed and I. Zadik. It was "all" for "nothing": sharp phase transitions for noiseless discrete channels. In *Proceedings of 34th Conference on Learning Theory (COLT'21)*, volume 134, pages 3546–3547, 2021.

[RSS18] P. Raghavendra, T. Schramm, and D. Steurer. High dimensional estimation via sum-of-squares proofs. In *Proceedings of the International Congress of Mathematicians: Rio de Janeiro 2018*, pages 3389–3423. World Scientific, 2018.

[RXZ19a] G. Reeves, J. Xu, and I. Zadik. All-or-nothing phenomena: From single-letter to high dimensions. In *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pages 654–658, 2019.

[RXZ19b] G. Reeves, J. Xu, and I. Zadik. The all-or-nothing phenomenon in sparse linear regression. In *Proceedings of the Thirty-Second Conference on Learning Theory (COLT'19)*, volume 99, pages 2652–2663, 2019.

[SC16] J. Scarlett and V. Cevher. Phase transitions in group testing. *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'16)*, 1:40–53, 2016.

[SC18] J. Scarlett and V. Cevher. Near-optimal noisy group testing via separate decoding of items. *IEEE Journal of Selected Topics in Signal Processing*, 12(5):902–915, 2018.

[Sri11] Murali K Srinivasan. Symmetric chains, Gelfand–Tsetlin chains, and the Terwilliger algebra of the binary Hamming scheme. *Journal of Algebraic Combinatorics*, 34(2):301–322, 2011.

[SW20] T. Schramm and A. Wein. Computational barriers to estimation from low-degree polynomials. *arXiv preprint arXiv:2008.02269*, 2020.

[TAS20] L. Truong, M. Aldridge, and J. Scarlett. On the all-or-nothing behavior of Bernoulli group testing. *IEEE Journal on Selected Areas in Information Theory*, 1(3):669–680, 2020.

[TM06] N. Thierry-Mieg. A new pooling strategy for high-throughput screening: the shifted transversal design. *BMC Bioinformatics*, 7:28, 2006.

[WLZ+11] L. Wang, X. Li, Y. Zhang, Y. Zhang, and K. Zhang. Evolution of scaling emergence in large-scale spatial epidemic spreading. *PloS one*, 6(7):e21197, 2011.

[WXY21] Y. Wu, J. Xu, and S. Yu. Settling the sharp reconstruction thresholds of random graph matching. *arXiv preprint arXiv:2102.00082*, 2021.

[ZSWB21] I. Zadik, M. Song, A. Wein, and J. Bruna. Lattice-based methods surpass sum-of-squares in clustering. *arXiv preprint arXiv:2112.03898*, 2021.

Near-Optimal Sparsity-Constrained Group Testing: Improved Bounds and Algorithms

Oliver Gebhard, Max Hahn-Klimroth, Olaf Parczyk, Manuel Penschuck, Maurice Rolvien, Jonathan Scarlett, and Nelvin Tan

Abstract

Recent advances in noiseless non-adaptive group testing have led to a precise asymptotic characterization of the number of tests required for high-probability recovery in the sublinear regime $k = n^\theta$ (with $\theta \in (0,1)$), with n individuals among which k are infected. However, the required number of tests may increase substantially under real-world practical constraints, notably including bounds on the maximum number Δ of tests an individual can be placed in, or the maximum number Γ of individuals in a given test. While previous works have given recovery guarantees for these settings, significant gaps remain between the achievability and converse bounds. In this paper, we substantially or completely close several of the most prominent gaps. In the case of Δ-divisible items, we show that the definite defectives (DD) algorithm coupled with a random regular design is asymptotically optimal in dense scaling regimes, and optimal to within a factor of e more generally; we establish this by strengthening both the best known achievability and converse bounds. In the case of Γ-sized tests, we provide a comprehensive analysis of the regime $\Gamma = \Theta(1)$, and again establish a precise threshold proving the asymptotic optimality of SCOMP (a slight refinement of DD) equipped with a tailored pooling scheme. Finally, for each of these two settings, we provide near-optimal adaptive algorithms based on sequential splitting, and provably demonstrate gaps between the performance of optimal adaptive and non-adaptive algorithms.

Oliver Gebhard, oliver.gebhard@tu-dortmund.de, Faculty of Computer Science, TU Dortmund University, Dortmund, Germany, 44227.

Max Hahn-Klimroth, maximilian.hahnklimroth@tu-dortmund.de, Faculty of Computer Science, TU Dortmund University, Dortmund, Germany, 44227.

Olaf Parczyk, parczyk@mi.fu-berlin.de, Department of Mathematics and Computer Science, FU Berlin, Berlin, Germany, 14195.

Manuel Penschuck, manuel@ae.cs.uni-frankfurt.de, Institute of Computer Science, Goethe University Frankfurt, Frankfurt, Germany, 60325.

Maurice Rolvien, maurice.rolvien@tu-dortmund.de, Faculty of Computer Science, TU Dortmund University, Dortmund, Germany, 44227.

Jonathan Scarlett, scarlett@comp.nus.edu.sg, Department of Computer Science, National University of Singapore, Singapore, 117418.

Nelvin Tan, tcnt2@cam.ac.uk, Department of Engineering, University of Cambridge, UK, CB2 1PZ.

The authors are listed alphabetically. This work was presented in part at the IEEE International Symposium on Information Theory (ISIT), 2020 [1] and is accepted for publication at IEEE Transactions on Information Theory.

I. INTRODUCTION

The group testing problem, originally introduced by Dorfman [2], is a prominent example of a classical inference problem that has recently regained considerable attention [3], [4], [5]. Briefly, the problem is posed as follows: Among a population of n individuals, a small subset of k individuals is infected with a rare disease. We are able to test groups of individuals at once, and each test result returns positive if (and only if) there is at least one infected individual in the test group. The challenge is to develop strategies for pooling individuals into tests such that the status of every individual can be recovered reliably from the outcomes, and to do so using as few tests as possible.

While the preceding terminology corresponds to medical applications, group testing also has many other key applications [3, Sec. 1.7], ranging from DNA sequencing [6], [7] to protein interaction experiments [8], [9]. Particular attention has been paid to group testing as a tool for the containment of an epidemic crisis. On the one hand, mass testing appears to be an essential tool to face pandemic spread [10], while on the other hand, the capability of efficiently identifying infected individuals fast and at a low cost is indispensable [11]. For the sake of pandemic control, risk surveillance plans aim at an early, fast and efficient identification of infected individuals to prevent diseases from spreading [12], [13], [14].

The group testing problem includes many variants, depending on the presence/absence of noise, possible adaptivity of the tests, recovery requirements, and so on. Our focus in this paper is on the following setup, which has been the focus of numerous recent works (see [3] for a survey):

- The tests are *non-adaptive*, meaning they must all be designed in advance before observing any outcomes. This is highly desirable in applications, as it permits the tests to be implemented in parallel.
- The tests are *noiseless*; this assumption is more realistic in some applications than others, but serves as an important starting point for understanding the problem.
- The goal is *high-probability* identification of each individual's defectivity status (i.e., probability approaching one as $n \to \infty$). While a deterministic (probability-one) recovery guarantee is also feasible in the noiseless setting [5], it requires considerably more tests, incurring a

k^2 dependence on the number of infected individuals (whenever $k \le O(\sqrt{n})$) instead of k.
- The number of infected individuals k is taken to equal n^θ for some $\theta \in (0,1)$,[1] i.e., the *sublinear regime*. Heaps' law of epidemics [15], [16] indicates that this regime is of major interest. In addition, recent hardness results preclude non-trivial recovery guarantees in the linear regime $k = \Theta(n)$ [17], at least under the most widely-adopted recovery criterion.

Under this setup, Coja-Oghlan et al. [18], [4] recently established the exact information-theoretic threshold on the number of tests, in an asymptotic sense including the implied constant. This threshold was originally attained using a *random regular testing* design [18] (see also [19]), improving on earlier results for *Bernoulli testing* [20], [21]. While the recovery algorithm used in [18] is not computationally efficient, the subsequent work [4] attained the same threshold using a *spatially coupled random regular design* and a computationally efficient recovery algorithm.

All of the preceding test designs have in common that each individual takes part in $O(\ln n)$ tests, and each test contains $O(n/k)$ individuals. As a result, these designs face limitations in real-world applications. Firstly, one may face dilution effects: If an infected individual gets tested within a group of many uninfected individuals, the signal of the infection (e.g., concentration of the relevant molecules) might be too low. For instance, a testing scheme for HIV typically should not contain more than 80 individual samples per test [22]. More recently, evidence was found that certain laboratory tests allow pooling of up to 5 individuals [23] or 64 individuals [24] per test for reliably detecting COVID-19 infections. Secondly, it is often the case that each individual can only be tested a certain number of times, due to the limited volume of the sample taken. More generally, test designs with few tests-per-individual and/or individuals-per-test may be favorable due to resource limitations, difficulties in manually placing samples into tests, and so on.

In light of these practical issues, there is substantial motivation to study the group testing problem under the following constraints on the test design:

- Under the Δ-*divisible items constraint* (or *bounded resource model*), any given individual can only be tested at most Δ times;
- Under the Γ-*sized tests constraint* (or *bounded test-size model*), any given test can only contain at most Γ individuals.

Previous studies of group testing under these constraints [25], [26], [27], [1] are surveyed in Section I-A. We note that some of the above practical motivations may warrant more sophisticated models (e.g., random noise models for dilution effects), but nevertheless, noiseless group testing under the preceding constraints serves as an important starting point towards a full understanding. In addition, as with previous works, we only consider the above two constraints separately, though the case that both are present simultaneously may be of interest for future studies.

A. Related Work

As outlined above, the asymptotically optimal performance limits are well-understood in the case of unconstrained test designs, with optimal designs placing each item in $\Delta = \Theta(\ln n)$ tests, and each test containing $\Gamma = \Theta\big(\frac{n}{k}\big)$ items. We refer the reader to [3] for a more detailed survey, and subsequently focus our attention on the (much more limited) prior work considering the constrained variants with $\Delta = o(\ln n)$ and $\Gamma = o\big(\frac{n}{k}\big)$.

The most relevant prior work is that of Gandikota et al. [25], who gave information-theoretic lower bounds on the number of tests under both kinds of constraint, as well as upper bounds via the simple COMP algorithm [28].[2] The main results therein are summarised as follows, assuming the sublinear regime $k = n^\theta$ with $\theta \in (0,1)$ throughout (we sometimes refer to θ as the *density parameter*):

- Δ-divisible items setting:
  - **(Converse)** For $\Delta = o(\ln n)$, any non-adaptive design with error probability at most ξ requires $m \ge \Delta k \big(\frac{n}{k}\big)^{\frac{1-5\xi}{\Delta}}$, for sufficiently small ξ and sufficiently large n . (Theorem 4.1 in [25])
  - **(Achievability)** Under a suitably-chosen random test design and the COMP algorithm, the error probability is at most ξ provided that $m \ge \lceil e\Delta k \big(\frac{n}{\xi}\big)^{\frac{1}{\Delta}} \rceil$. (Theorem 4.2 in [25])
- Γ-sized tests setting:
  - **(Converse)** For $\Gamma = \Theta\big(\big(\frac{n}{k}\big)^\beta\big)$ with $\beta \in [0,1)$, any non-adaptive design with error probability at most ξ requires $m \ge \frac{1-6\xi}{1-\beta} \cdot \frac{n}{\Gamma}$, for sufficiently large n. (Theorem 4.5 in [25])
  - **(Achievability)** Under a suitably-chosen random test design and COMP recovery, for $\Gamma = \Theta\big(\big(\frac{n}{k}\big)^\beta\big)$ with $\beta \in [0,1)$ and $\xi = n^{-\zeta}$ with $\zeta > 0$, the error probability is at most ξ when $m \ge \lceil \frac{1+\zeta}{(1-\theta)(1-\beta)} \rceil \cdot \lceil \frac{n}{\Gamma} \rceil$. (Theorem 4.6 in [25])

A sizable gap remains between the achievability and converse bounds in the case of Δ-divisible items, since typically $\big(\frac{1}{\xi}\big)^{\frac{1}{\Delta}} \gg \big(\frac{1}{k}\big)^{\frac{1}{\Delta}}$. For Γ-sized tests, the bounds match to within a constant factor, but the optimal constant remains unknown. In particular, the two differ by at least a multiplicative $\frac{1}{1-\theta}$ factor, and even for θ close to zero, the two can differ by a factor of 2 due to the rounding in the achievability part.

As we outline further below, we nearly completely close these gaps for Δ-divisible items, and we close them completely for Γ-sized tests in the special case $\beta = 0$ (i.e., $\Gamma = \Theta(1)$) for all $\theta \in (0,1)$. We achieve these results using both the DD and SCOMP algorithms introduced in [20]. While the regime $\beta \in (0,1)$ is also of interest, it appears to require different techniques, and is deferred to future work.

[1] To simplify notation, we assume that $k = n^\theta$ exactly, but all of our analysis and results extend easily to the more general case that $k = cn^\theta$ for any $c = \Theta(1)$.

[2] The COMP algorithm declares any individual in a negative test as uninfected, and all other individuals as infected. It is called Column Matching Algorithm in [25].

	Reference	Number of tests
Δ-div.	Lower Bound [25]	$\Delta k^{1+(1-\theta)/(\Delta\theta)}$
	Lower Bound (Theorem 3.2)	$\max\{e^{-1}\Delta k^{1+(1-\theta)/(\Delta\theta)}, \Delta k^{1+1/\Delta}\}$
	COMP [25]	$e\Delta k n^{\frac{1}{\Delta}}$
	DD (Theorem 3.3)	$\max\{\Delta k^{1+(1-\theta)/(\Delta\theta)}, \Delta k^{1+1/\Delta}\}$
Γ-sized	Lower Bound [25]	$\frac{n}{\Gamma}$
	Lower Bound (Theorem 4.1)	$\max\{(1+\lfloor\frac{\theta}{1-\theta}\rfloor)\frac{n}{\Gamma}, \frac{2n}{\Gamma+1}\}$
	COMP [25]	$\lceil\frac{1}{1-\theta}\rceil\lceil\frac{n}{\Gamma}\rceil$
	SCOMP (Theorems 4.10 and 4.18)	$\max\{(1+\lfloor\frac{\theta}{1-\theta}\rfloor)\frac{n}{\Gamma}, \frac{2n}{\Gamma+1}\}$

TABLE I: Overview of noiseless non-adaptive sparsity-constrained group testing results under the scaling $k = n^\theta$ ($\theta \in (0,1)$). For the setting of Γ-sized tests, this table only corresponds to $\Gamma = \Theta(1)$, and in both settings we neglect higher-order terms and the dependence on the error probability. See the main text for more complete and precise statements.

Gandikota et al. [25] additionally gave explicit designs (i.e., test matrices that can be deterministically constructed in polynomial time), but these give worse scaling laws, and are therefore of less relevance to our results based on random designs. In a distinct but related line of works, Macula [27] and Inan et al. [26], [29] developed designs for the much stronger guarantee of *uniform recovery*, i.e., a single test matrix that uniquely recovers any infected set of size at most k, without allowing any error probability. This stronger guarantee comes at the price of requiring considerably more tests, and we thus omit a direct comparison and refer the interested reader to [27], [26], [29] for details.

B. Contributions

Our main contributions are informally outlined as follows (with $k = n^\theta$ for $\theta \in (0,1)$, and ε being an arbitrarily small constant throughout), with "w.h.p." meaning probability approaching one as $n \to \infty$. The formal statements are given in the theorems referenced. The results are also summarised in Table I (non-adaptive only), and exemplified in Figure 1 (Δ-divisible) and Figure 2 (Γ-sparse).

- Δ-divisible items setting. Assuming that $\Delta = (\ln n)^{1-\Omega(1)}$ (and in some cases, any $\Delta = o(\ln n)$ is allowed), we have the following:
  - **(General converse – Theorem 3.1)** If $m \le (1-\varepsilon)e^{-1}\Delta k^{1+\frac{1-\theta}{\Delta\theta}}$, then w.h.p. any (possibly adaptive) group testing strategy fails.[3]
  - **(Non-adaptive converse – Theorem 3.2)** Under any non-adaptive test design, if $\Delta > \theta/(1-\theta)$ and $m \le (1-\varepsilon)\Delta k^{1+\frac{1}{\Delta}}$, then w.h.p. any inference algorithm fails. Combining with the general lower bound, the same holds for $m \le (1-\varepsilon)\max\left\{e^{-1}\Delta k^{1+\frac{1-\theta}{\Delta\theta}}, \Delta k^{1+\frac{1}{\Delta}}\right\}$.
  - **(Non-adaptive achievability via DD – Theorem 3.3)** Under a random regular test design, DD succeeds when $m \ge (1+\varepsilon)\max\left\{\Delta k^{1+\frac{1-\theta}{\Delta\theta}}, \Delta k^{1+\frac{1}{\Delta}}\right\}$ (w.h.p. when $\Delta = \omega(1)$, and with probability $\Omega(1)$ when $\Delta = \Theta(1)$).
  - **(DD-specific converse – Theorem 3.4)** Under random regular testing, DD fails when m is slightly below the achievability bound (w.h.p. when $\Delta = \omega(1)$, and with $\Omega(1)$ probability when $\Delta = \Theta(1)$).
  - **(Adaptive achievability – Theorem 5.1)** There exists an efficient adaptive algorithm succeeding with probability one when $m \ge (1+\varepsilon)\Delta k^{1+\frac{1-\theta}{\Delta\theta}}$.
- Γ-sized tests setting: Assuming that $\Gamma = \Theta(1)$ in the non-adaptive setting (whereas the adaptive results allow general $\Gamma = o(\frac{n}{k})$), we have the following:
  - **(Non-adaptive converse – Theorem 4.1)** If $m \le (1-\varepsilon)\max\left\{(1+\lfloor\frac{\theta}{1-\theta}\rfloor)\frac{n}{\Gamma}, \frac{2n}{\Gamma+1}\right\}$ and $\Gamma \ge 1 + \left\lfloor\frac{\theta}{1-\theta}\right\rfloor$, then any non-adaptive group testing strategy fails (w.h.p. if $\frac{\theta}{1-\theta}$ is non-integer, and with $\Omega(1)$ probability if $\frac{\theta}{1-\theta}$ is an integer).
  - **(Non-adaptive achievability via SCOMP – Theorems 4.10 and 4.18)** Under a suitably-chosen random test design, SCOMP succeeds w.h.p. when $m \ge \max\left\{(1+\lfloor\frac{\theta}{1-\theta}\rfloor)\frac{n}{\Gamma}, \frac{2n}{\Gamma+1}\right\}$. We use different test designs and analyses for the dense regime $\theta \ge \frac{1}{2}$ (Theorem 4.10) and sparse regime $\theta < \frac{1}{2}$ (Theorem 4.18), and combine the two results to get the overall condition in m in Section IV-F. For the dense regime, our analysis shows that DD has the same guarantee, whereas for the sparse regime, we crucially require the refined SCOMP algorithm.
  - **(Adaptive achievability – Theorem 6.1)** There exists an efficient adaptive algorithm succeeding with probability one when $m \ge (1+\varepsilon)\frac{n}{\Gamma} + k\log_2\Gamma$. In particular, when $\Gamma = o(\frac{n}{k\ln n})$, it suffices that $m \ge (1+\varepsilon)\frac{n}{\Gamma}$.
  - **(General converse – Theorem 6.2)** If $m \le (1-\varepsilon)\frac{n}{\Gamma}$, then the error probability is bounded away from zero for any (possibly adaptive) group testing strategy.

These results have several interesting implications, which we discuss as follows. In the Δ-divisible setting, our first converse bound strengthens that of [25] (removing the -5ξ term in the exponent) and extends it to the adaptive setting, and our second converse provides a further improvement for non-adaptive designs. Our DD achievability result scales as $O\big(\Delta k\big(\max\{k, \frac{n}{k}\}\big)^{\frac{1}{\Delta}}\big)$, which is strictly better than the $O\big(\Delta k\big(\frac{n}{\xi}\big)^{\frac{1}{\Delta}}\big)$ scaling of COMP [25] for all $\theta \in (0,1)$. In fact,

[3] These expressions are obtained after substituting $k = n^\theta$. In the more general case that k equals a positive constant times n^θ, the results remain unchanged upon replacing $k^{1+\frac{1-\theta}{\Delta\theta}}$ by $k\big(\frac{n}{k}\big)^{\frac{1}{\Delta}}$ everywhere. Note also that the achievability bounds may exceed n in some scaling regimes, but in such cases $m = n$ tests still suffice, since one can instead resort to one-by-one testing.

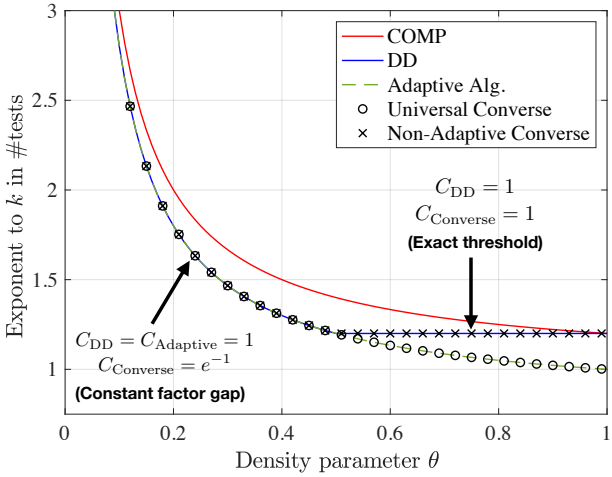

Fig. 1: Illustration of values of η (vertical axis) and C (labeled with text) such that $m = C\Delta k^{\eta}(1+o(1))$ under Δ-divisible item constraints, with $\Delta = 5$.

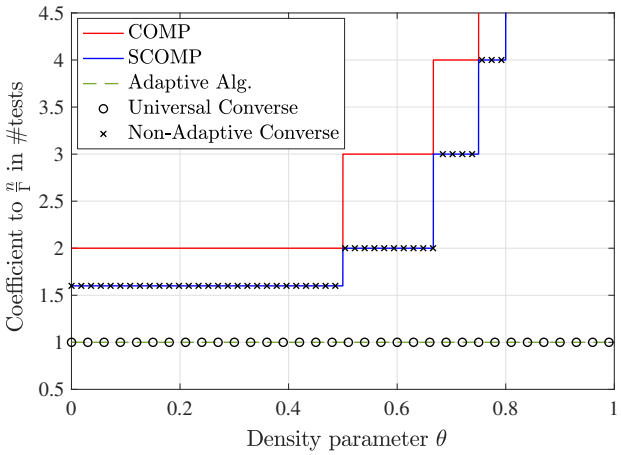

Fig. 2: Illustration of threshold C such that $m = (C+o(1))\frac{n}{\Gamma}$ under Γ-sized test constraints, with $\Gamma = 4$.

for $\theta > \frac{1}{2}$ and $\Delta = \omega(1)$, our results demonstrate that DD is asymptotically optimal among non-adaptive strategies, with a precise phase transition between success and failure at $m \approx \Delta k^{1+\frac{1}{\Delta}}$. For $\theta < \frac{1}{2}$, while establishing a precise phase transition remains an open problem, our results establish DD's optimality up to a multiplicative factor of e, and demonstrate that one cannot reduce the number of tests further under DD and the random regular design. Finally, our results prove a strict adaptivity gap for $\theta > \frac{1}{2}$, and demonstrate that our adaptive algorithm is optimal to within a factor of e for all $\theta \in (0,1)$.

In the Γ-sized tests setting, our results provide an exact asymptotic threshold on the number of tests in the $\Gamma = \Theta(1)$ regime, and we establish the asymptotic optimality of SCOMP in all such cases. To achieve this, we adopt novel analysis techniques specific to this scaling, including a novel test design in the case $\theta < \frac{1}{2}$, as described in the next section. This case of $\theta < \frac{1}{2}$ also has the interesting feature that using SCOMP instead of DD appears to be crucial, in stark contrast with other settings in which the two algorithms tend to have identical asymptotic performance [18]. We note that the distinction between integer and non-integer valued $\frac{\theta}{1-\theta}$ arises due to rounding issues in the analysis, e.g., counting the number of individuals appearing in at most $\lfloor \frac{\theta}{1-\theta} \rfloor$ tests. Our results again demonstrate a strict adaptivity gap (this time for all $\theta \in (0,1)$), and we provide a precise phase transition at $\frac{n}{\Gamma}$ for adaptive algorithms under most scalings of Γ. Finally, in Section VIII, we present numerical results for small population sizes to support our theoretical findings.

II. Fundamentals of Non-Adaptive Group Testing

A. General Notation

Given the number of individuals n, the number of infected individuals $k \sim n^{\theta}(\theta \in (0,1))$, and the number of tests m, we let $\mathscr{G} = (V \cup F, E)$ be a random bipartite (multi-)graph with $|F| = m$ *factor* nodes $(a_1, ..., a_m)$ and $|V| = n$ variable nodes $(x_1, ..., x_n)$. The variable nodes represent individuals, the factor nodes represent tests, and an edge between individual x_i and test a_j indicates, that x_i takes part in test a_j. Furthermore, let $(\partial_{\mathscr{G}} a_1, ..., \partial_{\mathscr{G}} a_m)$ and $(\partial_{\mathscr{G}} x_1, ..., , \partial_{\mathscr{G}} x_n)$ denote the neighbourhoods in $\mathscr{G}$. Whenever the context clarifies what $\mathscr{G}$ is, we will drop the subscript. The test-node degrees are given by $\Gamma_i(\mathscr{G}) = |\partial_{\mathscr{G}} a_i|$, and the individual-node degrees by $\Delta_i(\mathscr{G}) = |\partial_{\mathscr{G}} x_i|$. We can visualise any non-adaptive group testing instance by a *pooling scheme* in the form of such a graph $\mathscr{G}$.

We indicate the infection status of each individual of the population by $\boldsymbol{\sigma} \in \{0,1\}^n$, a uniformly chosen vector of Hamming weight k. Formally, $\boldsymbol{\sigma}_x = 1$ iff x is infected. Then, we let $\hat{\boldsymbol{\sigma}} = \hat{\boldsymbol{\sigma}}(\mathscr{G}, \boldsymbol{\sigma}) \in \{0,1\}^m$ denote the sequence of test results, such that $\hat{\boldsymbol{\sigma}}_a = 1$ iff test a contains at least one infected individual, that is

$$\hat{\boldsymbol{\sigma}}_a = \max_{x \in \partial a} \boldsymbol{\sigma}_x.$$

Throughout the paper, we use standard Landau notation, e.g., $o(1)$ is a function converging to 0 while $\omega(1)$ stands for an arbitrarily slowly diverging function. Moreover, we say that a property $\mathscr{P}$ holds *with high probability (w.h.p.)*, if $\mathbb{P}(\mathscr{P}) = 1 - o(1)$ as $n \to \infty$.

B. Pooling Schemes

The random (almost-)regular bipartite pooling scheme is known to be information-theoretically optimal in the unconstrained variant of group testing [18], and is conceptually simple and easy to implement. In this work, depending on the setup, we sometimes require less standard schemes, as described in the following. It is important to note that in each of these designs, we are constructing a multi-graph rather than a graph, and every multi-edge is counted when referring to a node degree. In the following we will define our choices of the restricted pooling scheme and denote them $\mathscr{G}_{\Delta}$ and $\tilde{\mathscr{G}}_{\Gamma}$

1) Δ-divisible: In this setup, we adopt the design of [18], [19], but with fewer tests per individual in accordance with the problem constraint: Each individual chooses Δ tests uniformly at random with replacement; thus, an individual may be placed in the same test more than once. By construction of $\mathcal{G}_\Delta$, any individual has degree *exactly* Δ, whereas the test degrees fluctuate. We denote by $\mathbf{\Gamma}(\mathcal{G}_\Delta) = \{\mathbf{\Gamma}_1(\mathcal{G}_\Delta), \dots, \mathbf{\Gamma}_m(\mathcal{G}_\Delta)\}$ the (random) sequence of test-degrees.

2) Γ-sparse: In the Γ-sparse case, our choice of pooling scheme requires additional care; we define $\tilde{\mathcal{G}_\Gamma}(\theta)$ separately for two cases:

$$\tilde{\mathcal{G}_\Gamma}(\theta) = \begin{cases} \mathcal{G}_\Gamma & \text{if } \theta \geq 1/2 \\ \mathcal{G}^*_\Gamma & \text{otherwise} \end{cases} \tag{1}$$

with $\mathcal{G}_\Gamma$ and $\mathcal{G}^*_\Gamma$ defined in the following. Throughout the paper, we will always clarify which of the cases we assume, and we will therefore refer to $\tilde{\mathcal{G}_\Gamma}(\theta)$ as $\tilde{\mathcal{G}_\Gamma}$. Starting with $\mathcal{G}_\Gamma$, we employ the *configuration model* [30]. Given n, m, Γ, set $\Delta = m\Gamma/n$ and create for each individual $x \in [n]$ exactly Δ clones $\{x\} \times \{1\}, \dots, \{x\} \times \{\Delta\}$. We assume throughout, that Δ, Γ, n, m are integers, thus all divisibility requirements are fulfilled.[4] Analogously, create Γ clones $\{a\} \times \{1\} \dots \{a\} \times \{\Gamma\}$ for each test $a \in [m]$. Then, choose a perfect matching uniformly at random between the individual-clones and the test-clones and construct a random multi-graph by merging the clones to vertices and adding an edge (x, a) whenever there are $i \in [\Delta], j \in [\Gamma]$ such that the edge $(\{x\} \times \{i\}, \{a\} \times \{j\})$ is part of the perfect matching (in other words, the edge (x, a) exists in the graph as a result of the i-th clone of x and the j-th clone of a being matched). We denote by $\mathcal{G}_\Gamma$ the random regular multi-graph that comes from this procedure.

For $\mathcal{G}^*_\Gamma$, we adopt a different approach. First, we select $\gamma \leq \frac{2n}{\Gamma+1}$ individuals randomly and put them apart for the moment (denote by $X = \{x_1 \dots x_\gamma\}$ the set of those vertices). The precise γ value is chosen such that we can create a random bipartite regular graph on the remaining vertices with each individual having degree 2 and each test having degree $\Gamma - 1$ (thus, an instance of $\mathcal{G}_{\Gamma-1}$). By a simple comparison of degrees, this is only possible if $m \geq 2\frac{n}{\Gamma+1}$. Now, we draw a uniformly random matching between the tests (of degree $\Gamma-1$) and the remaining individuals $x_1 \dots x_\gamma$. By definition, each of those individuals takes part in exactly one test.

In both cases above, $\mathcal{G}^*_\Gamma$ is an almost-regular bipartite graph with each test comprising at most Γ individuals.

C. Choice of recovery algorithm

We make use of the definite defectives (DD) and sequential combinatorial orthogonal matching pursuit (SCOMP) algorithms [20], which are described as follows. Note that SCOMP amounts to running DD and then performing greedy improvements.

[4] It will turn out in due course that $m\Gamma/n$ is an integer under the choice of Γ used in the analysis.

[5] A positive test is unexplained if it does not contain any individuals that have already been marked as infected.

1 Declare every individual x that appears in a negative test as non-infected; remove all such individuals.
2 Declare all individuals that are now the sole individual in a (positive) test as infected.
3 Proceed as follows depending on the algorithm:
- For DD, declare all remaining individuals as uninfected.
- For SCOMP, repeat the following step until no unexplained[5] positive tests remain:
  Declare as infected the (previously undeclared) individual in the largest number of unexplained positive tests.

Algorithm 1: The DD and SCOMP algorithms as defined by [20].

D. The combinatorics behind group testing

In this section, we introduce four types of individuals (see Figure 3) that might appear in any group testing instance and which the student can make use of. It turns out that the sizes of the sets of these individuals are the key to understanding group testing combinatorially. Given a pooling scheme $\mathcal{G}$, let

$$V_0(\mathcal{G}) = \{x \in V(\mathcal{G}) : \boldsymbol{\sigma}_x = 0\}$$
$$\text{and} \qquad V_1(\mathcal{G}) = \{x \in V(\mathcal{G}) : \boldsymbol{\sigma}_x = 1\}$$

be the uninfected and infected individuals, respectively. Then we can define *easy uninfected* individuals to be the uninfected individuals that appear in a negative test – clearly, they can easily be identified. We will call the set of such individuals V_{0-}; formally,

$$V_{0-}(\mathcal{G}) = \{x \in V_0(\mathcal{G}) : \exists a \in \partial_{\mathcal{G}} x : \hat{\boldsymbol{\sigma}}_a = 0\}. \tag{2}$$

Then, there the *easy infected* individuals (sometimes referred to as *definitive defectives*). These are those infected individuals that appear in at least one test with only easy uninfected individuals. Thus, upon removing the easy uninfected individuals, there will be at least one positive test with exactly one undeclared individual, and this individual has to be infected. We call this set

$$V_{1--}(\mathcal{G}) = \{x \in V_1(\mathcal{G}) : \exists a \in \partial_{\mathcal{G}} x : (\partial_{\mathcal{G}} a \setminus \{x\}) \subset V_{0-}(\mathcal{G})\}. \tag{3}$$

Subsequently, there might be *disguised uninfected* individuals, that are uninfected themselves but only appear in positive tests. It is well known [31], [18], [4] that since the prior probability of being uninfected is very large, a group testing instance can tolerate a certain number of individuals of this type. Formally,

$$V_{0+}(\mathcal{G}) = \{x \in V_0(\mathcal{G}) : \forall a \in \partial_{\mathcal{G}} x : \hat{\boldsymbol{\sigma}}_a = 1\}. \tag{4}$$

Finally, there might be *disguised infected* individuals, thus infected individuals appearing only in tests that contain at least one more infected individual. Formally,

$$V_{1+}(\mathcal{G}) = \{x \in V_1(\mathcal{G}) : \forall a \in \partial_{\mathcal{G}} x : (\partial_{\mathcal{G}} a \setminus \{x\}) \cap V_1(\mathcal{G}) \neq \emptyset\}. \tag{5}$$

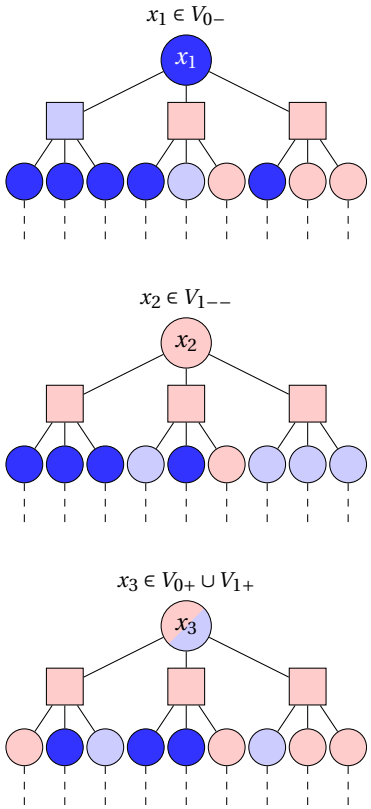

Fig. 3: Rectangles represent tests and circles individuals. Dark blue individuals are elements of V_{0-} and can be easily identified as uninfected. Light blue individuals are elements of V_{0+}, and even if uninfected themselves, they only appear in positive tests and might be hard to identify. Infected individuals (red) that appear only in such tests are impossible to identify. Finally, infected individuals of V_{1--} appear in at least one test with only elements of V_{0-}. Thus, after identifying all elements of V_{0-}, they can be identified. The dashed lines represent the fact that the individuals may also participate in other tests; these may include negative tests classifying their participants as uninfected (elements of V_{0-}) even though the particular test displayed is positive.

While the above types of individuals are not exhaustive, we will see in Section II-E that they are the relevant types for the information-theoretic and algorithmic analyses.

1) Remarks on information-theoretic and combinatorial bounds: It turns out that in the sparse group testing problem – as well as in the unrestricted version [20], [4] – the non-adaptive information-theoretic phase transition comes in two installments. First, there are universal information-theoretic bounds, e.g., counting bounds, that account for the fact that a given number of tests can carry only a certain amount of information. Such bounds directly apply to the non-adaptive as well as the adaptive setting. Second, there are combinatorial / graph theoretical restrictions: Given that there exist a large number of disguised infected individuals (i.e., individuals such that in each of its tests there is a second infected individual), any non-adaptive algorithm fails with high (conditional) probability [18], [4]. This non-adaptivity gap becomes stronger if we increase the infection density parameter θ, because for larger θ, the chance of finding multiple infected individuals in a small neighborhood increases as well. In this section we deal with the combinatorial part. In our setting, the transition where the combinatorial bound dominates the information-theoretic bound happens at $k \sim \sqrt{n}$, i.e., at the point where we find multiple infected individuals in a bounded neighborhood w.h.p..

E. The Nishimori property

Given a pooling scheme $\mathscr{G}$, a ground truth infection status vector $\boldsymbol{\sigma}$ (drawn uniformly from the vectors of Hamming weight k) and a sequence of test results $\hat{\boldsymbol{\sigma}}$, we denote by $S_k(\mathscr{G},\boldsymbol{\sigma})$ the set of all colorings (i.e., infection status assignments) of individuals $\tau \in \{0,1\}^n$ that would have led to the test outcomes $\hat{\boldsymbol{\sigma}}$ (clearly including $\boldsymbol{\sigma}$ itself). Furthermore, we define $Z_k(\mathscr{G},\boldsymbol{\sigma}) = |S_k(\mathscr{G},\boldsymbol{\sigma})|$. The following proposition states that all sets in $S_k(\mathscr{G},\boldsymbol{\sigma})$ are equally likely given the test outcomes.

Proposition 2.1: [Corollary 2.1 of [18]] For all $\tau \in \{0,1\}^n$ we have

$$\mathbb{P}(\boldsymbol{\sigma} = \tau | \mathscr{G}, \hat{\boldsymbol{\sigma}}) = \frac{\mathbb{1}\{\tau \in S_k(\mathscr{G},\boldsymbol{\sigma})\}}{Z_k(\mathscr{G},\boldsymbol{\sigma})}.$$

This immediately implies the following corollary.

Corollary 2.2: If $Z_k(\mathscr{G},\boldsymbol{\sigma}) \geq \ell$ w.h.p., then any inference algorithm recovers $\boldsymbol{\sigma}$ from $(\mathscr{G},\hat{\boldsymbol{\sigma}})$ with probability at most $\ell^{-1}(1+o(1))$.

In other words, as soon as multiple satisfying assignments exist, one cannot do any better than selecting one uniformly at random, as no further information is included in $\mathscr{G}$ and $\hat{\boldsymbol{\sigma}}$ [32]. The following claims will also be useful.

Claim 2.3: For any test design, we have $Z_k(\mathscr{G},\boldsymbol{\sigma}) \geq |V_{1+}(\mathscr{G})||V_{0+}(\mathscr{G})|$. Hence, conditioned on the sets $V_{1+}(\mathscr{G})$ and $V_{0+}(\mathscr{G})$, any inference algorithm fails with probability at least $1 - \frac{1}{|V_{1+}(\mathscr{G})||V_{0+}(\mathscr{G})|}$.

Proof: The first statement is straightforward and was already given in [18, Fact 3.3], and the second statement follows directly from Corollary 2.2. ∎

Finally, we have the following well-known result on the DD algorithm.

Claim 2.4: The DD algorithm succeeds if and only if $V_1(\mathscr{G}) = V_{1--}(\mathscr{G})$.

Proof: By definition, DD first classifies all $x \in V_{0-}(\mathscr{G})$ correctly. In the second step, DD classifies those individuals x as infected, which belong to a positive test a such that $\partial a \setminus \{x\} \subset V_{0-}(\mathscr{G})$. Thus, DD finds all $x \in V_1 \cap V_{1--}(\mathscr{G})$. As DD classifies the remaining individuals as uninfected, it fails as soon as there exists an individual $x \in V_1 \setminus V_{1--}(\mathscr{G})$. ∎

We note that even if $V_1(\mathscr{G}) \setminus V_{1--}(\mathscr{G}) \neq \emptyset$, the DD-algorithm does not produce any false positives but only false negatives. In addition, if DD succeeds then SCOMP is guaranteed to succeed [33], but unlike DD, in general SCOMP may produce both false positive and false negatives.

F. The two-round exposure technique

A key tool to deal with an arbitrary test design is to introduce certain levels of independent randomness. For example, the only randomness in $(\mathscr{G},\boldsymbol{\sigma})$ is the infection status of each individual. We will see in due course we can study an independent infection model (denoted by $\boldsymbol{\sigma}^*$) instead of dealing with exactly k infected individuals, specifically considering each individual as being infected independently from all others with probability $p = \frac{k-\sqrt{k}\ln n}{n}$ (see Corollary 3.6). For the purposes of establishing a converse, the main step is to show that $V_{1+}(\mathscr{G}) \neq \emptyset$, and we will establish this in two steps. We denote by $V_+(\mathscr{G})$ the set of *disguised* individuals, i.e., all tests containing this individual x contain at least one other individual (differing from x) that is infected, and hence

$$V_+(\mathscr{G}) = V_{1+}(\mathscr{G}) \cup V_{0+}(\mathscr{G}).$$

Once we find a large enough set $|V_+(\mathscr{G})| \gg n/k$, there will be some infected individuals in $V_+(\mathscr{G})$ w.h.p.. The main challenge is that in order to find the set of disguised individuals, one uses infected individuals, therefore the events $|V_+(\mathscr{G})|$ exceeding a specific size and infected individuals existing in $V_+(\mathscr{G})$ are not independent in $(\mathscr{G},\boldsymbol{\sigma}^*)$. This is where the two-round exposure technique, used very prominently in the study of random graphs [30], comes into account.

More specifically, our analysis will take the following steps in which individuals are randomly infected:

1) We first mark each individual as infected with probability $\alpha k/n$ for some fixed constant $\alpha \in (0,1)$ and find a set $\mathscr{K}_1$ of infected individuals whose neighbourhood (the tests they belong to) has certain properties.
2) Next, we mark the remaining individuals in the second neighbourhood of $\mathscr{K}_1$ (hence, we look at the individuals that are contained in the tests together with the vertices of $\mathscr{K}_1$) as infected independently with probability $(1-2\alpha)k/n$ for establishing the property of being disguised.
3) After the previous step, each individual has been infected with probability at most $\alpha k/n + (1-\alpha k/n)(1-2\alpha)k/n < p$. To attain the desired final distribution of $\boldsymbol{\sigma}^*$, we independently mark each individual $i \in [n]$ as infected with probability $p - p_i$, where p_i is the probability already incurred from the first two steps. By doing so, the overall distribution of $\boldsymbol{\sigma}^*$ is i.i.d. with probability p, as desired. While these extra infections are not actually analyzed, the idea is that they produce the desired overall distribution, while only enlarging (or keeping unchanged) the set of individuals that are disguised.

III. Non-Adaptive Group Testing with Δ-Divisible Individuals

In this section, we formally state and prove our main results regarding non-adaptive group testing with Δ-divisible individuals.

A. Model

As we highlighted earlier, optimal unconstrained designs are known that place each individual in $\Theta(\ln n)$ tests. Accordingly, we only consider the regime $\Delta = o(\ln n)$, and specifically suppose that $\Delta \le \ln^{1-\delta} n$ for some constant $\delta \in (0,1)$.

B. Results

Define

$$m_{\inf}(\Delta) = \Delta k \max\left\{e^{-1}k^{\frac{(1-\theta)}{\Delta\theta}}, k^{\frac{1}{\Delta}}\right\},$$
$$m_{\mathrm{DD}}(\Delta) = \Delta k \max\left\{k^{\frac{(1-\theta)}{\Delta\theta}}, k^{\frac{1}{\Delta}}\right\}, \tag{6}$$

which will represent the information theoretic converse bound for any non-adaptive group testing scheme and the algorithmic barrier for DD, respectively.

In the following, we assume that $\Delta \ge 2$ and $\Delta > \theta/(1-\theta)$. If the latter inequality is reversed, then we find that $m_{\mathrm{DD}}(\Delta) = \omega(n)$, in which case one is better off resorting to one-by-one testing.

Our first main result provides a simple counting-based converse bound for any adaptive or non-adaptive test design. This result, and all subsequent results, will be proved throughout the rest of the section. An overview of the proof strategy will be provided in Section III-B1

Theorem 3.1: Fix $\varepsilon \in (0,1)$, and suppose that $k = n^\theta$ with $\theta \in (0,1)$ and $\Delta = o(\ln(n))$. Then, if $m \le (1-\varepsilon)\mathrm{e}^{-1}\Delta k^{1+\frac{(1-\theta)}{\Delta\theta}}$ for fixed $\varepsilon > 0$, we have w.h.p. that any (possibly adaptive) group testing procedure that tests each individual at most Δ times fails to recover $\boldsymbol{\sigma}$.

This bound recovers the first term of $\max\{\cdot,\cdot\}$ appearing in the definition of $m_{\inf}(\Delta)$ above, which is dominant for $\theta \le 1/2$. For the second term (which is dominant for $\theta \ge 1/2$), we require a more sophisticated argument that only holds for non-adaptive designs; as we will see in Section V, adaptive designs can in fact go beyond this threshold. The proof of Theorem 3.1 is given in Section III-C.

Theorem 3.2: Given any non-adaptive pooling scheme $\mathscr{G}$ where any individual gets tested at most Δ times (with $\theta/(1-\theta) < \Delta \le (\ln n)^{1-\delta}$ for some $\delta > 0$), if $m \le (1-\varepsilon)\Delta k^{1+1/\Delta}$ for some $\varepsilon \in (0,1)$, any algorithm (efficient or not) fails at inferring $\boldsymbol{\sigma}$ from $(\mathscr{G},\hat{\boldsymbol{\sigma}})$, with probability $1-o(1)$ if $\Delta = \omega(1)$, and with probability $\Omega(1)$ if $\Delta = O(1)$.

Combining these results, we find that any non-adaptive group testing strategy using at most $(1-\varepsilon)m_{\inf}(\Delta)$ tests fails w.h.p. if $\Delta = \omega(1)$, and fails with constant non-zero probability if $\Delta = O(1)$. We provide the proof of Theorem 3.2 in Section III-D. Next, we state our main upper bound, corresponding to the random regular design and the DD algorithm.

Theorem 3.3: Suppose that $m = (1+\varepsilon)m_{\mathrm{DD}}(\Delta)$ for some $\varepsilon > 0$. Then, under the random regular design with parameter Δ, DD recovers $\boldsymbol{\sigma}$ from $(\mathscr{G}_\Delta,\hat{\boldsymbol{\sigma}})$ with probability at least $1-(1+\varepsilon)^{-\Delta}(1+o(1)) - O(n^{-\Omega(1)})$.

Note that the success probability tends to one as $\Delta \to \infty$; if $\Delta = O(1)$ then we need to take $\varepsilon \to \infty$ for the probability to approach one (but it can be close to one for finite

ε). The proof of Theorem 3.3 is given in Section III-E. Comparing this result with Theorem 3.1, we find that DD is asymptotically optimal for $\theta \geq 1/2$. On the other hand, a gap between $m_{\inf}(\Delta)$ and $m_{\mathrm{DD}}(\Delta)$ remains for $\theta < \frac{1}{2}$. In principle, this could be due to a weakness in the converse, a fundamental limitation of DD, or a weakness in our analysis of DD. However, the following theorem rules out the latter of these.

Theorem 3.4: Let $\theta < 1/2$. Given the random regular pooling scheme $\mathscr{G}_\Delta$ on $m = (1-\varepsilon)m_{\mathrm{DD}}(\Delta)$ tests for fixed $\varepsilon \in (0,1)$, we have the following:

1) If $\Delta = \Theta(1)$, then DD fails with positive probability bounded way from zero.
2) If $\Delta = (\ln n)^{1-\delta}$ for $\delta \in (0,1)$, then DD fails w.h.p..

Thus, Theorem 3.4 settles a coarse phase transition of DD in the random regular model when there are finitely many tests-per-individual, and a sharp phase transition when the number of tests-per-individual is diverging. The proof of Theorem 3.4 is provided in Section III-F. We expect that DD is in fact provably suboptimal for $\theta < \frac{1}{2}$, but leave this as an open problem.

1) Overview of proofs: Before proving Theorems 3.1–3.4, we provide a brief overview:

- To prove Theorem 3.1, we establish an upper bound on the probability that an arbitrary inference algorithm recovers $\boldsymbol{\sigma}$ correctly based on the amount of information provided by the test results (which is inherently limited due to the testing constraints). This already suffices to show that as soon as the number of tests crosses a certain lower bound, any inference algorithm must have an error probability approaching one.
- Theorem 3.2 deals with non-adaptive designs, which can be represented as a bipartite graph. The main argument is that when there are too many disguised infected and disguised uninfected individuals, perfect recovery becomes impossible, since interchanging these two types of individuals would not impact the test results. We carefully analyse the number of occurrences of these disguised individuals by the means of local structures in the graph (see Figure 3).
- Theorem 3.3 provides performance guarantees for the DD-algorithm in the $\Delta-$divisible setting. As this algorithm succeeds if and only if all infected individuals appear in one test containing only definitive uninfected individuals (c.f., Sections II-D and II-E), it suffices to analyse a carefully-chosen pooling scheme and pinpoint the number of tests required such that all infected individuals exhibit this property.
- Finally, we prove Theorem 3.4 by showing that as soon as the number of tests is too small, there exists a large number of infected individuals that fail to participate in any tests containing only definitive uninfected individuals.

C. Universal counting-based converse: Proof of Theorem 3.1

We first prove a counting-based upper bound on the success probability for any test design and inference algorithm. Afterwards, we will use this bound on the success probability to prove our main converse bound, providing a lower bound on m for attaining a given target error probability.

Let $\mathscr{A}(\mathscr{G},\hat{\boldsymbol{\sigma}},k)$ be the output of a group testing inference algorithm with input $\mathscr{G}$ (pooling scheme), $\hat{\boldsymbol{\sigma}}$ (test results), and k (number of infected individuals). The inference algorithm is successful if $\mathscr{A}(\mathscr{G},\hat{\boldsymbol{\sigma}},k) = \boldsymbol{\sigma}$, and $\mathbb{P}(\mathscr{A}(\mathscr{G},\hat{\boldsymbol{\sigma}},k) = \boldsymbol{\sigma})$ is the success probability. We first prove the following non-asymptotic counting-based bound via a similar approach to [34] with suitable adjustments, and also using the Nishimori property similarly to [18].

Lemma 3.5: Under the preceding setup, for any pooling scheme $\mathscr{G}$ and inference algorithm $\mathscr{A}(\mathscr{G},\hat{\boldsymbol{\sigma}},k)$, we have

$$\mathbb{P}(\mathscr{A}(\mathscr{G},\hat{\boldsymbol{\sigma}},k) = \boldsymbol{\sigma}) \leq \frac{\sum_{i=0}^{\Delta k}\binom{m}{i}}{\binom{n}{k}}. \tag{7}$$

Proof: Any given pooling scheme can be viewed as a deterministic mapping from an infection status vector $\boldsymbol{\sigma} \in \{0,1\}^n$ to an outcome vector $\hat{\boldsymbol{\sigma}} \in \{0,1\}^m$. Recall that in Proposition 2.1, $S_k(\mathscr{G},\boldsymbol{\sigma})$ is the set of all colorings of individuals that lead to the testing sequence $\hat{\boldsymbol{\sigma}}$, and $Z_k(\mathscr{G},\boldsymbol{\sigma})$ is its cardinality. In the following, we additionally let $\hat{Z}_k(\mathscr{G},\hat{\boldsymbol{\sigma}})$ denote $Z_k(\mathscr{G},\boldsymbol{\sigma})$ when the test outcomes produced by $(\mathscr{G},\boldsymbol{\sigma})$ are equal to $\hat{\boldsymbol{\sigma}}$, and let $\hat{S}_k(\mathscr{G},\hat{\boldsymbol{\sigma}})$ be the set of all $\boldsymbol{\sigma}$ sequences that produce test outcomes $\hat{\boldsymbol{\sigma}}$.

Proposition 2.1 shows that the optimal inference algorithm outputs an arbitrary element of $S_k(\mathscr{G},\boldsymbol{\sigma})$, and is correct with probability (conditioned on $\boldsymbol{\sigma}$) equal to $\frac{1}{Z_k(\mathscr{G},\boldsymbol{\sigma})}$. Thus, averaging over the $\binom{n}{k}$ possible k-sparse vectors $\boldsymbol{\sigma}$, we have the following:

$$\begin{aligned}
\mathbb{P}(\mathscr{A}(\mathscr{G},\hat{\boldsymbol{\sigma}},k) = \boldsymbol{\sigma}) &= \frac{1}{\binom{n}{k}}\sum_{\boldsymbol{\sigma}}\frac{1}{Z_k(\mathscr{G},\boldsymbol{\sigma})} \\
&= \frac{1}{\binom{n}{k}}\sum_{\hat{\boldsymbol{\sigma}}:\hat{Z}_k(\mathscr{G},\hat{\boldsymbol{\sigma}})\geq 1}\ \sum_{\boldsymbol{\sigma}\in\hat{S}_k(\mathscr{G},\hat{\boldsymbol{\sigma}})}\frac{1}{\hat{Z}_k(\mathscr{G},\hat{\boldsymbol{\sigma}})} \\
&\overset{(a)}{\leq} \frac{|\{\hat{\boldsymbol{\sigma}}\in\{0,1\}^m : \hat{Z}_k(\mathscr{G},\hat{\boldsymbol{\sigma}})\geq 1\}|}{\binom{n}{k}} \\
&\overset{(b)}{\leq} \frac{|\{\hat{\boldsymbol{\sigma}} \text{ with at most } \Delta k \text{ ones}\}|}{\binom{n}{k}} \\
&= \frac{\sum_{i=0}^{\Delta k}\binom{m}{i}}{\binom{n}{k}},
\end{aligned}$$

where (a) follows since there are $\hat{Z}_k(\mathscr{G},\hat{\boldsymbol{\sigma}})$ terms in the second summation, thus canceling the $\frac{1}{\hat{Z}_k(\mathscr{G},\hat{\boldsymbol{\sigma}})}$ term, and (b) uses the fact that at most Δk test outcomes can be positive, even in the adaptive setting; this is because adding another infected individual always introduces at most Δ additional positive tests. ∎

We now use the result in (7) to prove Theorem 3.1 - 3.1. In the following we want to provide a short overview of how we obtain these results

Proof of Theorem 3.1: Let $m_{\mathrm{count}}(\Delta) = \mathrm{e}^{-1}\Delta k^{1+\frac{(1-\theta)}{\Delta\theta}}$ denote the threshold in the theorem statement. It suffices to prove the claim for $m = (1-\varepsilon)m_{\mathrm{count}}(\Delta)$, since the inference

algorithm could choose to ignore tests. We use the non-asymptotic bound in Lemma 3.5, and upper bound the sum of binomial coefficients via [35, Section 4.7.1] to obtain the following for a fixed target success probability of $1-\xi$ (for some $\xi \in (0,1)$):

$$\mathbb{P}\left(\mathscr{A}(\mathscr{G}, \hat{\boldsymbol{\sigma}}, k) = \boldsymbol{\sigma}\right) \leq \frac{e^{mh(\frac{\Delta k}{m})}}{\binom{n}{k}} \equiv 1-\xi, \tag{8}$$

where $h(\cdot)$ is the binary entropy function in nats (logs to base e). From (8), we have $e^{mh(\frac{\Delta k}{m})}/\binom{n}{k} = 1-\xi$, which implies that

$$\begin{aligned}\ln\left((1-\xi)\binom{n}{k}\right) &= mh\left(\frac{\Delta k}{m}\right) \\ &= \Delta k \ln\frac{m}{\Delta k} + (m-\Delta k)\ln\frac{1}{1-\frac{\Delta k}{m}} \\ &\overset{(a)}{=} \Delta k \ln\frac{m}{\Delta k} + \Delta k(1+o(1)),\end{aligned} \tag{9}$$

where (a) uses a Taylor expansion and the fact that $\frac{\Delta k}{m} \in o(1)$ (due to $\Delta = o(\ln n)$ and $m = (1-\varepsilon)m_{\text{count}}(\Delta)$). Hence, we have $(1-\frac{\Delta k}{m})^{-1} = \exp(\frac{\Delta k}{m})(1+o(1))$ which is used to obtain the simplification. Rearranging (9), we obtain

$$\ln\frac{m}{\Delta k} = \frac{1}{\Delta k}\ln\left((1-\xi)\binom{n}{k}\right) - (1+o(1)),$$

which gives

$$\begin{aligned}m &= e^{-(1+o(1))}\Delta k\left((1-\xi)\binom{n}{k}\right)^{\frac{1}{\Delta k}} \\ &\overset{(a)}{\geq} e^{-(1+o(1))}(1-\xi)^{\frac{1}{\Delta k}}\Delta k^{1+\frac{1-\theta}{\theta\Delta}},\end{aligned} \tag{10}$$

where (a) follows from the fact that $\binom{n}{k} \geq \left(\frac{n}{k}\right)^k$ and $k = n^\theta$.

Finally, we note that $(1-\xi)^{1/(\Delta k)} \to 1$ for any fixed $\xi \in (0,1)$, since $k \to \infty$ by assumption. This means that m must be at least $(1-o(1))\mathrm{e}^{-1}\Delta k^{1+\frac{(1-\theta)}{\Delta\theta}}$ to obtain any arbitrarily small success probability, and hence, if m is instead a $(1-\varepsilon)$ factor below this threshold (as we have assumed) then the success probability must tend to zero. ∎

D. Universal converse for non-adaptive designs: Proof of Theorem 3.2

It suffices to prove the assertion of the theorem for $m = (1-\varepsilon)\Delta k^{1+1/\Delta}$, since extra tests can only help (or can be ignored). Let $\varepsilon, \theta, \delta \in (0,1)$, and $\theta/(1-\theta) \leq \Delta \leq \ln^{1-\delta} n$. Furthermore, let $\mathscr{G}$ be an arbitrary non-adaptive pooling scheme with $V(\mathscr{G})$ the set of n individuals and $F(\mathscr{G})$ the set of $m = (1-\varepsilon)\Delta k^{1+1/\Delta}$ tests such that each individual is tested at most Δ times. Let

$$\bar{\ell} = \frac{1}{1-\varepsilon}k^{-\frac{1}{\Delta}} \qquad \text{and} \qquad \bar{\Gamma} = \frac{1}{m}\sum_{a\in F(\mathscr{G})}\Gamma_a = \frac{n\Delta}{m} \geq \bar{\ell}\frac{n}{k}. \tag{11}$$

Thus, $\bar{\Gamma}$ represents the average degree of the tests in $F(\mathscr{G})$, where Γ_a is the size of test a. We pick a set of k infected individuals uniformly at random and let $\boldsymbol{\sigma}$ be the $\{0,1\}$-vector representing them. We introduce $p = \frac{k-\sqrt{k}\ln n}{n}$ and $\boldsymbol{\sigma}^*$ as a binomial $\{0,1\}$-vector, such that each entry represents one individual and equals 1 with probability p independently of the others. Our next result relates $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}^*$. As in [17], [4] the way to establish a lower bound is to establish that the underlying graph structure always contains a certain number of disguised infected as well as disguised uninfected individuals. We note that due to the Δ-divisibility condition, a straightforward application of the FKG inequality does not appear to provide a sufficiently strong bound, since the variances of the random variables of interest may become too large.

Corollary 3.6: Under the preceding setup, for fixed $\varepsilon \in (0,1)$ and n large enough, if there is a non-negative integer C (possibly $C=0$) such that

$$\begin{aligned}&\mathbb{P}\left(\left|V_{1+}(\mathscr{G}, \boldsymbol{\sigma}^*)\right| > 2C\right) \geq 1-\varepsilon \\ \text{and} \qquad &\mathbb{P}\left(\left|V_{0+}(\mathscr{G}, \boldsymbol{\sigma}^*)\right| > 2C\right) \geq 1-\varepsilon,\end{aligned}$$

then it also holds that

$$\begin{aligned}&\mathbb{P}\left(|V_{1+}(\mathscr{G}, \boldsymbol{\sigma})| > C\right) \geq 1-\varepsilon-o(1) \\ \text{and} \qquad &\mathbb{P}\left(|V_{0+}(\mathscr{G}, \boldsymbol{\sigma})| > C\right) \geq 1-\varepsilon-o(1).\end{aligned}$$

Proof: The proof follows along the lines of the proof of [4, Lemma 3.6]. Let $\mathscr{B}$ be the event that $|\boldsymbol{\sigma}^*| \in [k - 2\sqrt{k}\ln n, k]$. Then a standard application of the Chernoff bound guarantees that $\mathbb{P}(\mathscr{B}) = 1-o(1)$.

Given $\mathscr{B}$, we couple $\boldsymbol{\sigma}^*$ and $\boldsymbol{\sigma}$ by flipping at most $2\sqrt{k}\ln n$ uninfected individuals in $\boldsymbol{\sigma}^*$ to infected, uniformly at random. This yields the correct distribution, since by definition the set $I_1 = \{i : \boldsymbol{\sigma}^*_i = 1\}$ is a uniform subset of size $|\boldsymbol{\sigma}^*|$ (conditioned on $|\boldsymbol{\sigma}^*|$). Hence, when we infect another random subset of size $k - |I_1|$ uniformly at random, the overall infected set is uniform over the subsets of size k. Clearly, the number of disguised infected individuals can only increase, and hence

$$\left|V_{1+}(\mathscr{G}, \boldsymbol{\sigma}^*)\right| \leq |V_{1+}(\mathscr{G}, \boldsymbol{\sigma})|. \tag{12}$$

However, it might happen that previously disguised uninfected individuals do now contribute to $|V_{1+}(\mathscr{G}, \boldsymbol{\sigma})|$ instead of $|V_{0+}(\mathscr{G}, \boldsymbol{\sigma})|$. Let

$$\boldsymbol{V} := \left||V_{0+}(\mathscr{G}, \boldsymbol{\sigma})| - \left|V_{0+}(\mathscr{G}, \boldsymbol{\sigma}^*)\right|\right|.$$

By the above coupling argument, we have

$$\mathbb{E}[\boldsymbol{V} \mid \mathscr{B}] \leq \frac{2\sqrt{k}\ln n}{n-k}\left|V_{0+}(\mathscr{G}, \boldsymbol{\sigma}^*)\right| < n^{-(1-\theta)}\left|V_{0+}(\mathscr{G}, \boldsymbol{\sigma}^*)\right|.$$

Therefore, Markov's inequality implies

$$\begin{aligned}&\mathbb{P}\left(|V_{0+}(\mathscr{G}, \boldsymbol{\sigma})| \leq \left|V_{0+}(\mathscr{G}, \boldsymbol{\sigma}^*)\right|/2 \mid \mathscr{B}\right) \\ &\qquad \leq \mathbb{P}\left(\boldsymbol{V} \geq \frac{\mathbb{E}[\boldsymbol{V} \mid \mathscr{B}]}{2n^{-(1-\theta)}} \mid \mathscr{B}\right) = o(1).\end{aligned} \tag{13}$$

The desired result now follows directly from (12), (13), and $\mathbb{P}(\mathscr{B}) = 1-o(1)$. ∎

Corollary 3.7: Under the preceding setup, we have the following:

(i) If $\mathbb{P}(|V_{1+}(\mathscr{G}, \boldsymbol{\sigma}^*)| > 0) = 1-o(1)$, then it also holds that $\mathbb{P}(|V_{0+}(\mathscr{G}, \boldsymbol{\sigma})| > \ln n) = 1-o(1)$.

(ii) If $\mathbb{P}(|V_{1+}(\mathscr{G}, \boldsymbol{\sigma}^*)| > 0) = \Omega(1)$, then it also holds that $\mathbb{P}(|V_{0+}(\mathscr{G}, \boldsymbol{\sigma})| > \ln n) = \Omega(1)$.

Proof: We use the fact that the property of being disguised is independent of the infection status. Indeed, given the number of disguised individuals $|V_+(\mathcal{G},\boldsymbol{\sigma}^*)|$, we have $|V_{1+}(\mathcal{G},\boldsymbol{\sigma}^*)| \sim \mathrm{Bin}(|V_+(\mathcal{G},\boldsymbol{\sigma}^*)|, k/n)$ and $|V_{0+}(\mathcal{G},\boldsymbol{\sigma}^*)| \sim \mathrm{Bin}(|V_+(\mathcal{G},\boldsymbol{\sigma}^*)|, 1-k/n)$. Let $\delta > 0$ be such that, by assumption, $\mathbb{P}(|V_{1+}(\mathcal{G},\boldsymbol{\sigma}^*)| > 0) = 1-\delta$. Therefore,

$$\begin{aligned}\delta &= \sum_{n'=1}^{n} \mathbb{P}\left(\left|V_+(\mathcal{G},\boldsymbol{\sigma}^*)\right| = n'\right) \\ &\qquad \cdot \mathbb{P}\left(\left|V_{1+}(\mathcal{G},\boldsymbol{\sigma}^*)\right| = 0 \,\middle|\, \left|V_+(\mathcal{G},\boldsymbol{\sigma}^*)\right| = n'\right) \\ &= \sum_{n'=1}^{n} \mathbb{P}\left(\left|V_+(\mathcal{G},\boldsymbol{\sigma}^*)\right| = n'\right)\left(1-\frac{k}{n}\right)^{n'}. \end{aligned} \tag{14}$$

Observe that if $n' < \frac{n}{k\ln n}$, then we have $(1-k/n)^{n'} = 1-o(1)$. Therefore, due to (14), we require

$$\sum_{n'=1}^{n/(k\ln n)} \mathbb{P}\left(\left|V_+(\mathcal{G},\boldsymbol{\sigma}^*)\right| = n'\right) \le \delta + o(1) \tag{15}$$

and we conclude that $|V_+(\mathcal{G},\boldsymbol{\sigma}^*)| = \tilde{\Omega}\left(\frac{n}{k}\right) = n^{\Omega(1)}$ with probability at least $1-\delta-o(1)$. Moreover, conditioned on $|V_+(\mathcal{G},\boldsymbol{\sigma}^*)| = n^{\Omega(1)}$, the Chernoff bound yields w.h.p. that $|V_{0+}(\mathcal{G},\boldsymbol{\sigma}^*)| = n^{\Omega(1)} > 2\ln n$. The desired result then follows directly from Corollary 3.6, distinguishing between $\delta = o(1)$ and $\delta \in (\Omega(1), 1-\Omega(1))$. ∎

By adopting the two-round exposure technique from Section II-F, Theorem 3.2 will follow from the next lemma, which establishes the conditions in Corollary 3.7 regarding $\boldsymbol{\sigma}^*$.

Lemma 3.8: For any $\varepsilon,\theta,\delta \in (0,1)$ the following holds. Consider the i.i.d. infection model $\boldsymbol{\sigma}^*$, and let $\mathcal{G}$ be a test design such that any of the $n = V(\mathcal{G})$ individuals is tested at most Δ times (with $\theta/(1-\theta) < \Delta \le (\ln n)^{1-\delta}$) and $m = |F(\mathcal{G})| = (1-\varepsilon)\Delta k^{1+1/\Delta}$, where $k = n^\theta$. Then, if $\Delta = \omega(1)$ we have w.h.p. that $|V_{1+}(\mathcal{G},\boldsymbol{\sigma}^*)| > 0$, whereas if $\Delta = O(1)$, we have with $\Omega(1)$ probability that $|V_{1+}(\mathcal{G},\boldsymbol{\sigma}^*)| > 0$.

Proof: We first give a brief overview of the proof:

- We first establish that there must be no tests in $\mathcal{G}$ with too few individuals (Claim 3.9).
- Second, we apply the two-round exposure technique described in Section II to create a set $\mathcal{K}_1$ of infected individuals of size roughly αk.
- Third, we remove any tests that already contain two infected individuals, since individuals of $\mathcal{K}_1$ are disguised if and only if they are disguised upon the removal of such tests (Fact 3.10).
- Next, we show that, upon applying the second stage of the two-round exposure technique to the second neighbourhood of the individuals of $\mathcal{K}_1$ in the remaining graph, the probability an individual $x \in \mathcal{K}_1$ being disguised is minimised in the case that its tests are disjoint (Claim 3.11).
- The preceding result is used to lower bound the average probability of being disguised by employing a hypothetical model in which all tests are mutually disjoint and therefore independent (Claim 3.12).
- Finally, carefully applied concentration results are used complete the proof.

Proceeding more formally, we first show that $\mathcal{G}$ satisfies certain degree properties, namely, there cannot be any tests that are too small.

Claim 3.9: For any fixed integer D, we can assume without loss of generality (for proving Lemma 3.8) that, for n large enough, every test has size at least D.

Proof of Claim 3.9: We obtain an alternative design $\mathcal{G}'$ from $\mathcal{G}$ by iteratively deleting a test of size less than D and all individuals contained in the test, until all tests have size at least D. In each step, we remove one test, between one and D individuals, and at most ΔD edges. Without loss of generality, assume that in $\mathcal{G}$ there are only $o(n)$ individuals that are not contained in any tests (otherwise, the error probability would trivially tend to one). Therefore, the test-design $\mathcal{G}'$ contains at least $(1-o(1))n - m\Delta D = (1-o(1))n$ edges, and since the individual degree is still at most Δ, its number of individuals $n' = |V(\mathcal{G}')|$ satisfies $n' \ge (1-o(1))n/\Delta$. This lower bound on n' along with the assumption $\Delta \le (\ln n)^{1-\delta}$ additionally imply that $\Delta \le (\ln n')^{1-\delta/2}$ when n is sufficiently large.

As for the remaining number of tests $m' = |F(\mathcal{G}')|$, we claim that for all large enough n,

$$m' \le (1-\varepsilon)\Delta n^{\theta+\theta/\Delta} - (n-n')/D \le (1-\varepsilon/2)\Delta (n')^{\theta+\theta/\Delta}. \tag{16}$$

Indeed, the first inequality follows since $m \le (1-\varepsilon)\Delta k^{1+1/\Delta} = (1-\varepsilon)\Delta n^{\theta+\theta/\Delta}$ and the fact that we delete at least one test per D deleted individuals. For the second inequality, let $\zeta := \theta + \theta/\Delta$, which yields $\zeta < 1$ by our assumption $\Delta > \theta/(1-\theta)$. Then, we distinguish two cases:

- If $n - n' \ge \sqrt{n}$, then we have the following:

$$\begin{aligned}(n-n')/D &\ge \Delta (n-n')^\zeta \\ &\ge \Delta\left(n^\zeta - (n')^\zeta\right) \\ &\ge (1-\varepsilon)\Delta\left(n^\zeta - (n')^\zeta\right),\end{aligned}$$

  where the first inequality holds for sufficiently large n since D is constant, $\zeta \in (0,1)$, and Δ is at most logarithmic, and the second inequality holds because the function $f(x) = x^\zeta$ (for $\zeta \in (0,1)$) is concave and monotone, so for any $\delta > 0$ it holds that $f(x+\delta) - f(x)$ is largest when $x = 0$. Substituting the above finding yields the desired second inequality in (16).
- On the other hand, if $n - n' < \sqrt{n}$, we have the following for large enough n:

$$\begin{aligned}(1-\varepsilon)\Delta n^\zeta &< (1-\varepsilon)\Delta (n')^\zeta \cdot (1+\sqrt{n}/n')^\zeta \\ &\le (1-\varepsilon)\Delta (n')^\zeta \cdot (1+\sqrt{n}/n') \\ &\le (1-\varepsilon/2)\Delta (n')^\zeta,\end{aligned}$$

  since $n - n' < \sqrt{n}$ implies that $\sqrt{n}/n' = o(1)$. Hence, in this case we get the desired result even after trivially bounding $(n-n')/D$ by zero.

Since $V_{1+}(\mathcal{G}') \subseteq V_{1+}(\mathcal{G})$, we can continue working with $\mathcal{G}'$ and the desired claim holds. ∎

Recall that in the multi-step argument in Section II-F, for some $\alpha > 0$, the first step is to infect each individual independently with probability $\alpha k/n$, and denote the resulting

set of infected individuals by $\mathscr{K}_1$. We seek to characterize the number of disguised individuals in $\mathscr{K}_1$ following a second step of infections, in which each previously-uninfected individual is infected with probability $(1-2\alpha)k/n$. Given $\mathscr{K}_1$, let $\boldsymbol{X}_v^*$ be the probability that $v \in \mathscr{K}_1$ is disguised after this second step, and let $\boldsymbol{X}^* = \sum_{v\in\mathscr{K}_1} \boldsymbol{X}_v^*$. To prove that $\boldsymbol{X}^*$ is large, we need the following two statements.

Fact 3.10: Let a be a test such that $|\partial a \cap \mathscr{K}_1| \geq 2$. Then any individual in $\mathscr{K}_1$ is disguised if and only if it is disguised when removing the test a.

This fact is immediate as any infected individual is disguised in a by definition. Furthermore, to get a handle on the subtle dependencies between overlapping tests, we prove that the probability for an individual to be disguised in two tests is minimised when the tests are disjoint. For this, denote by $\partial^{(x)} a = \partial a \setminus \{x\}$ the individuals in test a without x.

Claim 3.11: Consider marking each individual in $\partial^{(x)} a \cup \partial^{(x)} a'$ as infected with some probability q independent of the others. Then, for any integer $z > 0$, any individual $x \in V(\mathscr{G})$ and any two tests $a, a' \in \partial x$, we have

$$\begin{aligned}&\mathbb{P}\Big(\partial^{(x)} a \cap V_1(\mathscr{G}) \neq \emptyset, \partial^{(x)} a' \cap V_1(\mathscr{G}) \neq \emptyset \mid \partial^{(x)} a \cap \partial^{(x)} a' = \emptyset\Big)\\ &\quad\leq \mathbb{P}\Big(\partial^{(x)} a \cap V_1(\mathscr{G}) \neq \emptyset,\\ &\qquad\qquad \partial^{(x)} a' \cap V_1(\mathscr{G}) \neq \emptyset \mid \big|\partial^{(x)} a \cap \partial^{(x)} a'\big| = z\Big).\end{aligned}$$

Proof: We first note that

$$\begin{aligned}&\mathbb{P}\left(\partial^{(x)} a \cap V_1(\mathscr{G}) \neq \emptyset, \partial^{(x)} a' \cap V_1(\mathscr{G}) \neq \emptyset \mid \partial^{(x)} a \cap \partial^{(x)} a' = \emptyset\right)\\ &\quad= \left(1-(1-q)^{|\partial^{(x)} a|}\right)\left(1-(1-q)^{|\partial^{(x)} a'|}\right),\end{aligned} \tag{17}$$

as the infected individuals in the two tests are independent due to the conditioning event.

On the other hand, suppose that $\big|\partial^{(x)} a \cap \partial^{(x)} a'\big| = z > 0$. In order to make both tests contain at least one infected individual that is not x, we can either have at least one of the z common individuals which is infected (happening with probability $\left(1-(1-q)^z\right)$), or we need both tests to contain an infected individual outside of the intersection. Hence,

$$\begin{aligned}&\mathbb{P}\left(\partial^{(x)} a \cap V_1(\mathscr{G}) \neq \emptyset, \partial^{(x)} a' \cap V_1(\mathscr{G}) \neq \emptyset \mid \big|\partial^{(x)} a \cap \partial^{(x)} a'\big| = z\right)\\ &\quad= \left(1-\left(1-q\right)^z\right) + \left(1-q\right)^z \left(1-(1-q)^{|\partial^{(x)} a|-z}\right)\\ &\qquad\qquad \cdot \left(1-(1-q)^{|\partial^{(x)} a'|-z}\right)\end{aligned} \tag{18}$$

Using (17) and (18), we conclude the proof with a short calculation:

$$\begin{aligned}&\mathbb{P}\left(\partial^{(x)} a \cap V_1(\mathscr{G}) \neq \emptyset, \partial^{(x)} a' \cap V_1(\mathscr{G}) \neq \emptyset \mid \big|\partial^{(x)} a \cap \partial^{(x)} a'\big| = z\right)\\ &-\mathbb{P}\left(\partial^{(x)} a \cap V_1(\mathscr{G}) \neq \emptyset, \partial^{(x)} a' \cap V_1(\mathscr{G}) \neq \emptyset \mid \partial^{(x)} a \cap \partial^{(x)} a' = \emptyset\right)\\ &= \left(1-\left(1-q\right)^z\right)\\ &\quad + \left(1-q\right)^z \left(1-(1-q)^{|\partial^{(x)} a|-z}\right)\left(1-(1-q)^{|\partial^{(x)} a'|-z}\right)\\ &\quad - \left(1-(1-q)^{|\partial^{(x)} a|}\right)\left(1-(1-q)^{|\partial^{(x)} a'|}\right)\\ &= \left(1-\left(1-q\right)^z\right)\left(1-q\right)^{|\partial^{(x)} a|+|\partial^{(x)} a'|-z} \geq 0,\end{aligned}$$

where the last step follows by expanding and simplifying. ■

With this in mind, we can consider a simplified model in which the test degrees are unchanged, but the *tests are all disjoint*.[6] More precisely, we define the following: Given an infection rate $q \in (0,1)$, we let $\boldsymbol{Y}_a = \boldsymbol{Y}_a(q) := \left(1-\left(1-q\right)^{\Gamma_a - 1}\right)$ be the probability that in a test a of size Γ_a with one fixed individual x, there is at least one infected individual that is not x. For any individual v, we then denote by $\boldsymbol{X}_v = \boldsymbol{X}_v(q) := \prod_{a\in\partial v} \boldsymbol{Y}_a(q)$ the probability that v is disguised in this model, where all tests are mutually disjoint. Observe that, by Claim 3.11, $\boldsymbol{X}_v^* \geq \boldsymbol{X}_v$, and therefore, $\boldsymbol{X}^* \geq \boldsymbol{X}$. The advantage is that in this model, $\boldsymbol{X}_v$ and $\boldsymbol{X}_u$ are independent for $v \neq u$. Recall that

$$\bar{\ell} = \frac{1}{1-\varepsilon} k^{-1/\Delta} = o(1),$$

because $\Delta = O(\ln^{1-\delta} n)$ and $k = n^\theta$, and let $\ell_a = \Gamma_a k/n$.

Note that $\boldsymbol{X}_v$ describes the probability of being disguised for one individual; we proceed by considering the entire set of individuals. The following lemma provides a useful lower bound on $n^{-1}\sum_{v\in V(\mathscr{G})} \boldsymbol{X}_v$.

Claim 3.12: Under the preceding setup with $q = (1-2\alpha)k/n$, we have

$$n^{-1} \sum_{v\in V(\mathscr{G})} \boldsymbol{X}_v \geq (1-\exp(-(1-3\alpha)\bar{\ell}))^\Delta.$$

Proof: By the inequality of arithmetic and geometric means, we have

$$n^{-1} \sum_{v\in V(\mathscr{G})} \boldsymbol{X}_v \geq \prod_{v\in V(\mathscr{G})} \left(\prod_{a\in\partial v} \boldsymbol{Y}_a\right)^{1/n} = \prod_{a\in F(\mathscr{G})} \boldsymbol{Y}_a^{\Gamma_a/n}. \tag{19}$$

Furthermore, by Claim 3.9, we may assume that $\Gamma_a \geq (3\alpha)^{-1}$, and we deduce that

$$\boldsymbol{Y}_a \geq 1-\exp\left(-q\left(\Gamma_a - 1\right)\right) \geq 1-\exp\left(-(1-3\alpha)\ell_a\right).$$

Hence, (19) yields

$$n^{-1} \sum_{v\in V(\mathscr{G})} \boldsymbol{X}_v \geq \prod_{a\in F(\mathscr{G})} \left(1-\exp(-(1-3\alpha)\ell_a)\right)^{\ell_a/k}. \tag{20}$$

Next, we note that $\sum_{a\in F(\mathscr{G})} \Gamma_a \leq \Delta n$ by the Δ-divisibility constraint, which further implies $\sum_{a\in F(\mathscr{G})} \ell_a \leq k\Delta$. The choice $m = (1-\varepsilon)\Delta k^{1+1/\Delta}$ also implies $\bar{\ell} = \Delta k m^{-1}$, and we can characterise the logarithm of the right-hand side of (20) as follows:

$$\begin{aligned}&k^{-1} \sum_{a\in F(\mathscr{G})} \ell_a \ln\left(1-\exp(-(1-3\alpha)\ell_a)\right)\\ &\quad= mk^{-1} \sum_{a\in F(\mathscr{G})} m^{-1}\left(\ell_a \ln\left(1-\exp(-(1-3\alpha)\ell_a)\right)\right)\\ &\quad\geq mk^{-1}\left(\sum_{a\in F(\mathscr{G})} m^{-1}\ell_a\right)\ln\left(1-\exp\left(-(1-3\alpha)m^{-1}\sum_{a\in F(\mathscr{G})}\ell_a\right)\right)\\ &\quad\geq \Delta \ln\left(1-\exp\left(-(1-3\alpha)\bar{\ell}\right)\right),\end{aligned} \tag{21}$$

where the first inequality applies Jensen's inequality applied to the convex function $f(x) = x\ln(1-\exp(-(1-3\alpha)x))$ on $(0,1)$, and the second inequality uses $\bar{\ell} \geq m^{-1}\sum_{a\in F(\mathscr{G})} \ell_a$

[6] This suggests an increase in the number of individuals, but the total number of individuals does not play a role in this part of the analysis.

(by the above calculations regarding $\bar{\ell}$ and ℓ_a above), along with the fact that $\bar{\ell} = \Delta k m^{-1} = o(1)$ and $f(x)$ is a decreasing function for small enough x. Finally, the assertion of the claim follows from (20) and (21). ∎

We note from this claim that if we let $\boldsymbol{v}$ be a uniformly random individual, we have (also using $\bar{\ell} = o(1)$) that

$$\begin{aligned}\mathbb{E}[\boldsymbol{X}_{\boldsymbol{v}}] &\geq \left(1-\exp\left(-(1-3\alpha)\bar{\ell}\right)\right)^{\Delta} \geq (1-4\alpha)^{\Delta}\bar{\ell}^{\Delta} \\ &= \frac{(1-4\alpha)^{\Delta}}{(1-\varepsilon)^{\Delta}k} \geq (1-\varepsilon/2)^{-\Delta}k^{-1},\end{aligned}$$

provided that $\alpha \leq \varepsilon/8$.

Now, recall that $\boldsymbol{X} = \sum_{v\in\mathscr{K}_1}\boldsymbol{X}_v$, and that each individual is in $\mathscr{K}_1$ with probability $\alpha k/n$. Then we deduce from the above that

$$\mathbb{E}[\boldsymbol{X}] = \alpha k\mathbb{E}[\boldsymbol{X}_{\boldsymbol{v}}] \geq \alpha(1-\varepsilon/2)^{-\Delta}.$$

As $\boldsymbol{X}_v$ and $\boldsymbol{X}_u$ are independent for $v \neq u$, we can apply the Chernoff bound (Lemma 7.1, or more precisely a one-sided version that saves a factor of 2) to obtain

$$\mathbb{P}\left(\boldsymbol{X} < \alpha(1-\varepsilon/2)^{-\Delta}/2\right) \leq \exp(-\alpha(1-\varepsilon/2)^{-\Delta}/12). \tag{22}$$

Now, as described earlier, consider infecting any uninfected individual with probability $q = (1-2\alpha)k/n$ independent of all the others. Then, as $\sum_{v\in\mathscr{K}_1}\mathbb{P}(v \in V_{1+}(\mathscr{G})) = \boldsymbol{X}^* \geq \boldsymbol{X}$, we find that conditioned on $\mathscr{K}_1$ and $\boldsymbol{X}$, it holds with probability at least

$$1-\prod_{v\in\mathscr{K}_1}(1-\mathbb{P}(v\in V_{1+}(\mathscr{G}))) \geq 1-\left(1-\frac{\boldsymbol{X}}{|\mathscr{K}_1|}\right)^{|\mathscr{K}_1|} \geq \frac{\boldsymbol{X}}{1+\boldsymbol{X}}$$

that at least one individual from $\mathscr{K}_1$ is disguised. Here we used the inequality of arithmetic and geometric means to upper bound the product, and the last step uses Bernoulli's inequality to write $(1-x/c)^c \leq 1-x \leq \frac{1}{1+x}$. With $\alpha = \varepsilon/8$ and the upper bound (22) on the probability that $\boldsymbol{X} < \varepsilon(1-\varepsilon/2)^{-\Delta}/16$, it follows that there exists a disguised individual in $\mathscr{K}_1$ with probability at least

$$\nu = \nu(\Delta,\varepsilon) := (1-\exp(-\varepsilon(1-\varepsilon/2)^{-\Delta}/96))$$

which yields the statement of Lemma 3.8; note that $\nu = 1-o(1)$ when $\Delta = \omega(1)$, and that $\nu = \Omega(1)$ when $\Delta = O(1)$. The latter assertion holds via the Taylor expansion $1-\exp(-x) = x+\Theta(x^2)$ as $x \to 0$.

Recall that $p = \frac{k-\sqrt{k}\ln n}{n}$, and note that any individual is infected with probability at most

$$\tilde{p} = \alpha k/n + (1-\alpha k/n)(1-2\alpha)k/n < p,$$

independent of all the others. As discussed in Section II-F we can in hindsight raise the infection probability of each individual to p, which can only increase the size of the set $V_{1+}(\mathscr{G})$ (i.e., the number of disguised infected individuals). This yields the assertion of Lemma 3.8 for the i.i.d. infection model. ∎

Proof of Theorem 3.2: The theorem now follows easily by combining Lemma 3.8 with Corollary 3.7: With at least one disguised infected individual and at least $\ln n$ disguised uninfected individuals, the conditional error probability is $1-o(1)$ due to Claim 2.3. ∎

E. Algorithmic achievability on the random regular model: Proof of Theorem 3.3

1) Further notation: Recall the random regular model $\mathscr{G}_\Delta$ from Section II-B1. We let $(\boldsymbol{\Gamma}_1,\dots,\boldsymbol{\Gamma}_m)$ be the (random) sequence of test-degrees, which satisfies the following by construction:

$$\sum_{i=1}^{m}\boldsymbol{\Gamma}_i = n\Delta. \tag{23}$$

Furthermore, given the sequence $(\boldsymbol{\Gamma}_i)_{i\in[m]}$, we define

$$\Gamma_{\min} = \min_{i\in[m]}\boldsymbol{\Gamma}_i, \qquad \bar{\Gamma} = \frac{1}{m}\sum_{i=1}^{m}\boldsymbol{\Gamma}_i = \frac{n\Delta}{m}$$

and

$$\Gamma_{\max} = \max_{i\in[m]}\boldsymbol{\Gamma}_i.$$

We stress at this point that the construction of $\mathscr{G}_\Delta$ allows for multi-edges, and hence one individual might take part in a test multiple times and contribute more than one to its degree.

Moreover, we parametrise the average degree as $\bar{\Gamma} = \ell n/k$, such that ℓ denotes the expected number of infected individuals a test would contain in a binomial random bipartite graph. The definition of $\bar{\Gamma}$ implies $\ell = \frac{k\Delta}{m}$, and substituting $m = (1+\varepsilon)m_{\mathrm{DD}}$ yields

$$\ell = (1+\varepsilon)^{-1}\left(\min\left\{n^{-(1-\theta)/\Delta}, n^{-\theta/\Delta}\right\}\right). \tag{24}$$

Note that with $\frac{\theta}{1-\theta} < \Delta \leq (\ln n)^{1-\Omega(1)}$, we have $\omega(n^{1-\theta}) \leq \ell \leq o(1)$. We will make use of a stronger version of the left inequality stating that $\frac{\ell}{n^{1-\theta}} \geq n^{\Omega(1)}$, which follows from $\Delta > \frac{\theta}{1-\theta}$ and checking both cases of which term in (24) attains the minimum.

We first argue that each test degree is tightly concentrated with high probability, defining the concentration event $\mathscr{C}_\Gamma$ as follows:

$$\begin{aligned}\mathscr{C}_\Gamma = \Big\{\left(1-O\left(n^{-\Omega(1)}\right)\right)\frac{\ell n}{k} \leq \Gamma_{\min} \leq \bar{\Gamma} \leq \Gamma_{\max} \\ \leq \left(1+O\left(n^{-\Omega(1)}\right)\right)\frac{\ell n}{k}\Big\}.\end{aligned} \tag{25}$$

Lemma 3.13: For ℓ given in (24), we have $\mathbb{P}(\mathscr{C}_\Gamma) = 1-\tilde{O}(n^{-3})$.

Proof: Each individual chooses Δ tests with replacement. Hence, each individual has a chance of picking a given test Δ times independently, yielding

$$\boldsymbol{\Gamma}_i = \sum_{j=1}^{n}\sum_{h=1}^{\Delta}\mathbf{1}\{x_j \text{ chooses } a_i \text{ in } h\text{-th selection}\}$$

and

$$\boldsymbol{\Gamma}_i \sim \mathrm{Bin}(n\Delta, 1/m).$$

Thus, we have $\mathbb{E}[\boldsymbol{\Gamma}_i] = \ell n/k$, which scales as $\omega(1)$ since we have established $\ell \geq \omega(n^{1-\theta})$.

Applying the Chernoff bound (Lemma 7.1) and the above-established fact $\frac{\ell}{n^{1-\theta}} \geq n^{\Omega(1)}$, we obtain

$$\mathbb{P}(\boldsymbol{\Gamma}_i < (1-t)\ell n/k) \leq \exp\left(-t^2\ell n^{1-\theta}/3\right) \leq \exp\left(-\Omega\left(t^2 n^{\Omega(1)}\right)\right).$$

Hence, we can choose t of the form $O(n^{-\Omega(1)}\ln n) = O(n^{-\Omega(1)})$ to attain

$$\mathbb{P}(\boldsymbol{\Gamma}_i < (1-t)\ell n/k) = \tilde{O}(n^{-4}). \tag{26}$$

An analogous calculation shows

$$\mathbb{P}(\boldsymbol{\Gamma}_i > (1+t)\ell n/k) = \tilde{O}(n^{-4}). \tag{27}$$

Therefore, the lemma follows from (26), (27), and a union bound over all $m \le n$ tests. ∎

2) Analysis of the different types of individuals: Let $\boldsymbol{Y}_i$ denote the number of infected individuals (including all multi-edges) in test a_i (for $i = 1\dots m$). These variables are not mutually independent, as a single individual takes part in multiple tests. Luckily, it turns out that the family of the $\boldsymbol{Y}_i$ can be approximated by a family of mutually independent random variables sufficiently well. Given $\boldsymbol{\Gamma}_1\dots\boldsymbol{\Gamma}_m$, let $(\boldsymbol{X}_i)_{i\in[m]}$ be a sequence of mutually independent $\mathrm{Bin}(\boldsymbol{\Gamma}_i, k/n)$ variables. Furthermore, let

$$\mathscr{E}_\Delta = \left\{\sum_{i=1}^m \boldsymbol{X}_i = k\Delta\right\} \tag{28}$$

be the event that the sequence $(\boldsymbol{X}_i)$ renders the correct number of infected individuals. Stirling's approximation (Lemma 7.2) guarantees that $\mathscr{E}_\Delta$ is not too unlikely; specifically, $\mathbb{P}(\mathscr{E}_\Delta \mid (\boldsymbol{\Gamma}_i)_i) = \Omega((n\Delta)^{-1/2})$. Furthermore, the $\boldsymbol{X}_i$ are indeed a good local approximation to the correct distribution, as stated in the following known result.

Lemma 3.14: [18, Appendix B.2] Conditioned on $(\boldsymbol{\Gamma}_i)_i$ and $\mathscr{E}_\Delta$, the sequences $(\boldsymbol{Y}_i)_{i\in[m]}$ and $(\boldsymbol{X}_i)_{i\in[m]}$ are identically distributed. ∎

Next, we establish that the number of negative tests $\boldsymbol{m}_0 = \boldsymbol{m}_0(\mathscr{G}_\Delta, \boldsymbol{\sigma})$ and the number of positive tests $\boldsymbol{m}_1 = m - \boldsymbol{m}_0$ are highly concentrated.

Lemma 3.15: With probability at least $1 - o(n^{-2})$ we have

$$\boldsymbol{m}_0 = \big(1 + O\big(n^{-\Omega(1)}\big)\big)\, m\exp(-\ell)$$

and

$$\boldsymbol{m}_1 = \big(1 + O\big(n^{-\Omega(1)}\big)\big)\, m\big(1-\exp(-\ell)\big).$$

Proof: Let $\boldsymbol{m}_0' = \big|\{(\boldsymbol{X}_i)_{i\in[m]} : \boldsymbol{X}_i = 0\}\big|$. Combining the definition of $\boldsymbol{X}_i$ with (7.5), we get

$$\mathbb{E}\big[\boldsymbol{m}_0' \mid (\boldsymbol{\Gamma}_i)_i\big] = \sum_{i=1}^m \mathbb{P}(\boldsymbol{X}_i = 0 \mid \boldsymbol{\Gamma}_i) = \sum_{i=1}^m (1-k/n)^{\boldsymbol{\Gamma}_i},$$

which represents the expected number of negative tests approximated through $(\boldsymbol{X}_i)_i$. Hence, when $(\boldsymbol{\Gamma}_i)_i$ satisfies the concentration event defining $\mathscr{C}_\Gamma$ (see (25)), a second order Taylor expansion (Lemma 7.4) yields

$$\mathbb{E}\big[\boldsymbol{m}_0' \mid (\boldsymbol{\Gamma}_i)_i\big] = \big(1 + O\big(n^{-\Omega(1)}\big)\big)\, m\exp(-\ell). \tag{29}$$

Then, conditioned on $(\boldsymbol{\Gamma}_i)_i$, the Chernoff bound implies implies with probability at least $1 - o(n^{-10})$ that

$$\boldsymbol{m}_0' = \mathbb{E}\big[\boldsymbol{m}_0' \mid (\boldsymbol{\Gamma}_i)_i\big]\big(1 + O(m^{-1/4})\big). \tag{30}$$

The first assertion of the lemma now follows from (29), (30), Lemma 3.13, Lemma 3.14, and the fact that $\mathscr{E}_\Delta$ has probability $\Omega((n\Delta)^{-1/2})$: Letting $\mathscr{A}$ be the above probability-$o(n^{-10})$ event, we simply write $\mathbb{P}(\mathscr{A}|\mathscr{E}_\Delta) \le \frac{\mathbb{P}(\mathscr{A})}{\mathbb{P}(\mathscr{E}_\Delta)}$, and substitute the upper bound on the numerator and lower bound on the denominator.

For the second assertion of the lemma, we need to additionally take note of the fact that $\ell = o(1)$ and hence $m(1-e^{-\ell}) = O(m\ell) \ll m$. But since $m\ell = k\Delta$, this only amounts to replacing $m^{-1/4}$ by $k^{-1/4}$ in the counterpart of (30), and otherwise has no impact. ∎

Next, we provide a characterization of the size of $V_{0+}(\mathscr{G}_\Delta)$, i.e., the number of disguised uninfected individuals.

Lemma 3.16: We have with probability at least $1 - O\big(n^{-\Omega(1)}\big)$ that

$$|V_{0+}(\mathscr{G}_\Delta)| = \big(1 + O\big(n^{-\Omega(1)}\big)\big)\, n\big(1-\exp(-\ell)\big)^\Delta.$$

Proof: Without loss of generality, given $\boldsymbol{m}_1$ and $\mathscr{C}_\Gamma$, we suppose that tests $a_1\dots a_{\boldsymbol{m}_1}$ are the positive tests. By the degree bounds in (25) and Lemma 3.15, the total number of edges connected to a positive test is w.h.p. given by

$$\sum_{i=1}^{\boldsymbol{m}_1} \boldsymbol{\Gamma}_i = \big(1 + O\big(n^{-\Omega(1)}\big)\big)\, m\bar{\Gamma}\big(1-\exp(-\ell)\big). \tag{31}$$

We need to calculate the probability that a given uninfected individual belongs to $V_{0+}(\mathscr{G}_\Delta)$, i.e., each of its Δ edges is connected to a positive test. By a counting argument, we have

$$\begin{aligned}
&\mathbb{P}_{\mathscr{G}_\Delta}\big(x \in V_{0+}(\mathscr{G}_\Delta) \mid x \in V_0(\mathscr{G}_\Delta), \boldsymbol{m}_1, \mathscr{C}_\Gamma, (\boldsymbol{\Gamma}_i)_i\big) \\
&\quad= \binom{\sum_{i=1}^{\boldsymbol{m}_1}\boldsymbol{\Gamma}_i}{\Delta}\binom{\sum_{i=1}^{m}\boldsymbol{\Gamma}_i}{\Delta}^{-1} \\
&\quad= \big(1 + O\big(n^{-\Omega(1)}\big)\big)\big(1-\exp(-\ell)\big)^\Delta,
\end{aligned}$$

where the simplification follows via Claim 7.3 along with (31) and $\sum_{i=1}^m \boldsymbol{\Gamma}_i = m\bar{\Gamma}$.

Therefore,

$$\mathbb{E}_{\mathscr{G}_\Delta}\big[|V_{0+}(\mathscr{G}_\Delta)| \mid \mathscr{C}_\Gamma\big] = \big(1 + O\big(n^{-\Omega(1)}\big)\big)\, n\big(1-\exp(-\ell)\big)^\Delta. \tag{32}$$

Analogously, the second moment turns out to be

$$\begin{aligned}
\mathbb{E}_{\mathscr{G}_\Delta}\big[|V_{0+}(\mathscr{G}_\Delta)|^2 \mid \mathscr{C}_\Gamma\big] &\le \frac{\binom{n-k}{2}\binom{\left(1+O\left(n^{-\Omega(1)}\right)\right)m\bar{\Gamma}\left(1-\exp(-\ell)\right)}{2\Delta}}{\binom{\left(1+O\left(n^{-\Omega(1)}\right)\right)m\bar{\Gamma}}{2\Delta}} \\
&= \big(1 + O\big(n^{-\Omega(1)}\big)\big)\, n^2\big(1-\exp(-\ell)\big)^{2\Delta}.
\end{aligned} \tag{33}$$

The idea of the first line of (33) is to consider pairs of uninfected individuals whose 2Δ combined edges only participate in positive tests.[7] The second line of (33) follows from Stirling's approximation in the form of Claim 7.3. We lemma is now obtained using (32), (33), and Chebyshev's inequality, and noting that $n(1-e^{-\ell})^\Delta = n^{-\Omega(1)}$ (which is seen by using $\ell = o(1)$ to approximate $(1-e^{-\ell})^\Delta$ by ℓ^Δ, and applying (24)). ∎

Let $\boldsymbol{A}$ denote the number of infected individuals that do not belong to the easy uninfected set $V_{1-\!-}(\mathscr{G}_\Delta)$. The following lemma allows us to bound its size.

Lemma 3.17: If $m = (1+\varepsilon)m_{\mathrm{DD}}(\Delta)$, then $\boldsymbol{A} = 0$ with probability at least $1 - (1+\varepsilon)^{-\Delta}(1+o(1)) - O(n^{-\Omega(1)})$.

[7] The contribution of "self-pairs" where a individual just chooses its own Δ edges from the corresponding set is strictly smaller, which is why the expression given is an upper bound rather than an equality.

Proof: We can split (24) into two cases, depending on the sparsity level θ:

$$\ell = \begin{cases} (1+\varepsilon)^{-1} n^{-(1-\theta)/\Delta}, & \text{if} \quad \theta \le 1/2 \\ (1+\varepsilon)^{-1} k^{-1/\Delta}, & \text{if} \quad \theta > 1/2. \end{cases} \tag{34}$$

Recall that $\boldsymbol{m}_1$ is the number of positive tests, and define

$$\begin{aligned} \mathscr{F}_\Delta = \big\{ \boldsymbol{m}_1 = \big(1 + O\big(n^{-\Omega(1)}\big)\big)\, m\big(1-\exp(-\ell)\big)\big\} \\ \cap \Big\{ |V_{0+}(\mathscr{G}_\Delta)| = \big(1 + O\big(n^{-\Omega(1)}\big)\big)\, n\big(1-\exp(-\ell)\big)^\Delta \Big\} \end{aligned} \tag{35}$$

as the event that both the number of positive tests as well as the size of $V_{0+}(\mathscr{G}_\Delta)$ behave as expected. Lemmas 3.15 and 3.16 guarantee that $\mathscr{F}_\Delta$ is a high probability event, namely, $\mathbb{P}\{\mathscr{F}_\Delta\} \ge 1 - \tilde{O}(n^{-1})$. Given $\boldsymbol{m}_1$, we suppose without loss of generality that $a_1 \dots a_{\boldsymbol{m}_1}$ are the tests rendering a positive result.

We describe the number of occurrences of different types of individuals by introducing two sequences of random variables. Define $\boldsymbol{R}_i = (\boldsymbol{R}_i^1, \boldsymbol{R}_i^{0+}, \boldsymbol{R}_i^{0-})_{i \in [\boldsymbol{m}_1]}$ as the number of infected individuals, disguised uninfected individuals of $V_{0+}(\mathscr{G}_\Delta)$, and non-disguised uninfected individuals (those of $V_{0-}(\mathscr{G}_\Delta)$) appearing in test i, respectively. By construction, we have $\boldsymbol{R}_i^{0-} = \Gamma_i - \boldsymbol{R}_i^{0+} - \boldsymbol{R}_i^1$.

Given $|V_{0+}(\mathscr{G}_\Delta)|$ and $\boldsymbol{m}_1$, we approximate these variables by a sequence of mutually independent multinomials. Specifically, let

$$\begin{aligned} \boldsymbol{H}_i &= (\boldsymbol{H}_i^1, \boldsymbol{H}_i^{0+}, \boldsymbol{H}_i^{0-})_{i \in [\boldsymbol{m}_1]} \\ &\overset{\text{i.i.d.}}{\sim} \mathrm{Mult}_{\ge(1,0,0)}\left(\boldsymbol{\Gamma}_i, \left(\frac{k}{n}, \frac{|V_{0+}(\mathscr{G}_\Delta)|}{n}, 1 - \frac{k + |V_{0+}(\mathscr{G}_\Delta)|}{n}\right)\right), \end{aligned} \tag{36}$$

where $\mathrm{Mult}_{\ge(1,0,0)}$ means multinomial conditioned on the first coordinate being at least one. We introduce the event

$$\mathscr{D}_\Delta = \left\{ \sum_{i=1}^{\boldsymbol{m}_1} \boldsymbol{H}_i^1 = k\Delta, \quad \sum_{i=1}^{\boldsymbol{m}_1} \boldsymbol{H}_i^{0+} = |V_{0+}(\mathscr{G}_\Delta)|\,\Delta \right\},$$

and make use of the following.

Claim 3.18: Given $(\Gamma_i)_i$, $|V_{0+}(\mathscr{G}_\Delta)|$, and $\boldsymbol{m}_1$, the distribution of $\boldsymbol{R}_i$ equals the distribution of $\boldsymbol{H}_i$ given $\mathscr{D}_\Delta$. Furthermore, $\mathbb{P}(\mathscr{D}_\Delta) \ge \Omega(n^{-2})$.

Proof: Let $(r_i)_{i \in [\boldsymbol{m}_1]}$ be a sequence with $r_i = (r_i^1, r_i^{0+}, r_i^{0-})$ satisfying

$$S_1 := \sum_{i=1}^{\boldsymbol{m}_1} r_i^1 = k\Delta, \qquad S_{0+} := \sum_{i=1}^{\boldsymbol{m}_1} r_i^{0+} = |V_{0+}(\mathscr{G}_\Delta)|\,\Delta$$

and

$$r_i^{0-} = \Gamma_i - r_i^1 - r_i^{0+}.$$

In addition, let

$$S_{0-} := \sum_{i=1}^{\boldsymbol{m}_1} r_i^{0-}$$

denote the number of connections from individuals in $V_{0-}(\mathscr{G}_\Delta)$ to positive tests. Then, a counting argument gives

$$\begin{aligned} &\mathbb{P}_{\mathscr{G}_\Delta}(\forall\, i \in [\boldsymbol{m}_1] : \boldsymbol{R}_i = r_i \mid (\Gamma_i)_i, |V_{0+}(\mathscr{G}_\Delta)|, \boldsymbol{m}_1) \\ &\quad = \frac{\binom{S_1}{r_1^1 \dots r_{\boldsymbol{m}_1}^1} \binom{S_{0+}}{r_1^{0+} \dots r_{\boldsymbol{m}_1}^{0+}} \binom{S_{0-}}{r_1^{0-} \dots r_{\boldsymbol{m}_1}^{0-}}}{\binom{S_1 + S_{0+} + S_{0-}}{\Gamma_1, \dots, \Gamma_{\boldsymbol{m}_1}}} \\ &\quad = \binom{S_1 + S_{0+} + S_{0-}}{S_1, S_{0+}, S_{0-}}^{-1} \prod_{i=1}^{\boldsymbol{m}_1} \binom{\Gamma_i}{r_i^1, r_i^{0+}, r_i^{0-}}. \end{aligned}$$

Letting $(r_i')_{i \in [\boldsymbol{m}_1]}$ be a second sequence as above, it follows that

$$\frac{\mathbb{P}_{\mathscr{G}_\Delta}(\forall\, i \in [\boldsymbol{m}_1] : \boldsymbol{R}_i = r_i \mid (\Gamma_i)_i, |V_{0+}(\mathscr{G}_\Delta)|, \boldsymbol{m}_1)}{\mathbb{P}_{\mathscr{G}_\Delta}(\forall\, i \in [\boldsymbol{m}_1] : \boldsymbol{R}_i = r_i' \mid (\Gamma_i)_i, |V_{0+}(\mathscr{G}_\Delta)|, \boldsymbol{m}_1)} \tag{37}$$

$$= \prod_{i=1}^{\boldsymbol{m}_1} \frac{\binom{\Gamma_i}{r_i^1 r_i^{0+} r_i^{0-}}}{\binom{\Gamma_i}{r_i^{1\prime} r_i^{0+\prime} r_i^{0-\prime}}}. \tag{38}$$

Next, define

$$R_1 = \sum_{i=1}^{\boldsymbol{m}_1} r_i^1, \qquad R_+ = \sum_{i=1}^{\boldsymbol{m}_1} r_i^{0+}, \qquad \text{and} \qquad R_- = \sum_{i=1}^{\boldsymbol{m}_1} r_i^{0-}$$

and analogously for R_1', R_+', R_-'. By definition, we have

$$R_1 = R_1', \qquad R_+ = R_+' \qquad \text{and} \qquad R_- = R_-'.$$

Then, by the definition of $\boldsymbol{H}$, we have

$$\begin{aligned} &\frac{\mathbb{P}_{\mathscr{G}_\Delta}(\forall\, i \in [\boldsymbol{m}_1] : \boldsymbol{H}_i = r_i \mid (\boldsymbol{\Gamma}_i)_i, |V_{0+}(\mathscr{G}_\Delta)|, \boldsymbol{m}_1, \mathscr{D}_\Delta)}{\mathbb{P}_{\mathscr{G}_\Delta}(\forall\, i \in [\boldsymbol{m}_1] : \boldsymbol{H}_i = r_i' \mid (\boldsymbol{\Gamma}_i)_i, |V_{0+}(\mathscr{G}_\Delta)|, \boldsymbol{m}_1, \mathscr{D}_\Delta)} \\ &= \frac{(k/n)^{R_1} (|V_{0+}(\mathscr{G}_\Delta)|/n)^{R_+} (1 - k/n - |V_{0+}(\mathscr{G}_\Delta)|/n)^{R_-}}{(k/n)^{R_1'} (|V_{0+}(\mathscr{G}_\Delta)|/n)^{R_+'} (1 - k/n - |V_{0+}(\mathscr{G}_\Delta)|/n)^{R_-'}} \\ &\quad \cdot \prod_{i=1}^{\boldsymbol{m}_1} \frac{\binom{\Gamma_i}{r_i^1, r_i^{0+}, r_i^{0-}}}{\binom{\Gamma_i}{r_i'^1, r_i'^{0+}, r_i'^{0-}}} = \prod_{i=1}^{\boldsymbol{m}_1} \frac{\binom{\Gamma_i}{r_i^1, r_i^{0+}, r_i^{0-}}}{\binom{\Gamma_i}{r_i'^1, r_i'^{0+}, r_i'^{0-}}}. \end{aligned} \tag{39}$$

Thus, the first statement of Claim 3.18 follows from (37) and (39), and the second statement follows from Claim 7.5 ∎

We now introduce a random variable that counts (positive) tests featuring only one infected individual and no disguised uninfected individuals. Formally, let

$$\begin{aligned} \boldsymbol{B} &= \sum_{i=1}^{\boldsymbol{m}_1} \mathbf{1}\{\boldsymbol{R}_i^1 + \boldsymbol{R}_i^{0+} = 1\} \\ \text{and} \qquad \boldsymbol{B}' &= \sum_{i=1}^{\boldsymbol{m}_1} \mathbf{1}\{\boldsymbol{H}_i^1 + \boldsymbol{H}_i^{0+} = 1\}. \end{aligned} \tag{40}$$

By the definition of $\boldsymbol{H}_i$ (see (36)), we have

$$\begin{aligned} &\mathbb{E}_{\mathscr{G}_\Delta}\left[\boldsymbol{B}' \mid (\boldsymbol{\Gamma}_i)_i, |V_{0+}(\mathscr{G}_\Delta)|, \boldsymbol{m}_1\right] \\ &\quad = \sum_{i=1}^{\boldsymbol{m}_1} \binom{\boldsymbol{\Gamma_i}}{1, 0, \Gamma_i - 1} \frac{k/n(1 - k/n - |V_{0+}(\mathscr{G}_\Delta)|/n)^{\boldsymbol{\Gamma}_i - 1}}{1 - (1 - k/n)^{\boldsymbol{\Gamma}_i}}. \end{aligned} \tag{41}$$

In the following, we suppose that $\boldsymbol{\Gamma}_i$ satisfies the concentration around $\bar{\Gamma}$ defining event $\mathscr{C}_\Gamma$ (see (25)), and $\boldsymbol{m}_1$ and $|V_{0+}(\mathscr{G}_\Delta)|$ satisfy the concentration defining event $\mathscr{F}_\Delta$ (see

(35)). Using the concentration of $\mathbf{\Gamma}_i$ and the asymptotic expansion $(1-k/n)^{\bar{\Gamma}} = \exp(-\ell(1+O(n^{-\Omega(1)})))$, we find that

$$\mathbb{E}_{\mathscr{G}_\Delta}\left[\boldsymbol{B}' \mid \mathscr{C}_\Gamma, |V_{0+}(\mathscr{G}_\Delta)|, \boldsymbol{m}_1\right] = \left(1+O\left(n^{-\Omega(1)}\right)\right)\boldsymbol{m}_1\bar{\Gamma}\frac{n^{-(1-\theta)}(1-n^{-(1-\theta)}-|V_{0+}(\mathscr{G}_\Delta)|/n)^{\bar{\Gamma}}}{1-\exp(-\ell(1+O(n^{-\Omega(1)})))}, \tag{42}$$

and further applying $\bar{\Gamma} = \frac{n\Delta}{m}$, $k = n^\theta$, and the concentration of $\boldsymbol{m}_1$, we obtain

$$\mathbb{E}_{\mathscr{G}_\Delta}\left[\boldsymbol{B}' \mid \mathscr{C}_\Gamma, |V_{0+}(\mathscr{G}_\Delta)|, \boldsymbol{m}_1\right] = \left(1+O\left(n^{-\Omega(1)}\right)\right)k\Delta\left(1-\frac{k+|V_{0+}(\mathscr{G}_\Delta)|}{n}\right)^{\bar{\Gamma}}. \tag{43}$$

Now, let us distinguish between the cases $\theta \le 1/2$ and $\theta > 1/2$.

Case 1: $\theta > 1/2$: In this case, we have $n/k = o(k)$, and $\ell = (1+\varepsilon)^{-1}k^{-1/\Delta}$. We recall the event $\mathscr{F}_\Delta$ from (35) that gives a concentration condition for $|V_{0+}(\mathscr{G}_\Delta)|$ and $\boldsymbol{m}_1$. Substituting ℓ into (35), we find that given $\mathscr{F}_\Delta$, there is some $\gamma \in (0,1)$ such that

$$|V_{0+}(\mathscr{G}_\Delta)| = \Theta\left((1+\varepsilon)^{-\Delta}n/k\right) = O\left(k^{1-\gamma}\right).$$

Hence, using (43) and applying $\bar{\Gamma} = \ell n/k$, we obtain

$$\begin{aligned}\mathbb{E}_{\mathscr{G}_\Delta}\left[\boldsymbol{B}' \mid \mathscr{C}_\Gamma, \mathscr{F}_\Delta\right] &= \left(1+O\left(n^{-\Omega(1)}\right)\right)k\Delta\left(1-\frac{\left(1+O\left(n^{-\Omega(1)}\right)\right)k}{n}\right)^{\bar{\Gamma}}\\ &= \left(1+O\left(n^{-\Omega(1)}\right)\right)k\Delta\exp(-\ell)\\ &= \left(1+O\left(n^{-\Omega(1)}\right)\right)k\Delta(1-\ell+O(\ell^2)),\end{aligned} \tag{44}$$

by a second-order Taylor expansion of $e^{-\ell}$. Now, $\boldsymbol{B}'$ is a binomial random variable with a random number of trials and a random probability parameter. Clearly, when conditioning on a specific number of trials and a specific probability, $\boldsymbol{B}'$ is a binomial random variable. Therefore, recalling the expression for ℓ in (34), the Chernoff bound guarantees that under the concentration events $\mathscr{C}_\Gamma$ and $\mathscr{F}_\Delta$, we have

$$\boldsymbol{B}' = \left(1+O\left(n^{-\Omega(1)}\right)\right)\Delta k\cdot(1-(1+\varepsilon)^{-1}k^{-1/\Delta}+O(k^{-2/\Delta}))$$

with probability at least $o(n^{-10})$. Then, similar to the proof of Lemma 3.15, Claim 3.18 yields that

$$\boldsymbol{B} = \left(1+O\left(n^{-\Omega(1)}\right)\right)\Delta k\cdot(1-(1+\varepsilon)^{-1}k^{-1/\Delta}+O(k^{-2/\Delta})) \tag{45}$$

with probability $1-O(n^{-\Omega(1)})$. Thus, we can calculate the probability of an infected individual not belonging to $V_{1--}(\mathscr{G})$ (i.e., not being in the easily-identified infected set) as follows. Such an individual has to choose all of its Δ edges out of the $k\Delta-\boldsymbol{B}$ edges that would lead to a test in which the individual could be identified by DD. Hence, we have

$$\mathbb{P}\left(x \notin V_{1--}(\mathscr{G}_\Delta) \mid x \in V_1(\mathscr{G}), \boldsymbol{B}\right) = \binom{k\Delta-\boldsymbol{B}}{\Delta}\binom{k\Delta}{\Delta}^{-1} = (1+o(1))\left((1+\varepsilon)^{-1}k^{-1/\Delta}\right)^{\Delta}, \tag{46}$$

where the simplification holds using (45) and Claim 7.3.[8] Interpreting the average of $\boldsymbol{A}$ as a sum of k probabilities, it follows that

$$\mathbb{E}_{\mathscr{G}_\Delta}\left[\boldsymbol{A} \mid \boldsymbol{B}\right] \le (1+o(1))(1+\varepsilon)^{-\Delta}. \tag{47}$$

Case 2: $\theta \le 1/2$: In this case, we have $\ell = (1+\varepsilon)^{-1}n^{-(1-\theta)/\Delta}$. Hence, given $\mathscr{F}_\Delta$,

$$|V_{0+}(\mathscr{G}_\Delta)| = \left(1-O\left(n^{-\Omega(1)}\right)\right)k(1+\varepsilon)^{-\Delta}. \tag{48}$$

In contrast to the first case, here we find that the influence of the size of disguised uninfected individuals does not vanish asymptotically in relation to the number of infected individuals in (43).

By a similar argument as the first case, (48) and (43) imply

$$\begin{aligned}\mathbb{E}_{\mathscr{G}_\Delta}\left[\boldsymbol{B}' \mid \mathscr{C}_\Gamma, \mathscr{F}_\Delta\right] &= \left(1+O\left(n^{-\Omega(1)}\right)\right)k\Delta\left(1-\frac{k}{n}\left(1-(1+\varepsilon)^{-\Delta}-O\left(\frac{n^{-\Omega(1)}}{(1+\varepsilon)^{\Delta}}\right)\right)\right)^{\bar{\Gamma}} &(49)\\ &= \left(1+O\left(n^{-\Omega(1)}\right)\right)k\Delta\exp\left(-\left(1-(1+\varepsilon)^{-\Delta}-O\left(\frac{n^{-\Omega(1)}}{(1+\varepsilon)^{\Delta}}\right)\right)\ell\right)\\ &= \left(1+O\left(n^{-\Omega(1)}\right)\right)\\ &\quad\cdot\Delta k\left(1-\left(1-(1+\varepsilon)^{-\Delta}-O\left(n^{-\Omega(1)}(1+\varepsilon)^{-\Delta}\right)\right)(\ell+O(\ell^2))\right), &(50)\end{aligned}$$

and similarly to (45), combining this with the Chernoff bound and Claim 3.18 yields that

$$\boldsymbol{B} = \left(1+O\left(n^{-\Omega(1)}\right)\right)\Delta k\cdot(1-(1+\varepsilon)^{-1}n^{-(1-\theta)/\Delta}+O(n^{-2(1-\theta)/\Delta})) \tag{51}$$

with probability $1-O(n^{-\Omega(1)})$. Therefore, the probability of an infected individual not belonging to $V_{1--}(\mathscr{G})$ satisfies the following analog of (46):

$$\mathbb{P}\left(x \notin V_{1--}(\mathscr{G}_\Delta) \mid x \in V_1(\mathscr{G}), \boldsymbol{B}\right) = \binom{k\Delta-\boldsymbol{B}}{\Delta}\binom{k\Delta}{\Delta}^{-1} = (1+o(1))(1+\varepsilon)^{-\Delta}n^{-(1-\theta)}.$$

Since $2\theta-1 \le 0$ by assumption, it follows that

$$\mathbb{E}\left[\boldsymbol{A} \mid \boldsymbol{B}\right] = (1+o(1))(1+\varepsilon)^{-\Delta}n^{\theta}n^{-(1-\theta)} \le (1+o(1))(1+\varepsilon)^{-\Delta}. \tag{52}$$

Thus, Lemma 3.17 follows from (47) and (52) followed by Markov's inequality. ∎

Theorem 3.3 now follows directly from Lemma 3.17 and Claim 2.4.

F. A converse for DD in the sparse regime: Proof of Theorem 3.4

In accordance with Claim 2.4, we first provide a lemma bounding the size of $V_{1--}(\mathscr{G}_\Delta)$, the set of infected individuals appearing in at least one test with only easy uninfected individuals.

[8]The $O(k^{-2/\Delta})$ term in (45) amounts to multiplying by $(1+O(k^{-1/\Delta}))^{\Delta}$ in (46). This simplifies to $1+o(1)$, since $k^{1/\Delta} = \omega(\Delta)$ due to our assumptions $\Delta \le (\ln n)^{1-\Omega(1)}$ and $k = n^\theta$ (this is verified by comparing the logarithms).

Lemma 3.19: For $\theta < 1/2$ and $m = (1-\varepsilon)m_{\mathrm{DD}}(\Delta)$, we have under the random regular design that

$$\mathbb{E}\left[|V_{1--}(\mathscr{G}_\Delta)|\right] = \left(1+O\left(n^{-\Omega(1)}\right)\right)k\left(1-\left(1-\exp\left(-(1-\varepsilon)^{-\Delta}(1-1/\Delta)\right)\right)^{\Delta}\right).$$

Proof: We re-use the notations $\bar{\ell}$ and $\bar{\Gamma}$ in (11), but their expressions are modified as follows in accordance with the choice $m = (1-\varepsilon)\Delta k^{1+\frac{(1-\theta)}{\Delta\theta}}$ associated with $\theta < \frac{1}{2}$:

$$\bar{\ell} = (1-\varepsilon)^{-1}n^{-(1-\theta)/\Delta} \qquad \text{and} \qquad \bar{\Gamma} = (1-\varepsilon)^{-1}n^{(1-\theta)(1-1/\Delta)}. \tag{53}$$

We additionally recall $\boldsymbol{B}$ from (40) as the number of tests featuring exactly one infected individual and no elements of V_{0+}. By the same calculation as in (50) and (51) with ℓ and $\bar{\Gamma}$ replaced by the values in (53), we obtain

$$\begin{aligned}\boldsymbol{B} &= \left(1+O\left(n^{-\Omega(1)}\right)\right)k\Delta\left(1-(1-\varepsilon)^{-\Delta}kn^{-1}\right)^{\bar{\Gamma}} \\ &= \left(1+O\left(n^{-\Omega(1)}\right)\right)k\Delta\exp\left(-(1-\varepsilon)^{-\Delta}(1-1/\Delta)\right) \end{aligned}\tag{54}$$

with probability at least $1-o(n^{-8})$. Therefore, we can calculate the probability that an infected individual does not belong to $V_{1--}(\mathscr{G}_\Delta)$ via Claim 7.3 as follows:

$$\begin{aligned}&\mathbb{P}\left(x \notin V_{1--}(\mathscr{G}_\Delta) \mid x \in V_1(\mathscr{G})\right) \\ &= \left(1+O\left(n^{-\Omega(1)}\right)\right)\frac{\binom{k\Delta-\boldsymbol{B}}{\Delta}}{\binom{k\Delta}{\Delta}} \\ &= \left(1+O\left(n^{-\Omega(1)}\right)\right)\left(1-\exp\left(-(1-\varepsilon)^{-\Delta}(1-1/\Delta)\right)\right)^{\Delta}.\end{aligned}$$

Since there are k individuals in $x \in V_1(\mathscr{G})$ by assumption, we obtain

$$\mathbb{E}\left[|V_1(\mathscr{G}) \setminus V_{1--}(\mathscr{G}_\Delta)|\right] \tag{55}$$

$$= \left(1+O\left(n^{-\Omega(1)}\right)\right)k\left(1-\exp\left(-(1-\varepsilon)^{-\Delta}(1-1/\Delta)\right)\right)^{\Delta} \tag{56}$$

and the lemma follows using $|V_{1--}(\mathscr{G}_\Delta)| = k - |V_1(\mathscr{G}) \setminus V_{1--}(\mathscr{G}_\Delta)|$. ∎

Knowing the expected size of $|V_{1--}(\mathscr{G}_\Delta)|$, Markov's inequality leads to the following.

Corollary 3.20: Let $\theta < 1/2$ and $m = (1-\varepsilon)m_{\mathrm{DD}}(\Delta)$ and $\Delta = \Theta(1)$. Then, with probability at least

$$1 - \frac{1-\left(1-\exp\left(-(1-\varepsilon)^{-\Delta}(1-1/\Delta)\right)\right)^{\Delta}}{1-\gamma} \tag{57}$$

there are at least γk infected individuals $x \in V_1(\mathscr{G}) \setminus V_{1--}(\mathscr{G}_\Delta)$.

Claim 2.4 and Corollary 3.20 immediately imply Theorem 3.4, since (57) is always positive for sufficiently small γ, and approaches one as $\Delta \to \infty$.

IV. Non-Adaptive Group Testing with Γ-Sized Tests

In this section, we formally state and prove our main results concerning non-adaptive group testing Γ-sized tests, namely, a universal lower bound and an algorithmic upper bound that matches the lower bound. Recall that we focus on the regime $\Gamma = \Theta(1)$. Within this section, $\mathscr{G}$ denotes an arbitrary non-adaptive pooling scheme with respect to the Γ-sparsity constraint. The section contains two main parts, outlined as follows:

- Theorem 4.1 states our universal lower bound for non-adaptive designs. The proof is based on a careful analysis of the appearance of disguised individuals (see Section II-D), with the idea being that too many such individuals leads to failure. For $\theta < \frac{1}{2}$, we additionally use the idea of identifying sufficiently many tests with multiple individuals of degree one, prohibiting reliable inference.
- Theorems 4.10 and 4.18 analyze the performance of the DD and SCOMP algorithms. The proofs are again based on the idea that in the underlying pooling scheme, any infected individual appears in at least one test with only definitive uninfected individuals (elements of $V_{0-}(\mathscr{G})$). We refer the reader to Sections II-D and II-E for further insights on these properties. The test size constraints pose additional technical challenges compared to the unconstrained setting [18], in particular leading us to adopt a less standard matching-based test design when $\theta < \frac{1}{2}$.

A. A universal information-theoretic bound

The first statement that we prove is an information-theoretic converse that applies to *any* non-adaptive group testing scheme with maximum test size Γ. Denote by

$$m_{\inf,\Gamma} = \max\left\{\left(1+\left\lfloor\frac{\theta}{1-\theta}\right\rfloor\right)\frac{n}{\Gamma}, 2\frac{n}{\Gamma+1}\right\}, \tag{58}$$

which we will show to be the sharp information-theoretic phase transition point when $\Gamma \geq 1+\left\lfloor\frac{\theta}{1-\theta}\right\rfloor$; note that if this inequality is reversed, then $m_{\inf,\Gamma} > n$, whereas n tests trivially suffice via one-by-one testing. In [25] a lower bound of $(n/\Gamma)(1+o(1))$ was proved, and we see that in the regime $\Gamma = \Theta(1)$, our lower bound improves on this for all $\theta \in (0,1)$.

Theorem 4.1: Let $\theta \in (0,1)$, $\Gamma \geq 1+\left\lfloor\frac{\theta}{1-\theta}\right\rfloor$, and $\delta > 0$. Furthermore, let $\mathscr{G}$ be any non-adaptive pooling scheme (deterministic or randomised) with $m = (1-\delta)m_{\inf,\Gamma}$ tests such that each test contains at most Γ individuals. Then any inference algorithm $\mathscr{A}$ fails in recovering $\boldsymbol{\sigma}$ from $(\hat{\boldsymbol{\sigma}}, \mathscr{G})$

- with probability $1-o(1)$ if $\theta/(1-\theta) \notin \mathbb{Z}$,
- with probability $\Omega(1)$ if $\theta/(1-\theta) \in \mathbb{Z}$.

Thus, even with unlimited computational power, there cannot be any algorithm with a maximum test size of Γ that is able to infer the infected individuals correctly w.h.p. once the number of tests drops below (58). The distinction between integer vs. non-integer values of $\theta/(1-\theta)$ arises for technical reasons (e.g., counting the number of nodes with degree at most $\lfloor\theta/(1-\theta)\rfloor$), and we found it difficult to prove a high-probability (rather than constant-probability) failure result in the integer case.

The proof of the universal information-theoretic converse resembles the proof of [4] for the existence of a universal information-theoretic bound for unrestricted non-adaptive group testing, but several modifications are required to handle the test size constraint. We provide the details in the following subsection.

B. Proof of Theorem 4.1

We start by defining

$$d^+ = 1 + \left\lfloor \frac{\theta}{1-\theta} \right\rfloor \quad \text{and} \quad d^- = \left\lfloor \frac{\theta}{1-\theta} \right\rfloor.$$

For the proof, we distinguish two different regimes for θ, as stated in Proposition 4.2 and Proposition 4.6. We start with the following proposition addressing the existence of disguised individuals.

Proposition 4.2: Let $1/2 \le \theta < 1$, $\Gamma \ge d^+$, and let $\mathcal{G}$ be an arbitrary pooling scheme with tests of size at most Γ. For any $\varepsilon \in (0,1)$, if $m = (1-\varepsilon)d^+\frac{n}{\Gamma}$, then

- $\mathbb{P}(|V_{1+}(\mathcal{G})| > \ln n) \ge 1 - o(1)$ and $\mathbb{P}(|V_{0+}(\mathcal{G})| > \ln n) \ge 1 - o(1)$ if $\frac{\theta}{1-\theta} \notin \mathbb{Z}$
- $\mathbb{P}(|V_{1+}(\mathcal{G})| \ge 1) = \Omega(1)$ and $\mathbb{P}(|V_{0+}(\mathcal{G})| > \ln n) \ge 1 - o(1)$ if $\frac{\theta}{1-\theta} \in \mathbb{Z}$

1) Proof of Proposition 4.2: Let $\mathcal{G}$ be an arbitrary pooling scheme such that each test contains at most Γ individuals. We denote by $V(\mathcal{G})$ the set of individuals, and by $F(\mathcal{G})$ the set of tests in $\mathcal{G}$ (by the identification of $\mathcal{G}$ with a bipartite graph). Instead of analysing $(\mathcal{G}, \hat{\boldsymbol{\sigma}})$, similarly to in the Δ-divisible case, we analyse a related model that eliminates nuisance dependencies between the infection status of different individuals.

Specifically, let $p = \frac{k-\sqrt{k}\ln n}{n}$, and let $\boldsymbol{\sigma}^*$ be a $\{0,1\}$-valued vector, where every entry is one with probability p. Corollary 3.6 guarantees that if the modified model satisfies

$$\mathbb{P}\left(\left|V_{1+}(\mathcal{G}, \boldsymbol{\sigma}^*)\right| > 2C\right) \ge 1 - o(1)$$
$$\text{and} \qquad \mathbb{P}\left(\left|V_{0+}(\mathcal{G}, \boldsymbol{\sigma}^*)\right| > 2C\right) \ge 1 - o(1),$$

then the original model satisfies

$$\mathbb{P}(|V_{1+}(\mathcal{G}, \boldsymbol{\sigma})| > C) \ge 1 - o(1)$$
$$\text{and} \qquad \mathbb{P}(|V_{0+}(\mathcal{G}, \boldsymbol{\sigma})| > C) \ge 1 - o(1)$$

Thus, working with the modified model is sufficient. For the sake of brevity, we henceforth write $\mathcal{G}_\Gamma$ in place of $(\mathcal{G}, \boldsymbol{\sigma}^*)$, leaving the dependencies on $\boldsymbol{\sigma}^*$ implicit.

We proceed by finding a set of (many) individuals, that have a high probability of being disguised. We will apply the probabilistic method iteratively to create the desired set. Creating this set turns out to be delicate due to the dependencies in an arbitrary pooling scheme. Luckily, it will suffice for our purposes to note that whenever individuals have distance at least 6 (i.e., the shortest path between two individuals has at least 6 edges) in the underlying graph, the events of being disguised are independent [4]. To see this, note that we can identify whether an individual is disguised by looking at the tests it is in, and the defectivity status of all other individuals in those tests. This procedure only reaches the second neighborhood, so a separation of 6 is enough to ensure there is no overlap when doing this for two different nodes (which implies independence under an i.i.d. defectivity model).

In the following, we denote the set of all disguised individuals by

$$V^+(\mathcal{G}) = V_{0+}(\mathcal{G}) \cup V_{1+}(\mathcal{G}).$$

We first present a claim establishing that we may safely assume that each individual gets tested $\Theta(1)$ times.

Claim 4.3: Given any pooling scheme $\mathcal{G}'$ with $m = (1-2\varepsilon)d^+\frac{n}{\Gamma}$ (for some $\varepsilon > 0$) such that each test contains at most $\Gamma = \Theta(1)$ individuals, there is another pooling scheme $\mathcal{G}$ such that each test contains at most $\Gamma = \Theta(1)$ individuals with $m = (1-\varepsilon)d^+\frac{n}{\Gamma}$, while also satisfying the following:

- Each individual is contained in at most $C = \Theta(1)$ tests;
- Recovery of $\boldsymbol{\sigma}$ from $(\mathcal{G}', \hat{\boldsymbol{\sigma}}')$ implies recovery from $(\mathcal{G}, \hat{\boldsymbol{\sigma}})$.

Proof: Given $\mathcal{G}'$ and a constant $C \in \mathbb{N}$, there is $C' \in \mathbb{N}$ such that there are at most n/C individuals of degree at least C' in $\mathcal{G}'$, which is an immediate consequence of m being linear in n (due to $\Gamma = \Theta(1)$). Design $\mathcal{G}$ such that each individual of $\mathcal{G}'$ with degree larger than C' gets tested individually (causing n/C additional tests) and all other individuals and tests stay the same as under $\mathcal{G}'$. Clearly, if recovery in $\mathcal{G}'$ was possible, then it is possible in $\mathcal{G}$ as well. Setting $C = \frac{\Gamma}{\varepsilon d^+}$, the claim follows. ∎

In addition to being able to assume there are no individuals with an overly high degree, we can also prove that there cannot be too many individuals with an overly low degree.

Lemma 4.4: Let $\mathcal{G}$ be the given pooling scheme and $m \le (1-\varepsilon)d^+\frac{n}{\Gamma}$, where $\Gamma \ge d^+$. If there is a constant $\alpha > 0$ such that the number of individuals of degree at most d^- is αn, then we have the following:

- $|V_{1+}(\mathcal{G})| > 2\ln n$ w.h.p. if $\theta/(1-\theta) \notin \mathbb{Z}$,
- $|V_{1+}(\mathcal{G})| > 0$ with probability $\Omega(1)$ if $\theta/(1-\theta) \in \mathbb{Z}$.

Proof: Suppose that the number of individuals with degree at most d^- is αn, and recall that $p = \frac{k-\sqrt{k}\ln(n)}{n}$. Without loss of generality, we can assume that there are no tests of degree zero or one. Otherwise, remove them and each connected individual from the testing scheme and note that, by the assumed lower bound $\Gamma \ge d^+$, there are at least εn individuals left. This manipulated graph satisfies the same inequality between the number of individuals and number of tests and, clearly, if the inference of $\boldsymbol{\sigma}$ does not succeed on this manipulated graph, then it cannot succeed in $\mathcal{G}$. Before proceeding, we introduce the following auxiliary result.

Claim 4.5: Under the preceding setup, suppose that there exists a set $I^- \subset V$ of individuals of degree at most d^- with $|I^-| \le \alpha n$ ($\alpha \in (0,1)$). Then, there exists $\beta \in (0,\alpha)$ (depending only on d^- and Γ) such that there must also exist $I^+ \subset I^-$ with $|I^+| = \beta n$, having the property that for all pairs $x \ne y$ in I^+ it holds that $\mathrm{dist}(x,y) \ge 6$.

Proof: First recall from Claim 4.3 that all degrees in the graph are bounded. Consider the procedure of *iterating through all individuals $x \in I^-$, and deleting all $y \in I^-$ of distance at most four from x*, and repeating until no individuals remain. Let I^+ denote set of x's visited by this process. Since the degrees in the graph are finite, each removal only decreases the size of the set I^- by at most a constant, and the assertion of the claim follows. ∎

Let B be the largest possible subset of individuals satisfying the requirements of Claim 4.5. Thus, B is a set of βn

individuals such that for all $x \neq x' \in B$ we have

(B1) $\deg(x) \leq d^-$
(B2) $\mathrm{dist}(x, x') \geq 6$.

We analyze a single individual $x \in B$ using the FKG inequality (e.g., see [36, Proposition 1]); as noted in [17, Lemma 4], the events of x being disguised in each of its tests are increasing with respect to $\boldsymbol{\sigma}^*$ (in the sense that marking additional individuals as infected in $\boldsymbol{\sigma}^*$ can only increase the probability that an individual x is disguised). Hence, the FKG inequality yields the following, recalling that we are considering the case that $\deg(a) \geq 2$ for all a:

$$\mathbb{P}\left(x \in V^+(\mathscr{G})\right) \geq \prod_{a \in \partial x} \left(1 - \left(1-p\right)^{\deg(a)-1}\right).$$

Then, by the fact that $\deg(x) \leq d^- = O(1)$ within B, Claim 7.4 guarantees that

$$\prod_{a \in \partial x} \left(1 - \left(1-p\right)^{\deg(a)-1}\right) \geq C p^{d^-}$$

for some constant C depending on θ and Γ.

We now turn to the total number of disguised individuals in B. As noted above, for two individuals $x, x' \in B$, the events of being disguised are independent due to the pairwise distances being at least 6, as described above. Thus, the number of disguised infected individuals $|V_{1+}(\mathscr{G})|$ dominates a binomial random variable $\mathrm{Bin}(\beta n, p \cdot C p^{d^-})$. Since $np \sim k = n^\theta$, the mean of this binomial distribution scales as $\Theta(n^{\theta-(1-\theta)d^-})$. In particular, when $\frac{\theta}{1-\theta}$ is non-integer, the choice $d^- = \left\lfloor \frac{\theta}{1-\theta} \right\rfloor$ ensures that the exponent is positive, and the Chernoff bound gives w.h.p. that

$$|V_{1+}(\mathscr{G})| \geq n^{\Omega(1)}. \tag{59}$$

On the other hand, if $\frac{\theta}{1-\theta}$ is integer-valued, then the mean of the binomial is $\Theta(1)$, which is enough to ensure that $|V_{1+}(\mathscr{G})| > 0$ with $\Omega(1)$ probability. Combining these two cases completes the proof of Lemma 4.4. ∎

As an immediate consequence of Lemma 4.4, in any group testing instance that succeeds w.h.p., there are at most $o(n)$ individuals of degree up to d^-. However, if $m \leq (1-\varepsilon)d^+ n/\Gamma$ we find at least αn individuals of degree at most d^- (for some α depending on ε) by the handshaking lemma [37, Corollary 1.3], yielding a contradiction. Therefore, Proposition 4.2 is a direct consequence of Lemma 4.4, with the claims regarding $|V_{0+}(\mathscr{G})|$ following easily from those regarding $|V_{1+}(\mathscr{G})|$ in the same way as Corollary 3.7. ∎

We now turn to the sparse regime $\theta < \frac{1}{2}$, establishing the following proposition as a stepping stone to Theorem 4.1.

Proposition 4.6: Let $0 < \theta < 1/2$, and let $\mathscr{G}$ be an arbitrary pooling scheme with tests of size at most Γ. For all $\varepsilon_* > 0$ and sufficiently large n, if $m \leq (2-\varepsilon)\frac{n}{\Gamma+1}$, then any algorithm (efficient or not) fails at recovering $\boldsymbol{\sigma}$ from $\hat{\boldsymbol{\sigma}}$ and $\mathscr{G}$ w.h.p..

2) Proof of Proposition 4.6: The proof hinges on a fairly straightforward observation. We can again assume without loss of generality that there are no tests containing only one individual (otherwise, we remove them and their corresponding individuals from the testing scheme). By a simple counting argument, there can be only $o(n)$ such tests (since otherwise $m > 2n/\Gamma$, which is a contradiction). In addition, we can assume that there are no degree-zero individuals; if there were $\Omega(n)$ of them, high-probability correct inference would trivially be impossible, whereas with $o(n)$ of them, they can be removed and the subsequent analysis still holds for those remaining, with the $o(n)$ difference not impacting the final result.

Then, another counting argument leads to the fact that the number of individuals of degree 1 is large when $m < 2n/\Gamma$, as stated in the following.

Lemma 4.7: If $m = (2-\varepsilon)n/\Gamma$, then there are at least εn individuals of degree 1.

Proof: Denote by αn the number of individuals of degree 1, i.e., $\alpha > 0$ is the proportion of such individuals. Then the lemma follows by double counting edges (on the individual side and on the test-side):

$$(2-\varepsilon)n = m\Gamma \geq \sum_{a \in F(\mathscr{G})} \deg(a) = \sum_{x \in V(\mathscr{G})} \deg(x) \geq \alpha n + 2(1-\alpha)n.$$

Solving for α yields $\alpha \geq \varepsilon$, and the lemma follows. ∎

The next lemma shows that there can only be a small number of tests containing more than one individual of degree 1.

Lemma 4.8: If there is any algorithm recovering $\boldsymbol{\sigma}$ from the test results with $\Omega(1)$ probability, then the number of tests containing more than one individual of degree one is below $n/\sqrt{k} = o(n)$.

Proof: Suppose that at least $n/\sqrt{k}$ tests contain at least two individuals of degree one, and consider any resulting subset of $2n/\sqrt{k}$ individuals (two per test). The average number of infected individuals among these is $(2n/\sqrt{k}) \cdot (k/n) = 2\sqrt{k}$. Hence, by the Chernoff bound for the hypergeometric distribution, w.h.p. there are at least $\sqrt{k}/\ln n$ such infected individuals. On the other hand, among these tests, the average number in which both of these degree-one individuals are infected is $(n/\sqrt{k}) \cdot O((k/n)^2) = O(k\sqrt{k}/n)$, so Markov's inequality implies that w.h.p. the actual number is $O(\sqrt{k} n^{-\Omega(1)})$.

Hence, all but an $o(1)$ fraction of the above-mentioned $\sqrt{k}/\ln n$ infected individuals must be in a test with both a degree-one infected and a degree-one uninfected individual. For these tests, the inference algorithm cannot do better than guess which one is the infected one, but then the probability of all guesses being correct is $(1/2)^{\omega(1)} = o(1)$, from which the lemma follows. ∎

We are now in a position to prove Proposition 4.6. For $m = (2-\varepsilon)n/\Gamma$, we find by Lemma 4.7 that there are at least εn individuals of degree 1. By Lemma 4.8 and the fact that $\Gamma = \Theta(1)$, only $o(n)$ such individuals can be placed together in any tests, and hence, the total number of tests is at least $\varepsilon n - o(n)$. Formally,

$$(2-\varepsilon)n/\Gamma = m \geq \varepsilon n - o(n). \tag{60}$$

Solving (60) for ε, we find $\varepsilon \le \frac{2}{\Gamma+1} + o(1)$. Hence,

$$m \ge \left(2 - \frac{2}{\Gamma+1} - o(1)\right)\frac{n}{\Gamma} = 2\frac{n}{\Gamma+1} - o(n),$$

and the proposition follows. ∎

The universal lower bound in the considered regime is a direct consequence of Proposition 4.2, Proposition 4.6, and Claim 2.3. The proof of Theorem 4.1 is thus complete.

C. Algorithmic bound: Preliminaries and statement of result

We now turn to the problem of establishing an upper bound, with a suitably-chosen test design and an efficient inference algorithm, that matches the universal lower bound. We start by recalling the definition of $\tilde{\mathscr{G}}_\Gamma$ in Section II-B2:

$$\tilde{\mathscr{G}}_\Gamma(\theta) = \begin{cases} \mathscr{G}_\Gamma & \text{if } \theta \ge 1/2 \\ \mathscr{G}_\Gamma^* & \text{otherwise} \end{cases} \tag{61}$$

We equip this pooling scheme with the efficient DD algorithm (see Algorithm 1). In the following, we will see that the combination of these tools will lead to information-theoretically optimal performance in the Γ-sparse setting with $\Gamma = \Theta(1)$.

Proposition 4.9: Define

$$m_{\mathrm{SCOMP}}(\tilde{\mathscr{G}}_\Gamma) = \max\left\{\left(1 + \left\lfloor \frac{\theta}{1-\theta} \right\rfloor\right)\frac{n}{\Gamma}, 2\frac{n}{\Gamma+1}\right\}.$$

For $\Gamma = \Theta(1)$ and $m = (1+\varepsilon)m_{\mathrm{SCOMP}}$, we have

$$\mathbb{P}(\mathscr{A}_{\mathrm{DD}}(\tilde{\mathscr{G}}_\Gamma, \hat{\boldsymbol{\sigma}}, k) = \boldsymbol{\sigma}) = 1 - o(1).$$

To prove this result, we handle the dense regime $\theta > \frac{1}{2}$ in Theorem 4.10 below, the sparse regime $\theta < \frac{1}{2}$ in Theorem 4.18, and combine them in Section IV-F. We observe that $m_{\mathrm{SCOMP}} = m_{\mathrm{inf},\Gamma}$, i.e., the achievability and converse results match for all $\theta \in (0,1)$.

D. Algorithmic feasibility I: The configuration model

We first show that the DD algorithm succeeds with a slightly higher threshold, namely $\max\left\{2, 1 + \left\lfloor \frac{\theta}{1-\theta} \right\rfloor\right\}\frac{n}{\Gamma}$, employing the configuration model $\mathscr{G}_\Gamma$. We define

$$\Delta_{\mathrm{DD}}(\theta) = \max\left\{2, 1 + \left\lfloor \frac{\theta}{1-\theta} \right\rfloor\right\}, m_{\mathrm{DD}}(\mathscr{G}_\Gamma) = \Delta_{\mathrm{DD}}(\theta)\frac{n}{\Gamma}, \tag{62}$$

representing this achievability bound for DD in $\mathscr{G}_\Gamma$.

Theorem 4.10: Let $\varepsilon > 0$ and $m \ge m_{\mathrm{DD}}(\mathscr{G}_\Gamma)$. Then w.h.p. DD infers $\boldsymbol{\sigma}$ from $(\mathscr{G}_\Gamma, \hat{\boldsymbol{\sigma}})$ correctly.

We stress at this point that Theorem 4.10 gives a performance guarantee for the configuration model with any sparsity level, but it will turn out in due course that for $\theta < \frac{1}{2}$ a different model performs slightly better. Note also that for $\theta \ge \frac{1}{2}$, we can simplify $\max\left\{2, 1 + \left\lfloor \frac{\theta}{1-\theta} \right\rfloor\right\} = 1 + \left\lfloor \frac{\theta}{1-\theta} \right\rfloor$.

1) Proof of Theorem 4.10: The proof of Theorem 4.10 hinges on a slightly delicate combinatorial argument. Recall from Figure 3 that V_{1--} consists of those infected individuals that appear in at least one test with only individuals that are removed in the first step of DD (i.e., the easy uninfected individuals V_{0-}). By Claim 2.4, DD succeeds if and only if $V_1(\mathscr{G}) = V_{1--}$.

Lemma 4.11: Let $\boldsymbol{A} = |V_1(\mathscr{G}) \setminus V_{1--}(\mathscr{G}_\Gamma)|$ denote the number of infected individuals that are not identified in the second step of DD. If $m \ge m_{DD}$, then it holds w.h.p. that $\boldsymbol{A} = 0$.

The proof of Lemma 4.11, while conceptually not difficult and similar to [18], is technically challenging, as we have to deal with subtle dependencies in the pooling scheme, caused by the mutli-edges given through the configuration model. A heuristic argument with a (false) independence assumption can provide some intuition as follows: In order for an individual x to be part of a test containing no infected individual (besides possibly x itself) is roughly $(1-k/n)^{\Gamma-1}$. For x to be disguised, thus being element of $V_{0+}(\mathscr{G}_\Gamma)$ or $V_{1+}(\mathscr{G}_\Gamma)$, x may not be part of such a test. Hence, the probability of x being disguised would be roughly $\left(1-(1-k/n)^{\Gamma-1}\right)^{\Delta}$ if the associated Δ events were independent (recall that $\Delta = m\Gamma/n$ is the degree of each individual in the random regular design).

To formally deal with the dependencies in the graph, we proceed as follows. Denote by $(\boldsymbol{Y}_1, \dots, \boldsymbol{Y}_m)$ the number of infected individuals in the tests. There are $n\Delta$ edges connected to individuals, out of which exactly $k\Delta$ correspond to infected individuals. Each test chooses exactly Γ individuals without replacement, and hence, the number of infected individuals in any test follows a hypergeometric distribution. In order to get a handle on this distribution, we introduce a family $(\boldsymbol{X}_1, \dots, \boldsymbol{X}_m)$ of independent binomial variables, such that $\boldsymbol{X}_i \sim \mathrm{Bin}(\Gamma, k/n)$. These variables can accurately describe the local behaviour of how many infected individuals belong to test a_i. We define $\mathscr{E}_\Gamma$ to be the event that the overall number of edges containing infected individuals is correct, i.e.,

$$\mathscr{E}_\Gamma = \left\{\sum_{i=1}^m \boldsymbol{X}_i = k\Delta\right\}. \tag{63}$$

Claim 7.5 implies that $\mathbb{P}(\mathscr{E}_\Gamma) = \Omega((n\Delta)^{-1/2})$. In addition, we have the following.

Lemma 4.12: The sequence $(\boldsymbol{Y}_1, \dots, \boldsymbol{Y}_n)$ is identically distributed with $(\boldsymbol{X}_1, \dots, \boldsymbol{X}_n)$ given the event $\mathscr{E}_\Gamma$

Proof:

By the definition of $\boldsymbol{Y}_i$, we find for any $(y_i)_i$ satisfying $\sum_i y_i = k\Delta$ that

$$\begin{aligned} &\mathbb{P}(\boldsymbol{Y}_i = y_i, \forall i \in [m]) \\ &\quad = \binom{k\Delta}{y_1, \dots, y_m}\binom{(n-k)\Delta}{\Gamma - y_1, \dots, \Gamma - y_m}\binom{n\Delta}{\Gamma, \dots, \Gamma}^{-1} \\ &\quad = \frac{\prod_{i=1}^m \binom{\Gamma}{y_i}}{\binom{n\Delta}{k\Delta}}. \end{aligned}$$

where the equality follows by rewriting in terms of factorials and simplifying. Furthermore, given $\sum_i x_i = k\Delta$, we have

$$\mathbb{P}(\boldsymbol{X}_i = x_i, \forall\, i \in [m] | \mathscr{E}_\Gamma) = \prod_{i=1}^{m} \binom{\Gamma}{x_i} (k/n)^{x_i} (1-k/n)^{\Gamma - x_i} \left(\mathbb{P}(\mathscr{E}_\Gamma)\right)^{-1}.$$

Now, for two sequences $(y_i)_{i\in[m]}$ and $(y_i')_{i\in[m]}$ such that $\sum_{i=1}^m y_i = \sum_{i=1}^m y_i' = k\Delta$, we obtain

$$\frac{\mathbb{P}(\forall\, i \in [m] : \boldsymbol{Y}_i = y_i)}{\mathbb{P}(\forall\, i \in [m] : \boldsymbol{Y}_i = y_i')} = \prod_{i=1}^{m} \frac{\binom{\Gamma}{y_i}}{\binom{\Gamma}{y_i'}} = \frac{\mathbb{P}(\forall\, i \in [m] : \boldsymbol{X}_i = y_i | \mathscr{E}_\Gamma)}{\mathbb{P}(\forall\, i \in [m] : \boldsymbol{X}_i = y_i' | \mathscr{E}_\Gamma)}.$$

This implies the lemma. ∎

Thus, similarly to the analysis following Lemma 3.14, we are able to carry out all necessary calculations with respect to $(\boldsymbol{X}_1, ..., \boldsymbol{X}_n)$ and transfer the results to the original pooling scheme. For the next step, we need to get a handle on the number of positive and negative tests occurring in this setting. Let $\boldsymbol{m}_0 = \boldsymbol{m}_0(\mathscr{G}_\Gamma, \boldsymbol{\sigma})$ be the number of tests that render a negative result, and let $\boldsymbol{m}_1 = \boldsymbol{m}_1(\mathscr{G}_\Gamma, \boldsymbol{\sigma})$ be the number of tests that render a positive result. Then $\boldsymbol{m}_0$ and $\boldsymbol{m}_1$ are highly concentrated around their means as follows.

Lemma 4.13: With probability $1 - o(n^{-2})$, we have

$$\boldsymbol{m}_0 = \left(1 + n^{-\Omega(1)}\right) m (1-k/n)^\Gamma$$

and

$$\boldsymbol{m}_1 = \left(1 + n^{-\Omega(1)}\right) m \left(1 - (1-k/n)^\Gamma\right).$$

Proof: Recalling the definitions of $(\boldsymbol{Y}_i)_i$ and $(\boldsymbol{X}_i)_i$ from (63), we have

$$\boldsymbol{m}_0 = \sum_{i=1}^{m} \mathbf{1}\{\boldsymbol{Y}_i = 0\},$$

and we further denote by

$$\boldsymbol{m}_0' = \sum_{i=1}^{m} \mathbb{1}\{\boldsymbol{X}_i = 0\} \qquad \text{and} \qquad \boldsymbol{m}_1' = m - \boldsymbol{m}_0'$$

the number of negative and positive tests as modelled by the family of independent binomial variables $(\boldsymbol{X}_i)_i$. Clearly, as the $\boldsymbol{X}_i$ are mutually independent,

$$\mathbb{E}\left[\boldsymbol{m}_1'\right] = m \cdot (1 - \mathbb{P}(\mathrm{Bin}(\Gamma, k/n) = 0)) = m\left(1 - \left(1 - \frac{k}{n}\right)^\Gamma\right).$$

Observing that $\mathbb{E}\left[\boldsymbol{m}_1'\right] = \Theta(k)$ (since $m = \Theta(n)$ due to $\Gamma = \Theta(1)$), the Chernoff bound (Lemma 7.1) guarantees that

$$\mathbb{P}\left(\left|\boldsymbol{m}_1' - \mathbb{E}(\boldsymbol{m}_1')\right| > \sqrt{k}\ln(n) \mid \Gamma\right) = o(n^{-10})$$

and, similar to the proof of Lemma 3.15, by combining Lemma 4.12 with Claim 7.5, we obtain

$$\mathbb{P}\left(\left|\boldsymbol{m}_1 - \mathbb{E}(\boldsymbol{m}_1')\right| > \sqrt{k}\ln(n) \mid \Gamma\right) = o(n^{-8}).$$

Thus, the first part of the lemma follows. The second part is immediate, as $\boldsymbol{m}_0 + \boldsymbol{m}_1 = m$. ∎

The above-mentioned naive calculation (assuming independence) can now be rigorously justified, and we can establish the sizes of the disguised individuals w.h.p. as follows.

Lemma 4.14: Given n and $k = n^\theta$ as well as $\Gamma = \Theta(1)$ and $\Delta \geq 2$, we have w.h.p. that $|V_{0+}(\mathscr{G}_\Gamma)| = o(k)$.

Proof: By the definition of $\mathscr{G}_\Gamma$ via the configuration model, Lemma 4.13 guarantees that the total number of edges connected to a positive test is, with probability at least $1 - o(n^{-2})$, given by

$$\boldsymbol{m}_1\Gamma = \left(1 + O\left(n^{-\Omega(1)}\right)\right) m\Gamma\left(1 - (1-k/n)^\Gamma\right). \tag{64}$$

Let x be an uninfected individual. We can calculate the probability of x belonging to $V_{0+}(\mathscr{G}_\Gamma)$ (i.e., being disguised and uninfected) as follows: Each of the $\Delta = \Theta(1)$ edges[9] that are mapped to x in the configuration model have to be connected to a positive test. Thus, by (64) along with Claim 7.3, we obtain

$$\begin{aligned}\mathbb{P}\left(x \in V_{0+}(\mathscr{G}_\Gamma) \mid x \in V_0(\mathscr{G}_\Gamma), \boldsymbol{m}_1\right) &= \binom{\boldsymbol{m}_1\Gamma}{\Delta}\binom{m\Gamma}{\Delta}^{-1} = \left(1 + O\left(n^{-\Omega(1)}\right)\right)\left(1 - (1-k/n)^\Gamma\right)^\Delta \\ &= O\left(\left(\frac{k}{n}\right)^\Delta\right).\end{aligned}$$

Therefore,

$$\mathbb{E}\left[|V_{0+}(\mathscr{G}_\Gamma)|\right] = O\left((n-k)\left(\frac{k}{n}\right)^\Delta\right) = O\left(k\left(\frac{k}{n}\right)^{\Delta-1}\right) = o(k). \tag{65}$$

Combining (65), $\Delta \geq 2$, and Markov's inequality, we obtain the assertion of Lemma 4.14. ∎

Next, we define the event

$$\mathscr{F}_\Gamma = \left\{\boldsymbol{m}_1 = (1+o(1))\, m\left(1 - (1-k/n)^\Gamma\right)\right\} \cap \left\{|V_{0+}(\mathscr{G}_\Gamma)| = o(k)\right\}, \tag{66}$$

in which the number of positive tests and disguised uninfected individuals behave as expected. By Lemmas 4.13 and 4.14, we have $\mathbb{P}(\mathscr{F}_\Gamma) \geq 1 - o(1)$. We assume without loss of generality that the first $\boldsymbol{m}_1$ tests render a positive result.

Letting

$$\mathscr{D}_\Gamma = \left\{\sum_{i=1}^{\boldsymbol{m}_1} \boldsymbol{H}_i^1 = k\Delta, \quad \sum_{i=1}^{\boldsymbol{m}_1} \boldsymbol{H}_i^{0+} = |V_{0+}(\mathscr{G}_\Gamma)|\,\Delta\right\}$$

be the event that $\boldsymbol{H} = \sum_{i=1}^{\boldsymbol{m}_1} \boldsymbol{H}_i$ equals its expectation, we have the following analog of Corollary 3.18.

Claim 4.15: The distribution of $\boldsymbol{R}_i$ equals the distribution of $\boldsymbol{H}_i$ given $\mathscr{D}_\Gamma$ and Γ, and furthermore, $\mathbb{P}(\mathscr{D}_\Gamma) = \Omega(n^{-1})$.

Proof of Claim 4.15: Let $(r_i)_{i\in[\boldsymbol{m}_1]}$ be a sequence such that $r_i = (r_i^1, r_i^{0+}, r_i^{0-})$ and $\sum_i r_i^1 = k\Delta, \sum_i r_i^{0+} = |V_{0+}(\mathscr{G}_\Gamma)|\,\Delta$, and $r_i^{0-} = \Gamma - r_i^1 - r_i^{0+}$. Let

$$S_1 = k\Delta, \qquad S_{0+} = \Delta\,|V_{0+}(\mathscr{G}_\Gamma)| \qquad \text{and}$$
$$S_{0-} = n\Delta - n\Delta(1-(1-k/n)^\Gamma) - k\Delta.$$

[9] By counting degrees, we have $n\Delta = m\Gamma$, so the assumption $\Gamma = \Theta(1)$ leads to $m = \Theta(n\Delta)$. Since Δ is integer-valued and we are considering $m > 0$ and $m \leq n$ (otherwise, individual testing would be preferred), it follows that $\Delta = \Theta(1)$.

By the definition of $\boldsymbol{R}_i$, we have

$$\begin{aligned}&\mathbb{P}(\forall i\in[\boldsymbol{m}_1]:\boldsymbol{R}_i=r_i\,|\,|V_{0+}(\mathscr{G}_\Gamma)|,\boldsymbol{m}_1)\\&=\frac{\binom{S_1}{r_1^1\ldots r_{\boldsymbol{m}_1}^1}\binom{S_{0+}}{r_1^{0+}\ldots r_{\boldsymbol{m}_1}^{0+}}\binom{S_{0-}}{\Gamma-r_1^1-r_1^{0+}\ldots\Gamma-r_{\boldsymbol{m}_1}^1-r_{\boldsymbol{m}_1}^{0+}}}{\binom{n\Delta}{\Gamma,\ldots,\Gamma}}\\&=\binom{n\Delta}{S_1,S_{0+},S_{0-}}^{-1}\prod_{i=1}^{\boldsymbol{m}_1}\binom{\Gamma}{r_i^1,r_i^{0+},r_i^{0-}}.\end{aligned}$$

Letting $(r_i^{'})_{i\in[\boldsymbol{m}_1]}$ be a second sequence as above, it follows that

$$\frac{\mathbb{P}(\forall i\in[\boldsymbol{m}_1]:\boldsymbol{R}_i=y_i\,|\,|V_{0+}(\mathscr{G}_\Gamma)|,\boldsymbol{m}_1)}{\mathbb{P}(\forall i\in[\boldsymbol{m}_1]:\boldsymbol{R}_i=y_i^{'}\,|\,|V_{0+}(\mathscr{G}_\Gamma)|,\boldsymbol{m}_1)}=\prod_{i=1}^{\boldsymbol{m}_1}\frac{\binom{\Gamma}{r_i^1r_i^{0+}r_i^{0-}}}{\binom{\Gamma}{(r')_i^1(r')_i^{0+}(r')_i^{0-}}}.\tag{67}$$

Furthermore, by the definition of $\boldsymbol{X}$, we have

$$\begin{aligned}&\frac{\mathbb{P}(\forall i\in[\boldsymbol{m}_1]:\boldsymbol{H}_i=r_i\,|\,|V_{0+}(\mathscr{G}_\Gamma)|,\boldsymbol{m}_1,\mathscr{D}_\Gamma)}{\mathbb{P}(\forall i\in[\boldsymbol{m}_1]:\boldsymbol{H}_i=r_i^{'}\,|\,|V_{0+}(\mathscr{G}_\Gamma)|,\boldsymbol{m}_1,\mathscr{D}_\Gamma)}\\&=\frac{\left(\frac{k}{n}\right)^{\sum_{i=1}^{\boldsymbol{m}_1}r_i^1}\left(\frac{|V_{0+}(\mathscr{G}_\Gamma)|}{n}\right)^{\sum_{i=1}^{\boldsymbol{m}_1}r_i^{0+}}\left(\frac{n-k-|V_{0+}(\mathscr{G}_\Gamma)|}{n}\right)^{\sum_{i=1}^{\boldsymbol{m}_1}r_i^{0-}}}{\left(\frac{k}{n}\right)^{\sum_{i=1}^{\boldsymbol{m}_1}(r')_i^1}\left(\frac{|V_{0+}(\mathscr{G}_\Gamma)|}{n}\right)^{\sum_{i=1}^{\boldsymbol{m}_1}(r')_i^{0+}}\left(\frac{n-k-|V_{0+}(\mathscr{G}_\Gamma)|}{n}\right)^{\sum_{i=1}^{\boldsymbol{m}_1}(r')_i^{0-}}}\\&\cdot\prod_{i=1}^{\boldsymbol{m}_1}\frac{\binom{\Gamma}{r_i^1,r_i^{0+},r_i^{0-}}}{\binom{\Gamma}{(r')_i^1,(r')_i^{0+},)(r')_i^{0-}}}=\prod_{i=1}^{\boldsymbol{m}_1}\frac{\binom{\Gamma}{r_i^1,r_i^{0+},r_i^{0-}}}{\binom{\Gamma}{(r')_i^1,(r')_i^{0+},(r')_i^{0-}}}.\end{aligned}\tag{68}$$

The first part of the claim follows from Equations (67) and (68). The probability follows by applying Claim 7.5 for $\Delta=\Theta(1)$ ∎

We are interested in the number of positive tests that contain exactly one infected individual and no elements of $V_{0+}(\mathscr{G}_\Gamma)$. Therefore, we define

$$\boldsymbol{B}=\sum_{i=1}^{\boldsymbol{m}_1}\mathbf{1}\left\{\boldsymbol{R}_i^1+\boldsymbol{R}_i^{0+}=1\right\}\quad\text{and}\quad\boldsymbol{B}'=\sum_{i=1}^{\boldsymbol{m}_1}\mathbf{1}\left\{\boldsymbol{H}_i^1+\boldsymbol{H}_i^{0+}=1\right\}.$$

Claim 4.16: We have w.h.p. that

$$\boldsymbol{B}\le\Delta k\left(1-O\left(\Gamma n^{-(1-\theta)}\right)\right)$$

Proof of Claim 4.16: We use Claim 4.15 to simulate $\boldsymbol{B}$ through independent random variables as in $\boldsymbol{B}'$. Since $\boldsymbol{B}'$ is a sum of independent multinomial variables, we obtain its expectation by applying (66), Lemma 7.2 and Bayes Theorem:

$$\begin{aligned}&\mathbb{E}\left[\boldsymbol{B}'\,|\,|V_{0+}(\mathscr{G}_\Gamma)|,\boldsymbol{m}_1\right]\\&=\sum_{i=1}^{\boldsymbol{m}_1}\mathbb{P}\left(\boldsymbol{H}_i=(1,0,\Gamma-1)\,|\,|V_{0+}(\mathscr{G}_\Gamma)|\right)\\&=\boldsymbol{m}_1\Gamma\frac{k/n\cdot(1-(k+|V_{0+}(\mathscr{G}_\Gamma)|)/n)^{\Gamma-1}}{1-(1-k/n)^\Gamma}\\&=\left(1+O\left(\Gamma\frac{k}{n}\right)\right)\boldsymbol{m}_1\left(1-\frac{k+|V_{0+}(\mathscr{G}_\Gamma)|}{n}\right)^{\Gamma-1},\end{aligned}\tag{69}$$

where the last step follows from Lemma 7.4 and $\Gamma=\Theta(1)$. Conditioning on $\mathscr{F}_\Gamma$ defined in (66), we obtain

$$\begin{aligned}&\mathbb{E}\left[\boldsymbol{B}'\,|\,\mathscr{F}_\Gamma\right]\\&=\left(1+O\left(\frac{\Gamma k}{n}\right)\right)\frac{m\Gamma k}{n}\cdot\left(1-\frac{k+o(k)-O\left(n^{-\Omega(1)}\right)}{n}\right)^{\Gamma-1}\\&=\left(1+O\left(\frac{\Gamma k}{n}\right)\right)\frac{m\Gamma k}{n}\cdot\left(1-(\Gamma-1)\left(\frac{k+o(k)-O\left(n^{-\Omega(1)}\right)}{n}\right)\right)\\&=\left(1+O\left(\frac{\Gamma k}{n}\right)\right)\Delta k\left(1-(\Gamma-1)n^{-(1-\theta)}-o\left(n^{-(1-\theta)}\right)\right)\\&=\Delta k\left(1+O\left(\Gamma n^{-(1-\theta)}\right)\right),\end{aligned}\tag{70}$$

where the first line uses Lemma 4.14, the second line uses Claim 7.4 , and we additionally recall that $k=n^\theta$, $\Delta=\frac{m\Gamma}{n}$, and $\Gamma=\Theta(1)$. Moreover, since $\boldsymbol{B}'$ is a binomial random variable, the Chernoff bound (Lemma 7.1) yield with probability $o(n^{-10})$ that

$$\boldsymbol{B}'\le\Delta k\left(1+O\left(\Gamma n^{-(1-\theta)}\right)\right).$$

Thus, similar to the proof of Lemma 3.15, by Claim 4.15 we have w.h.p. that

$$\boldsymbol{B}\le\Delta k\left(1+O\left(\Gamma n^{-(1-\theta)}\right)\right).\tag{71}$$

∎

We are now in a position to characterize $\boldsymbol{A}=|V_1(\mathscr{G})\setminus V_{1--}(\mathscr{G}_\Gamma)|$.

Claim 4.17: Given $\boldsymbol{B}\le\Delta k\left(1-O\left(\Gamma n^{-(1-\theta)}\right)\right)$, we have for some constant $C>0$ that

$$\mathbb{E}\left[\boldsymbol{A}\,|\,\boldsymbol{B},\mathscr{F}_\Gamma\right]=k\binom{k\Delta-\boldsymbol{B}}{\Delta}\binom{k\Delta}{\Delta}^{-1}\le k(C\cdot\Gamma)^\Delta n^{-(1-\theta)\Delta}\tag{72}$$

Proof of Claim 4.17: The combinatorial expression follows by adding k probabilities, one per defective item. Each probability is the probability that an infected individual does not belong to V_{1--}, which equals the probability that all of its Δ connections are disjoint from the $k\Delta-\boldsymbol{B}$ connections to tests in which it would have been the only infected individual with no disguised uninfected individuals. The assertion then follows by combining the assumption $\boldsymbol{B}\le\Delta k\left(1-O\left(\Gamma n^{-(1-\theta)}\right)\right)$ with Claim 7.3. ∎

Proof of Lemma 4.11: We distinguish between $\theta/(1-\theta)\notin\mathbb{Z}$ and $\theta/(1-\theta)=T\in\mathbb{Z}$, and recall m_{DD} from (62) with $\Delta=\max\left\{2,1+\left\lfloor\frac{\theta}{1-\theta}\right\rfloor\right\}$. For simplicity, we assume that the inequality $m\ge m_{DD}$ holds with equality, but the general case is analogous.

Case A: $\theta/(1-\theta)\notin\mathbb{Z}$**.** In this case, for $m=m_{DD}$, we have $\Delta=\max\left\{2,1+\left\lfloor\frac{\theta}{1-\theta}\right\rfloor\right\}=\max\{2,\lceil\theta/(1-\theta)\rceil\}$. We distinguish the two cases $\theta<1/2$ and $\theta>1/2$ as follows:

- **Case A1:** $\theta>1/2$**.** In this case, we have $\Delta=\lceil\theta/(1-\theta)\rceil$. Defining $\eta=\theta-(1-\theta)\cdot\lceil\theta/(1-\theta)\rceil<0$, using (72) and $\Gamma,\Delta=\Theta(1)$, we find
$$\mathbb{E}\left[\boldsymbol{A}\,|\,\boldsymbol{B},\mathscr{F}_\Gamma\right]\le O(1)n^{\theta-(1-\theta)\cdot\lceil\theta/(1-\theta)\rceil}=O(n^\eta).\tag{73}$$
- **Case A2:** $\theta<1/2$**.** In this case, we have $\Delta=2$, and hence
$$\mathbb{E}\left[\boldsymbol{A}\,|\,\boldsymbol{B},\mathscr{F}_\Gamma\right]\le O(1)\Gamma^\Delta n^{3\theta-2}\le o(1).\tag{74}$$

Case B: $\theta/(1-\theta) = T \in \mathbb{Z}$**.** Again, we distinguish the cases $\theta = 1/2$ and $\theta > 1/2$:

- **Case B1:** $\theta > 1/2$**.** We have $\Delta = T+1$, so by (72) and $\Gamma, \Delta = \Theta(1)$, we find

$$\mathbb{E}[\boldsymbol{A} \mid \boldsymbol{B}, \mathscr{F}_\Gamma] \le O(1) n^{\theta-(1-\theta)\cdot(T+1)} = O(n^{-(1-\theta)}), \quad (75)$$

  where the last step uses $1-\theta \ge \theta$ and $T > 1$.
- **Case B2:** $\theta = 1/2$**.** We have $\Delta = 2$, and hence

$$\mathbb{E}[\boldsymbol{A} \mid \boldsymbol{B}, \mathscr{F}_\Gamma] \le O(n^{-1/2}). \quad (76)$$

Combining (73)–(76) with Markov's inequality and the fact that $\mathscr{F}_\Gamma$ occurs w.h.p., we deduce that $\boldsymbol{A} = 0$ w.h.p., completing the proof of Lemma 4.11. ∎

Theorem 4.10 now follows directly by combining Lemma 4.11 and Claim 2.4. So far, we have addressed the case where the test design is formed using the configuration model, and showed that the DD-algorithm is optimal in this regime if applied to the random regular pooling scheme $\mathscr{G}_\Gamma$. However, the preceding analysis does not provide a tight bound for the matching-based design.

E. Algorithmic feasibility II: Matching-based model

Recall from from Section II-B2 that the matching-based model with parameter γ is denoted by $\mathscr{G}^*_\Gamma$. While the DD algorithm does not appear to be optimal in this case, it turns out that turning to SCOMP (a slight refinement of DD) suffices for optimality.

Theorem 4.18: If $m \ge 2n/(\Gamma+1)$ and $0 < \theta < 1/2$, then w.h.p. SCOMP recovers $\boldsymbol{\sigma}$ from $\mathscr{G}^*_\Gamma$ and $\hat{\boldsymbol{\sigma}}$.

1) Proof of Theorem 4.18: We prove the theorem for $m = 2n/(\Gamma+1)$ (which implies $\gamma = \frac{2}{\Gamma+1}n$), but the more general case follows analogously; intuitively, a higher number of tests can only help. We analyse the DD algorithm on $\mathscr{G}^*_\Gamma$ in two steps, starting with the regular part of the graph. Denote by $\mathscr{G}^{*,r}_\Gamma$ the $(\Gamma-1,2)$ regular part, in which we select $n-\gamma$ individuals and pool them into two tests each. Denote by $\boldsymbol{\sigma}[\mathscr{G}^{*,r}_\Gamma]$ and $\hat{\boldsymbol{\sigma}}[\mathscr{G}^{*,r}_\Gamma]$ the infection status vector and outcome vector resulting from the regular part alone.

Lemma 4.19: If $m \ge 2n/(\Gamma+1)$, then w.h.p. DD recovers $\boldsymbol{\sigma}[\mathscr{G}^{*,r}_\Gamma]$ from $(\mathscr{G}^{*,r}_\Gamma, \hat{\boldsymbol{\sigma}}[\mathscr{G}^{*,r}_\Gamma])$ correctly.

Proof: This follows from Theorem 4.10, as $\mathscr{G}^{*,r}_\Gamma$ is identically distributed with $\mathscr{G}_{\Gamma-1}$ therein. With $\gamma = \frac{2}{\Gamma+1}n$ individuals removed from the population, we have $n' = \frac{\Gamma-1}{\Gamma+1}n$ individuals being tested in $\mathscr{G}^{*,r}_\Gamma$. Thus, we require at most $m' = 2\frac{n'}{\Gamma-1} = 2\frac{n}{\Gamma+1}$ tests in order for DD to succeed w.h.p. on $\mathscr{G}^{*,r}_\Gamma$. ∎

It remains to handle the second step, and specifically, argue that after adding the $\gamma = 2\frac{n}{\Gamma+1}$ individuals (one to each test) we can guarantee the success of SCOMP. We denote by k' the number of infected individuals under the remaining n' individuals, and let θ' be the value such that $k' = \Theta((n')^{\theta'})$, which is well-defined due to the following.

Claim 4.20: Under the preceding setup, we have w.h.p. that $\theta' = \theta$.

Proof: As we remove $\gamma = \frac{2}{\Gamma+1}n$ individuals randomly, the number of infected individuals in the remaining part is a hypergeometrically distributed random variable $\boldsymbol{K}' \sim H(n,k,n')$. Thus, the Chernoff bound for the hypergeometric distribution guarantees w.h.p. that

$$\boldsymbol{K}' = (1+o(1))kn'/n = (1+o(1))\frac{\Gamma-1}{\Gamma+1}k,$$

and the assertion follows. ∎

In the second step, we analyse the remaining part of the graph, in which the γ remaining individuals are placed into one test each. To do so, the following lemma turns out to be useful.

Lemma 4.21: Under the matching-based model $\mathscr{G}^*_\Gamma$ with $\theta < \frac{1}{2}$, it holds w.h.p. that there are no two infected individuals within distance 4 in the graph.

Proof: By construction, it holds with probability one that $\mathscr{G}^*_\Gamma$ has individual-degree at most two, and test-degree at most $\Gamma = \Theta(1)$. Hence, all degrees are bounded. This means that for any given individual x, the set of individuals x' with $\mathrm{dist}(x,x') \le 4$ has size $O(1)$. For any two individuals x and x', the probability of both being infected is $O((k/n)^2)$, and a union bound over the $O(n)$ possible pairs with $\mathrm{dist}(x,x') \le 4$ increases this probability to $O(n(k/n)^2)$. The assumption $\theta < \frac{1}{2}$ implies that $k = o(\sqrt{n})$, and thus, we have $O(n(k/n)^2) = o(1)$, which establishes the lemma. ∎

We now combine the preceding lemmas to establish the success of the DD algorithm.

Lemma 4.22: Conditioned on the DD algorithm recovering $\boldsymbol{\sigma}[\mathscr{G}^{*,r}_\Gamma]$ from $(\mathscr{G}^{*,r}_\Gamma, \hat{\boldsymbol{\sigma}}[\mathscr{G}^{*,r}_\Gamma])$, and on all infected individuals having pairwise distance exceeding 4, it holds with conditional probability one that the SCOMP algorithm recovers $\boldsymbol{\sigma}$ from $(\mathscr{G}^*_\Gamma, \hat{\boldsymbol{\sigma}})$.

Proof: By the construction of $\mathscr{G}^*_\Gamma$, there are $\gamma = \frac{2}{\Gamma+1}n$ individuals added to $\mathscr{G}^{*,r}_\Gamma$ to produce $\mathscr{G}^*_\Gamma$. Denote the set of these individuals by $X = \{x_1 \dots x_\gamma\}$. As $\gamma \le m$, there is a matching from X to the the m tests.

Having assumed success on the regular part $\mathscr{G}^{*,r}_\Gamma$, we only need to show that the newly added individuals in X are also correctly identified, and additionally do not impact the identifications in $\mathscr{G}^{*,r}_\Gamma$. Recall from Claim 2.4 that DD succeeds if and only if all infected individuals are easy infected (i.e., are in $V_{1--}(\mathscr{G}^*_\Gamma)$), and recall also that the success of DD implies the success of SCOMP [33]. We distinguish four different cases, which are illustrated in Figure 4.

Case A: Connecting to a negative test. Suppose that an individual $x \in X$ connects to a (previously) negative test a. Then, for all $y \in \partial_{\mathscr{G}^{*,r}_\Gamma}(a)$ we have $y \in V_{0-}(\mathscr{G}^{*,r}_\Gamma)$.

- **Case A-1:** $\boldsymbol{\sigma}_x = 0$**.** If x is uninfected and connects to a negative test, then the test remains negative. It follows immediately that $x \in V_{0-}(\mathscr{G}^*_\Gamma)$ (i.e., x is easy uninfected), which further implies that all other individuals in the test that were previously easy uninfected or easy infected in $\mathscr{G}^{*,r}_\Gamma$ remain so in $\mathscr{G}^*_\Gamma$, as desired.
- **Case A-2:** $\boldsymbol{\sigma}_x = 1$**.** In this case, we have $\hat{\boldsymbol{\sigma}}_a(\mathscr{G}^{*,r}_\Gamma) = 0$ but $\hat{\boldsymbol{\sigma}}_a(\mathscr{G}^*_\Gamma) = 1$. To maintain success, we need to show that all $y \in \partial_{\mathscr{G}^{*,r}_\Gamma}(a)$ (which were previously easy uninfected) remain easy uninfected in $\mathscr{G}^*_\Gamma$; this implies both that previous decisions are not affected, and that

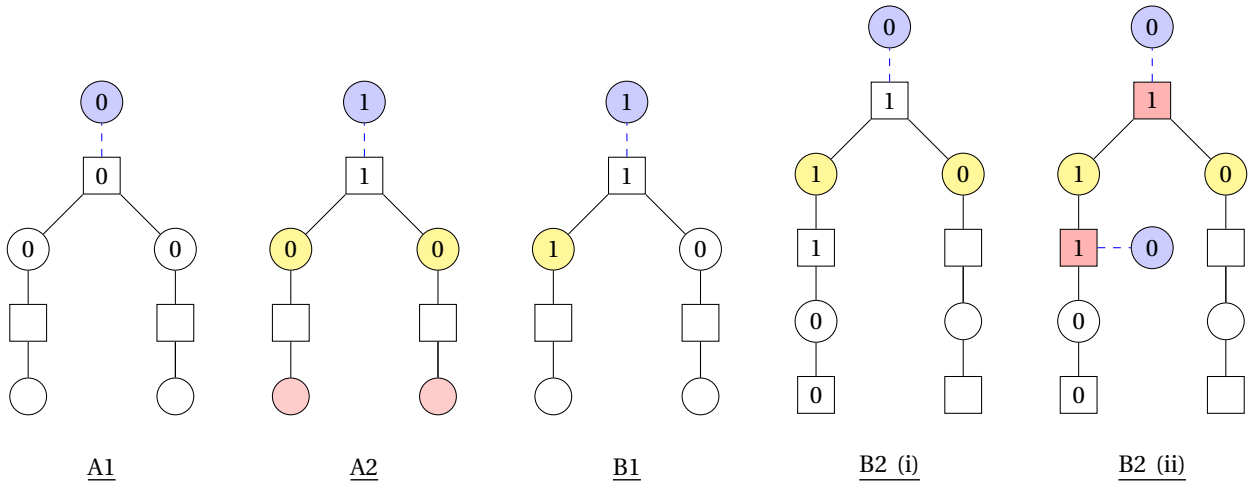

Fig. 4: The cases considered in our analysis. Round vertices are items, and square vertices are tests. The blue vertex is added in the second step of construction of $\mathscr{G}_\Gamma^*$, and the labels inside the vertices indicate the defectivity status or test outcome after adding the blue vertex. In case A1, recovery is clearly possible if and only if the same is true the remainder of the graph. In case A2, the yellow vertices may, in principle, no longer be identifiable as definite non-defectives. This happens if and only if the corresponding red individual is infected, which in turn implies a length-4 path between defectives, contradicting a high-probability event that we show. In case B1, there is a path of length two from the infected blue individual to the infected yellow individual, which is again a contradiction. In case B2(i), the infected yellow vertex can still be recovered as it is element of V_{1--} in the regular part and will be recovered successfully during the first two steps of SCOMP. In case B2(ii), the two red tests could either be explained by the yellow infected individual or by the two blue (uninfected) individuals, and due to its greedy selection rule, SCOMP declares the yellow individual as infected and the blue individuals as uninfected.

the decision for x is correct due to $x \in V_{1--}(\mathscr{G}_\Gamma^*)$. To establish that each $y \in \partial_{\mathscr{G}_\Gamma^{*,r}}(a)$ is easy uninfected, we argue that the second test that y belongs to is negative. Indeed, suppose for contradiction that y is in another positive test a' with an infected individual x'. Then, there is a path of length 4 in $\mathscr{G}_\Gamma^*$ from x to a to y to a' to x', and this contradicts Lemma 4.21.

Case B: Connecting to a positive test. Suppose that an individual $x \in X$ connects to a (previously) positive test a. Therefore, there exists at least one $y \in V_1(\mathscr{G}_\Gamma^{*,r}) \cap \partial_{\mathscr{G}_\Gamma^{*,r}}(a)$. As DD succeeds on $\mathscr{G}_\Gamma^{*,r}$ by assumption, we have $y \in V_{1--}(\mathscr{G}_\Gamma^{*,r})$.

- **Case B-1: $\boldsymbol{\sigma}_x = 1$.** This case does not occur, because it implies a length-2 path from x to y, both of which are infected, in contradiction with the lemma assumption.
- **Case B-2: $\boldsymbol{\sigma}_x = 0$.** Since the first two steps of SCOMP (Algorithm 1) never make mistakes, the only way that an error can occur in this case is that (i) x is added in some step of the final (sequential greedy) step, or (ii) $y \notin V_{1--}(\mathscr{G}_\Gamma^*)$ and y fails to be chosen throughout the final step. We argue that neither of these events occur. To see this, first note that in $\mathscr{G}_\Gamma^{*,r}$, y is not only part of $V_{1--}(\mathscr{G}_\Gamma^{*,r})$ because of a, but also because the second test that y belongs to consists only of y and individuals from $V_{0-}(\mathscr{G}_\Gamma^*)$: If this were not the case, then we could create a path from y to another infected individual using a path of at most 4. We then have the following:
  - If $y \in V_{1--}(\mathscr{G}_\Gamma^*)$ then y is trivially decoded correctly, and x is certainly not added in the final step (since its only test is already explained).
  - If $y \notin V_{1--}(\mathscr{G}_\Gamma^*)$ then the two tests containing y are unexplained at the start of the final step. Due to the above-established property of both of these tests leading to $y \in V_{1--}(\mathscr{G}_\Gamma^{*,r})$ in the regular part, we have that in $\mathscr{G}_\Gamma^*$, only y and/or the newly added elements of X can explain these two tests. But since y explains both of them, but the elements of X can only explain one each (since their degree is one), it is clearly y (and not x) that will be chosen, as desired.

■

We now have all the ingredients to prove Theorem 4.18.

Proof of Theorem 4.18: By construction, $\mathscr{G}_\Gamma^*$ consists of n individuals and $m = 2n/(\Gamma+1)$ tests. By Lemma 4.19, this m suffices for DD to succeed w.h.p. on the regular part of $\mathscr{G}_\Gamma^*$ (i.e., on $\mathscr{G}_\Gamma^{*,r}$). In addition, Lemma 4.21 gives the convenient distance-4 property w.h.p., and Lemma 4.22 guarantees that the preceding two findings suffice to ensure that SCOMP infers $\boldsymbol{\sigma}$ correctly from $\mathscr{G}_\Gamma^*$ and $\hat{\boldsymbol{\sigma}}$. Hence, the theorem follows. ■

F. Putting the pieces together

Theorem 4.10 proves that DD succeeds on the bi-regular graph $\mathscr{G}_\Gamma$ created by the configuration model using $\max\left\{2, 1+\left\lfloor \frac{\theta}{1-\theta} \right\rfloor\right\}$ tests, and hence so does SCOMP [33].

Furthermore, as Theorem 4.18 shows, for $\theta < 1/2$, $\frac{2n}{\Gamma+1}$ tests suffice employing $\mathscr{G}_\Gamma^*$ and using SCOMP.

Finally, we show that the results of Theorem 4.10 and Theorem 4.18 combine to match the information-theoretic lower bound (58), i.e., $\max\left\{\left(1+\left\lfloor\frac{\theta}{1-\theta}\right\rfloor\right)\frac{n}{\Gamma}, 2\frac{n}{\Gamma+1}\right\}$. On the one hand, for $\theta < \frac{1}{2}$, the lower bound simplifies to the desired quantity $\frac{2n}{\Gamma+1}$ due to the fact that $\left\lfloor\frac{\theta}{1-\theta}\right\rfloor = 0$ in this regime, and $\frac{2}{\Gamma+1} \geq \frac{1}{\Gamma}$ for $\Gamma \geq 1$. On the other hand, if $\theta \geq \frac{1}{2}$ then we have $\left\lfloor\frac{\theta}{1-\theta}\right\rfloor \geq 1$, and so the maximum in the lower bound is achieved by the first term (since $\frac{1}{\Gamma} \geq \frac{1}{\Gamma+1}$), thus again matching the upper bound. Hence, the SCOMP algorithm is information-theoretically optimal when used with the pooling scheme $\bar{\mathscr{G}}_\Gamma$.

V. Adaptive Group Testing with Δ-Divisible Individuals

In this section, we turn to adaptive testing strategies in the case of Δ-divisible individuals, and demonstrate that in certain cases the number of tests can be reduced significantly.

A. Converse

Recall that the converse bound proved in Theorem 3.1 already considered adaptive test designs. Thus, any adaptive strategy fails w.h.p.when $m \leq (1-\varepsilon)\mathrm{e}^{-1}\Delta k^{1+\frac{(1-\theta)}{\Delta\theta}}$ for fixed $\varepsilon > 0$.

B. Algorithm

We present an algorithm that can be viewed as an analog of Hwang's binary splitting algorithm [38], instead using *non-binary* splitting in order to ensure that each item is in at most Δ tests. Like with Hwang's algorithm, we assume that the size k of the infected set is known. In the case case that only an upper bound $k_{\max} \geq k$ is known, the same analysis and results apply with $k_{\max}$ in place of k. However, such bounds may somewhat loose, and care should be taken in using initial tests to estimate k as an initial step (e.g., see [39], [40], [41]), as this may use a significant portion of the Δ budget. For clarity, we only consider the case of known k in this section, and leave the case of unknown k to future work (see also [1] for some initial findings).

1) Recovering the infected Set: Our adaptive algorithm is described in Algorithm 2, where we assume for simplicity that $\left(\frac{n}{k}\right)^{1/\Delta}$ is an integer.[10] Using Algorithm 2, we have the following theorem, which is proved throughout the remainder of the subsection. We define

$$m_{\mathrm{ada}}(\Delta) = \Delta k^{1+\frac{1-\theta}{\theta\Delta}}. \tag{77}$$

Theorem 5.1: For $\Delta = o(\ln n)$ and $k = n^\theta$ with $\theta \in (0,1)$, the adaptive algorithm in Algorithm 2 tests each individual at most Δ times and uses at most $m_{\mathrm{ada}}(\Delta)(1+o(1))$ tests to recover the infected set exactly with zero error probability.

Require: Number of individuals n, number of infected individuals k, and divisibility of each individual Δ

Initialise $\tilde{n} \leftarrow \left(\frac{n}{k}\right)^{\frac{\Delta-1}{\Delta}}$ and the estimate $\widehat{\mathscr{K}} \leftarrow \emptyset$
Arbitrarily group the n individuals into $n/\tilde{n}$ groups of size $\tilde{n}$
Test each group and discard any that return negative
Label the remaining groups incrementally as $G_j^{(0)}$, where $j = 1, 2, \ldots$
for $i = 1$ to $\Delta - 1$ **do**
  for each group $G_j^{(i-1)}$ from the previous stage **do**
    Arbitrarily group all individuals in $G_j^{(i-1)}$ into $\tilde{n}^{1/(\Delta-1)}$ sub-groups of size $\tilde{n}^{1-i/(\Delta-1)}$
    Test each sub-group and discard any that return a negative outcome
    Label the remaining sub-groups incrementally as $G_j^{(i)}$
Add the individuals from all of the remaining singleton groups $G_j^{(\Delta-1)}$ to $\widehat{\mathscr{K}}$
return $\widehat{\mathscr{K}}$

Algorithm 2: Adaptive algorithm for Δ-divisible individuals

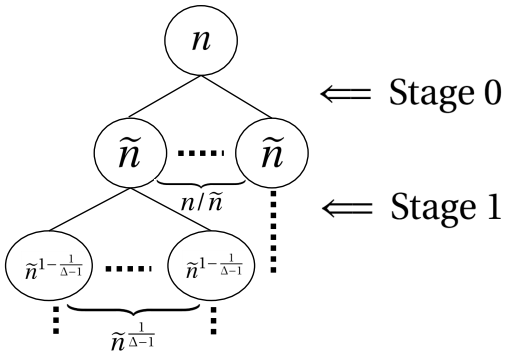

Fig. 5: Visualization of splitting in the adaptive algorithm.

Proof: Similar to Hwang's generalised binary splitting algorithm [38], the idea behind the parameter $\tilde{n}$ in Algorithm 2 is that when k becomes large, having large groups during the initial splitting stage is wasteful, as it results in each test having a high probability of being positive (not very informative). Hence, we want to find the appropriate group sizes that result in more informative tests to minimise the number of tests. Each stage (outermost for-loop in Algorithm 2) here refers to the process where all groups of the same sizes are split into smaller groups (e.g., see Figure 5). We let $\tilde{n}$ be the group size at the initial splitting stage of the algorithm. The algorithm first tests $n/\tilde{n}$ groups of size $\tilde{n}$ each,[11] then steadily decrease the sizes of each group down the stages: $\tilde{n} \to \tilde{n}^{1-1/(\Delta-1)} \to \tilde{n}^{1-2/(\Delta-1)} \to \cdots \to 1$ (see Figure 5). Hence, we have $n/\tilde{n}$ groups in the initial splitting and

[10] Note that we assume $k = o(n)$ and $\Delta = o\left(\ln\left(\frac{n}{k}\right)\right)$, meaning that $\left(\frac{n}{k}\right)^{1/\Delta} \to \infty$. Hence, the effect of rounding is asymptotically negligible, and is accounted for by the $1+o(1)$ term in Theorem 5.1.

[11] Note that $n/\tilde{n}$ is an integer for our chosen $\tilde{n}$ below, which gives $\frac{n}{\tilde{n}} = k\left(\frac{n}{k}\right)^{1/\Delta}$, and $\left(\frac{n}{k}\right)^{1/\Delta}$ was already assumed to be an integer.

Require: Number of individuals n, number of infected individuals k, and test size restriction Γ
Initialize infected set $\mathscr{K} \leftarrow \emptyset$
Randomly group n individuals into n/Γ groups of size Γ
for each group G_i where $i \in \mathbb{Z} : i \in [1, n/\Gamma]$ **do**
  while testing G_i returns a positive outcome **do**
    run Algorithm 4 on a copy of G_i, and add its one infected individual output k^* into $\mathscr{K}$
    $G_i \leftarrow G_i \setminus \{k^*\}$
return $\mathscr{K}$

Algorithm 3: Adaptive algorithm for Γ-sparse tests

$\tilde{n}^{\frac{1}{\Delta-1}}$ groups in all subsequent splits.

With the above observations, we can derive an upper bound on the total number of tests needed. We have $n/\tilde{n}$ tests in the first stage. Since we have k infected and split into $\tilde{n}^{\frac{1}{\Delta-1}}$ sub-groups in subsequent stages, the number of smaller groups that each stage can produce is at most $k\tilde{n}^{\frac{1}{\Delta-1}}$. This implies that the number of tests conducted at each stage is at most $k\tilde{n}^{\frac{1}{\Delta-1}}$, giving the following bound on m:

$$m \le \frac{n}{\tilde{n}} + (\Delta-1)k\tilde{n}^{\frac{1}{\Delta-1}}. \tag{78}$$

We optimise with respect to $\tilde{n}$ by differentiating the upper bound and setting it to zero. This gives $\tilde{n} = \left(\frac{n}{k}\right)^{\frac{\Delta-1}{\Delta}} = n^{\frac{(1-\theta)(\Delta-1)}{\Delta}}$, and substituting $\tilde{n} = \left(\frac{n}{k}\right)^{\frac{\Delta-1}{\Delta}}$ into the general upper bound in (78) gives the following upper bound:

$$m \le \frac{n}{(n/k)^{\frac{\Delta-1}{\Delta}}} + (\Delta-1)k\left(\left(\frac{n}{k}\right)^{\frac{\Delta-1}{\Delta}}\right)^{\frac{1}{\Delta-1}} = \Delta k\left(\frac{n}{k}\right)^{\frac{1}{\Delta}} = \Delta k^{1+\frac{1-\theta}{\theta\Delta}}. \tag{79}$$

■

Comparisons: We observe that $m_{\mathrm{ada}}(\Delta)$ matches the universal lower bound in Theorem 3.1 to within a factor of e for all $\theta \in (0,1)$. For $\theta < \frac{1}{2}$, we have $m_{\mathrm{ada}}(\Delta) = m_{\mathrm{DD}}(\Delta) = \Delta k^{1+\frac{(1-\theta)}{\Delta\theta}}$, meaning that the best known bounds for the adaptive and non-adaptive settings are identical (though the adaptive algorithm attains *zero* error probability). In contrast, for $\theta > \frac{1}{2}$, we have $m_{\mathrm{DD}}(\Delta) = \Delta k^{1+\frac{1}{\Delta}}$ and $m_{\mathrm{ada}}(\Delta) = \Delta k^{1+\frac{(1-\theta)}{\Delta\theta}}$. The former is significantly higher, and Theorem 3.2 reveals that this limitation is inherent to *any* non-adaptive test design and algorithm. Hence, for $\theta > \frac{1}{2}$, there is a significant gap between the number of tests required by adaptive and non-adaptive algorithms.

VI. Adaptive Group Testing with Γ-Sized Tests

Our adaptive algorithm with Γ-sparse tests, shown in Algorithm 3, is again a modification of Hwang's generalised binary splitting algorithm [38], where we initially divide the n individuals into $\frac{n}{\Gamma}$ groups of size Γ, instead of k groups of size $\frac{n}{k}$ as in the original algorithm.

Our main result is stated as follows, in which we define

$$m_{\mathrm{ada}}(\Gamma) = \frac{n}{\Gamma} + k\log_2\Gamma. \tag{80}$$

Require: a group of individuals $\tilde{G}$
while $\tilde{G}_i$ consists of multiple individuals **do**
  Pick half of the individuals in $\tilde{G}$ and call this set $\tilde{G}'$. Perform a single test on $\tilde{G}'$.
  If the test is positive, set $\tilde{G} \leftarrow \tilde{G}'$. Otherwise, set $\tilde{G} \leftarrow \tilde{G} \setminus \tilde{G}'$.
return single individual in $\tilde{G}$

Algorithm 4: Binary splitting

Theorem 6.1: For any $\Gamma = o\left(\frac{n}{k}\right)$, Algorithm 3 outputs the correct configuration of infection statuses with probability one, while using at most $m_{\mathrm{ada}}(\Gamma)(1+o(1))$ tests, each containing at most Γ items.

Proof: Let k_i be the number of infected individuals in each of the initial $\frac{n}{\Gamma}$ groups. Note that since $\Gamma = o\left(\frac{n}{k}\right)$ implies $k = o(\frac{n}{\Gamma})$, most groups will not have a infected individual. In the binary splitting stage of the algorithm, we can round the halves in either direction if they are not an integer. Hence, for each of the initial $\frac{n}{\Gamma}$ groups, we take at most $\lceil\log_2\Gamma\rceil$ adaptive tests to find a infected individual, or one test to confirm that there are no infected individuals. Therefore, for each of the initial $\frac{n}{\Gamma}$ groups, we need $\max\{1, k_i\log_2\Gamma + O(k_i)\}$ tests to find k_i infected individuals. Summing across all $\frac{n}{\Gamma}$ groups, we need a total of $m = \sum_{i=1}^{n/\Gamma}\max\{1, k_i\log_2\Gamma + O(k_i)\}$ tests. This has the following upper bound:

$$\begin{aligned} m &\le \frac{n}{\Gamma} + k\log_2\Gamma + O(k) \overset{(a)}{=} \frac{n}{\Gamma}(1+o(1)) + k\log_2\Gamma \\ &= m_{\mathrm{ada}}(\Gamma)(1+o(1)), \end{aligned} \tag{81}$$

where (a) uses $k = o\left(\frac{n}{\Gamma}\right)$. ■

If we slightly strengthen the requirement $\Gamma = o\left(\frac{n}{k}\right)$ to $\Gamma = o\left(\frac{n}{k\ln(n/k)}\right)$ (which, in particular, includes the regime $\Gamma = \left(\frac{n}{k}\right)^{1-\Omega(1)}$ studied in [25]), then we have $\frac{n}{\Gamma} = \omega\left(k\ln\left(\frac{n}{k}\right)\right)$ and hence $\frac{n}{\Gamma} = \omega(k\ln\Gamma)$. Thus, we obtain

$$m_{\mathrm{ada}}(\Gamma) = \frac{n}{\Gamma}(1+o(1)). \tag{82}$$

This simplified upper bound is tight, due the simple fact that $\frac{n}{\Gamma}(1-o(1))$ tests (of size at most Γ) are needed just to test a fraction $1-o(1)$ of the items at least once each (which is a minimal requirement for recovering $\boldsymbol{\sigma}$ w.h.p.). Formally, this argument reveals the following.

Theorem 6.2: In the setup of Γ-sparse tests with $k = n^\theta$ for some $\theta \in (0,1)$, any (possibly adaptive) group testing procedure that recovers $\boldsymbol{\sigma}$ w.h.p. must use at least $\frac{n}{\Gamma}(1-o(1))$ tests.

VII. Auxiliary Results

The following variant of the Chernoff bound is convenient to work with (e.g., see [42, Sec. 4.1]).

Lemma 7.1 (Multiplicative Chernoff Bound): Let $\boldsymbol{X}_1, \dots, \boldsymbol{X}_n$ be independent random variables such that $0 \le \boldsymbol{X}_i \le 1$ a.s., and fix $\delta \in (0,1)$. Then, we have

$$\mathbb{P}(|\boldsymbol{X} - \mathbb{E}[\boldsymbol{X}]| \ge \delta\mathbb{E}[\boldsymbol{X}]) \le 2\exp(-\delta^2\mathbb{E}[\boldsymbol{X}]/3).$$

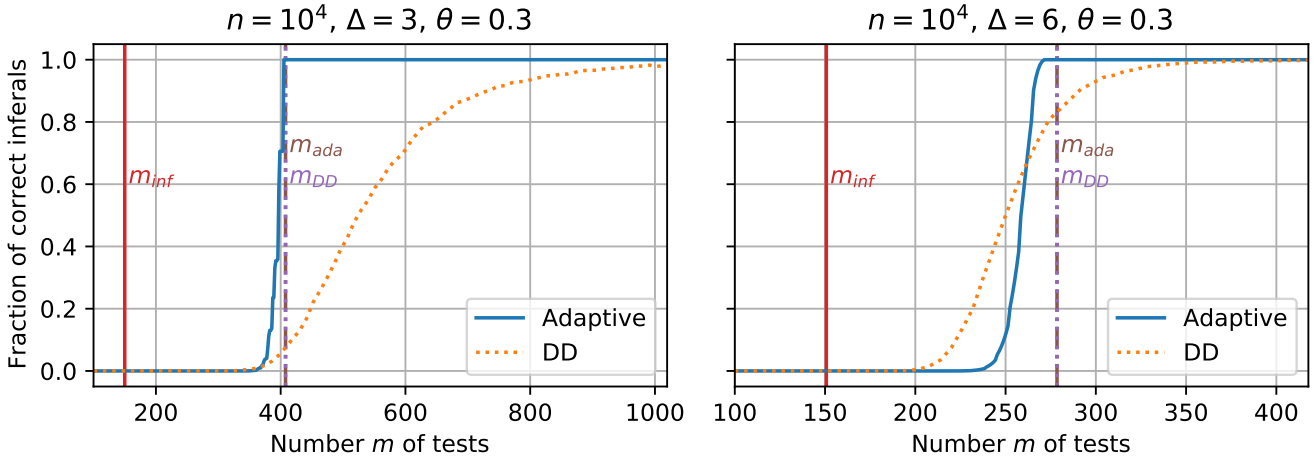

Fig. 6: Performance of adaptive and non-adaptive Δ-divisible algorithms as function of number of tests.

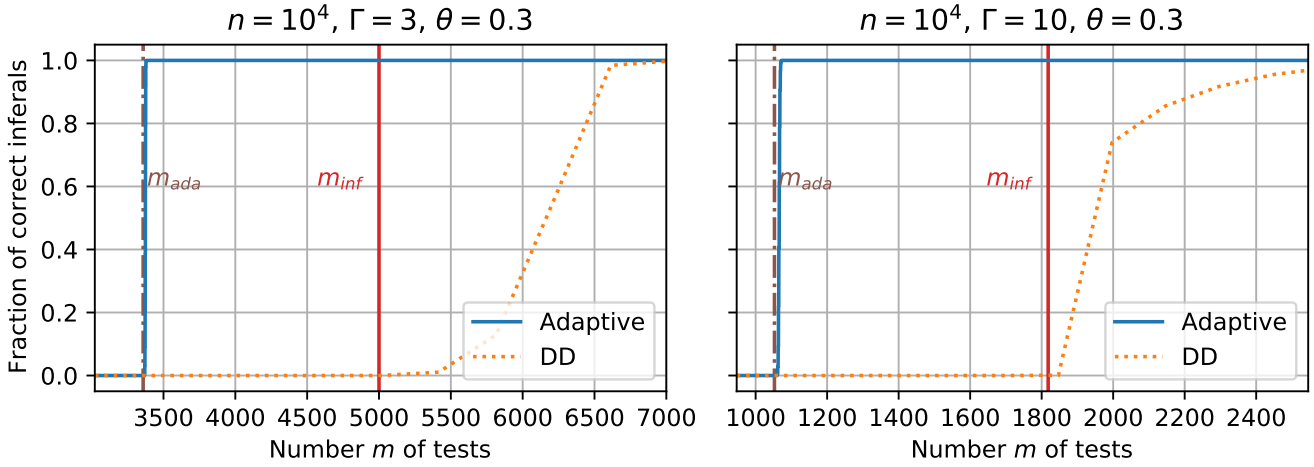

Fig. 7: Performance of adaptive and non-adaptive Γ-sparse algorithms as function of number of tests.

Lemma 7.2 (Stirling Approximation, [43]): We have for $n \to \infty$ that

$$n! = (1 + O(1/n))\sqrt{2\pi n} n^n \exp(-n).$$

Claim 7.3: Let $n > 0$, $\Delta = \ln^{O(1)} n$ be integers, and let $\alpha \in (0,1)$. Then

$$\binom{\alpha n}{\Delta}\binom{n}{\Delta}^{-1} = \left(1 + O\left(n^{-\Omega(1)}\right)\right)\alpha^{\Delta}.$$

Proof: By definition, we have

$$\frac{\binom{\alpha n}{\Delta}}{\binom{n}{\Delta}} = \frac{(\alpha n)!(n-\Delta)!}{n!(\alpha n - \Delta)!}.$$

Hence, applying Lemma 7.2 on each factor yields

$$\frac{\binom{\alpha n}{\Delta}}{\binom{n}{\Delta}} = (1 + O(n^{-1}))\exp(-\alpha n + (n-\Delta) - (\alpha n - \Delta) - n) \cdot (\alpha n)^{\alpha n}(n-\Delta)^{n-\Delta}(\alpha n - \Delta)^{-(\alpha n - \Delta)} n^{-n}\sqrt{\frac{(\alpha n)(n-\Delta)}{n(\alpha n - \Delta)}}. \quad (83)$$

As $\Delta = \ln^{O(1)} n$, we find that (83) equals

$$\frac{\binom{\alpha n}{\Delta}}{\binom{n}{\Delta}} = \left(1 + O\left(n^{-\Omega(1)}\right)\right)(\alpha n)^{\alpha n} n^n (\alpha n)^{-(\alpha n - \Delta)} n^{-n} = \left(1 + O\left(n^{-\Omega(1)}\right)\right)\alpha^{\Delta}, \quad (84)$$

and the assertion follows. ■

We also use the following direct consequence of the binomial expansion.

Claim 7.4: For any real number $x \geq -1$ and any integer $t \geq 0$ the following holds:

$$(1+x)^t = 1 + tx + O(t^2 x^2).$$

Finally, we state the following useful result relating to Stirling's approximation and the local limit theorem.

Claim 7.5: [Appendix B1 of [18]] For any $m, \Delta \in \mathbb{N}, \theta \in (0,1)$, $k \sim n^{\theta}$, let $(\boldsymbol{X}_i)_{i \in [m]}$ denote a sequence of independent $\mathrm{Bin}(\Gamma_i, k/n)$ and define

$$\mathscr{E} = \left\{ \sum_{i \in [m]} \boldsymbol{X}_i = k\Delta \right\}.$$

Then, we have $\mathbb{P}(\mathscr{E}) = \Omega(1/\sqrt{n\Delta})$.

VIII. Simulations

In Figures 6 and 7, we compare our theoretical findings to empirical results obtained as follows:

- In the non-adaptive case, we fix the number of individuals n, the infection parameter θ, and, depending on the setup considered, the individual degree Δ or test degree Γ. We vary the number m of tests (x-axis), and simulate 10^4 independent trials per parameter set. DD's performance (y-axis) is reported as the fraction of simulations per parameter point that inferred the infected set without errors.
- In the adaptive case, we cannot directly control the number of tests m *a priori*. Instead, we fix the same parameter set as in the non-adaptive case, and carry out 10^6 simulations. We then report the cumulative distribution of tests required, i.e., the y-value corresponding to some m is given as the fraction of runs that required at most m tests.

We observe that the empirical results are consistent with our theoretical thresholds in all cases. The adaptive testing strategies show a particularly rapid transition at $m_{\text{ada}}(\Delta)$ and $m_{\text{ada}}(\Gamma)$ respectively. We find that the non-adaptive DD algorithm requires more tests in comparison to the adaptive schemes, and has a much broader range of transient behaviour. This suggests that *convergence rates* to the first-order asymptotic threshold may reveal an even wider gap between adaptive and non-adaptive designs, in analogy with studies of channel coding [44]. Note that the change of slope in Figure 7 (right) at m=2000 is due to rounding of Δ.

IX. Conclusion

We have studied the information-theoretic and algorithmic thresholds of group testing with constraints on the number of items-per-test or test-per-item. For Δ-divisible items, we proved that at least for $\Delta = \omega(1)$, the DD algorithm is asymptotically optimal for $\theta > \frac{1}{2}$, and is optimal to within a factor of e for all $\theta \in (0,1)$, thus significantly improving on existing bounds for the COMP algorithm having suboptimal scaling laws. For Γ-sized tests with $\Gamma = \Theta(1)$, we improved on both the best known upper bounds and lower bounds, established a precise threshold for all $\theta \in (0,1)$, and introduced a new randomised test design for $\theta > \frac{1}{2}$. In both settings, we additionally provided near-optimal adaptive algorithms, and demonstrated a strict gap between the number of tests for adaptive and non-adaptive designs in broad scaling regimes.

Acknowledgments

OG was funded by DFG CO 646/3. MHK was partially funded by Stiftung Polytechnische Gesellschaft and DFG FOR 2975. OP was supported by the DFG (Grant PA 3513/1-1) and the London School of Economics and Political Science. MP was funded by ME 2088/4-2 and ME 2088/5-1 (DFG FOR 2975). JS was funded by an NUS Early Career Research Award.

References

[1] N. Tan and J. Scarlett, "Near-optimal sparse adaptive group testing," in *IEEE International Symposium on Information Theory (ISIT)*, 2020.

[2] R. Dorfman, "The detection of defective members of large populations," *Annals of Mathematical Statistics*, vol. 14, pp. 436–440, 1943.

[3] M. Aldridge, O. Johnson, and J. Scarlett, *Group testing: an information theory perspective*. Foundations and Trends in Communications and Information Theory, 2019.

[4] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick, "Optimal group testing," *Combinatorics, Probability and Computing*, pp. 1–38, 2021.

[5] D. Du and F. Hwang, *Combinatorial group testing and its applications*. Singapore: World Scientific, 1993.

[6] H. Kwang-Ming and D. Ding-Zhu, "Pooling designs and nonadaptive group testing: important tools for DNA sequencing," *World Scientific*, 2006.

[7] H. Ngo and D. Du, "A survey on combinatorial group testing algorithms with applications to dna library screening," *Discrete Mathematical Problems with Medical Applications*, vol. 7, pp. 171–182, 2000.

[8] R. Mourad, Z. Dawy, and F. Morcos, "Designing pooling systems for noisy high-throughput protein-protein interaction experiments using boolean compressed sensing," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 10, pp. 1478–1490, 2013.

[9] N. Thierry-Mieg, "A new pooling strategy for high-throughput screening: the shifted transversal design," *BMC Bioinformatics*, vol. 7, p. 28, 2006.

[10] I. Cheong, "The experience of south korea with covid-19," *Mitigating the COVID Economic Crisis: Act Fast and Do Whatever It Takes (CEPR Press)*, pp. 113–120, 2020.

[11] N. Madhav, B. Oppenheim, M. Gallivan, P. Mulembakani, E. Rubin, and N. Wolfe, "Pandemics: Risks, impacts and mitigation," *The World Bank: Disease control priorities*, vol. 9, pp. 315–345, 2017.

[12] E. C. for Disease Prevention and Control, "Surveillance and studies in a pandemic in europe," *ECDC Technical Report*, 2009.

[13] U. D. of Health and H. Services, "Pandemic influenza plan," *Planning and Preparedness Resources*, 2017.

[14] W. H. Origanisation, "Global surveillance during an influenza pandemic," *Global Influenza Program*, 2009.

[15] R. Benz, S. Swamidass, and P. Baldi, "Discovery of power-laws in chemical space," *Journal of Chemical Information and Modeling*, vol. 48, pp. 1138–1151, 2008.

[16] L. Wang, X. Li, Y. Zhang, and K. Zhang, "Evolution of scaling emergence in large-scale spatial epidemic spreading," *PLoS ONE*, vol. 6, 2011.

[17] M. Aldridge, "Individual testing is optimal for non-adaptive group testing in the linear regime," *IEEE Transactions on Information Theory*, vol. 65, p. 2058–2061, 2019.

[18] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick, "Information-theoretic and algorithmic thresholds for group testing," *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, vol. 132(43), pp. 1–14, 2019.

[19] O. Johnson, M. Aldridge, and J. Scarlett, "Performance of group testing algorithms with near-constant tests per item," *IEEE Transactions on Information Theory*, vol. 65, pp. 707–723, 2018.

[20] M. Aldridge, L. Baldassini, and O. Johnson, "Group testing algorithms: bounds and simulations," *IEEE Transactions on Information Theory*, vol. 60, pp. 3671–3687, 2014.

[21] J. Scarlett and V. Cevher, "Phase transitions in group testing," *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016)*, vol. 1, pp. 40–53, 2016.

[22] L. Wein and S. Zenios, "Pooled testing for HIV screening: Capturing the dilution effect," *Operations Research*, vol. 44, p. 543–569, 1996.

[23] E. S. S. Ciesek, "Pool testing of SARS-Cov-2 samples increases worldwide test capacities many times over," *https://www.bionity.com/en/news/1165636/pool-testing-of-sars-cov-02-samples-increases-worldwide-test-capacities-many-times-over.html, last accessed on 2020-04-08*, 2020.

[24] Y. Gefen, M. Szwarcwort-Cohen, and R. Kishony, "Pooling method for accelerated testing of covid-19," *https://www.technion.ac.il/en/2020/03/pooling-method-for-accelerated-testing-of-covid-19/*, 03/26/20.

[25] V. Gandikota, E. Grigorescu, S. Jaggi, and S. Zhou, "Nearly optimal sparse group testing," *IEEE Transactions on Information Theory*, vol. 65, no. 5, pp. 2760–2773, May 2019. [Online]. Available: https://doi.org/10.1109/tit.2019.2891651

[26] H. Inan, K. Kairouz, and A. Özgür, "Sparse group testing codes for low-energy massive random access," *55th Annual Allerton Conference*, vol. 1, pp. 658–665, 2017.
[27] A. Macula, "A simple construction of d-disjunct matrices with certain constant weights," *Discrete Mathematics*, vol. 162, pp. 311–312, 1996.
[28] C. Chan, P. Che, S. Jaggi, and V. Saligrama, "Non-adaptive probabilistic group testing with noisy measurements: near-optimal bounds with efficient algorithms," *49th Annual Allerton Conference on Communication, Control, and Computing*, vol. 1, pp. 1832–1839, 2011.
[29] H. A. Inan, P. Kairouz, and A. Özgür, "Sparse combinatorial group testing," *IEEE Transactions on Information Theory*, vol. 66, no. 5, pp. 2729–2742, 2020.
[30] S. Janson, T. Luczak, and A. Rucinski, *Random Graphs*. John Wiley and Sons, 2011.
[31] M. Aldridge, O. Johnson, and J. Scarlett, "Improved group testing rates with constant column weight designs," *IEEE Transactions on Information Theory*, vol. 65(2), pp. 1381–1385, 2016.
[32] L. Zdeborová and F. Krzakala, "Statistical physics of inference: thresholds and algorithms," *Advances in Physics*, vol. 65, no. 5, p. 453–552, Aug 2016. [Online]. Available: http://dx.doi.org/10.1080/00018732.2016.1211393
[33] M. Aldridge, "On the optimality of some group testing algorithms," in *IEEE International Symposium on Information Theory (ISIT)*, 2017.
[34] L. Baldassini, O. Johnson, and M. Aldridge, "The capacity of adaptive group testing," *Proc. ISIT*, vol. 1, pp. 2676–2680, 2013.
[35] R. Ash, *Information Theory*. Dover Publications Inc., New York, 1990.
[36] C. Fortuin, P. Kasteleyn, and J. Ginibre, "Correlation inequalities on some partially ordered sets," *Communications in Mathematical Physics*, vol. 22, pp. 89–103, 1971.
[37] B. Wu, "The weighted version of the handshaking lemma," *Journal of inequalities and application*, vol. 351, 2014.
[38] F. Hwang, "A method for detecting all defective members in a population by group testing," *Journal of the American Statistical Association*, vol. 67, pp. 605–608, 1972.
[39] P. Damaschke and A. Muhammad, "Competitive group testing and learning hidden vertex covers with minimum adaptivity," *Disc. Math., Algs. and Apps.*, vol. 2, no. 03, pp. 291–311, 2010.
[40] M. Falahatgar, A. Jafarpour, A. Orlitsky, V. Pichapati, and A. Suresh, "Estimating the number of defectives with group testing," in *IEEE Int. Symp. Inf. Theory*, 2016, pp. 1376–1380.
[41] N. Bshouty, V. Bshouty-Hurani, T. Hashem, and O. Sharafy, "Adaptive group testing algorithms to estimate the number of defectives," *Algorithmic Learning Theory*, 2018.
[42] R. Motwani and P. Raghavan, *Randomized Algorithms*. Chapman & Hall/CRC, 2010.
[43] A. J. Maria, "A remark on stirling's formula," *The American Mathematical Monthly*, vol. 72, no. 10, p. 1096, 1965.
[44] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Feedback in the non-asymptotic regime," *IEEE Transactions on Information Theory*, vol. 57, no. 8, pp. 4903–4925, 2011.

Oliver Gebhard studied Mathematics and Economics at Goethe University Frankfurt and University of Toronto. Currently, he is a PhD student under the supervision of Amin Coja-Oghlan.

Max Hahn-Klimroth is PostDoc at TU Dortmund University. He studied Mathematics and Computer Science at Goethe-University Frankfurt and obtained his PhD in Mathematics under the supervision of Amin Coja-Oghlan (Goethe-University Frankfurt) and Yury Person (TU Ilmenau).

Olaf Parczyk studied Mathematics at the Free University of Berlin and obtained his PhD at Goethe University Frankfurt under the supervision of Yury Person. He was a PostDoc at Technical University Ilmenau and the London School of Economics and Political Science.

Manuel Penschuck studied Computer Science at Goethe-University Frankfurt. He received his PhD in Computer Science from Goethe-University Frankfurt under the supervision of Ulrich Meyer.

Maurice Rolvien studied Mathematics at Johannes-Gutenberg University Mainz and Goethe-University Frankfurt. He is currently a PhD student in Mathematics under the supervision of Amin Coja-Oghlan.

Jonathan Scarlett (S'14 – M'15) received the B.Eng. degree in electrical engineering and the B.Sci. degree in computer science from the University of Melbourne, Australia. From October 2011 to August 2014, he was a Ph.D. student in the Signal Processing and Communications Group at the University of Cambridge, United Kingdom. From September 2014 to September 2017, he was post-doctoral researcher with the Laboratory for Information and Inference Systems at the École Polytechnique Fédérale de Lausanne, Switzerland. Since January 2018, he has been an assistant professor in the Department of Computer Science and Department of Mathematics, National University of Singapore. His research interests are in the areas of information theory, machine learning, signal processing, and high-dimensional statistics. He received the Singapore National Research Foundation (NRF) fellowship, and the NUS Presidential Young Professorship.

Nelvin Tan received the B.Comp. degree in computer science and statistics from the National University of Singapore, in 2021. He is currently pursuing the Ph.D. degree from the Signal Processing and Communications Group in the Department of Engineering, University of Cambridge. His research interests include information theory and statistical learning.

APPENDIX D. IMPROVED BOUNDS FOR NOISY GROUP TESTING WITH CONSTANT TESTS PER ITEM

IMPROVED BOUNDS FOR NOISY GROUP TESTING WITH CONSTANT TESTS PER ITEM

OLIVER GEBHARD, OLIVER JOHNSON, PHILIPP LOICK, MAURICE ROLVIEN

{Gebhard, Loick, Rolvien}@math.uni-frankfurt.de, *Goethe University, Mathematics Institute, 10 Robert Mayer St, Frankfurt 60325, Germany.*

O.Johnson@bristol.ac.uk, *University of Bristol, School of Mathematics, Woodland Road, Bristol, BS8 1UG, United Kingdom*

ABSTRACT. The group testing problem is concerned with identifying a small set of infected individuals in a large population. At our disposal is a testing procedure that allows us to test several individuals together. In an idealized setting, a test is positive if and only if at least one infected individual is included and negative otherwise. Significant progress was made in recent years towards understanding the information-theoretic and algorithmic properties in this noiseless setting. In this paper, we consider a noisy variant of group testing where test results are flipped with certain probability, including the realistic scenario where sensitivity and specificity can take arbitrary values. Using a test design where each individual is assigned to a fixed number of tests, we derive explicit algorithmic bounds for two commonly considered inference algorithms and thereby naturally extend the results of Scarlett & Cevher (2016) and Scarlett & Johnson (2020). We provide improved performance guarantees for the efficient algorithms in these noisy group testing models – indeed, for a large set of parameter choices the bounds provided in the paper are the strongest currently proved.

1. Introduction

1.1. **Motivation and background.** Suppose we have a large collection of n people, a small number k of whom are infected by some disease, and where only $m \ll n$ tests are available. In a landmark paper [16] from 1943, Dorfman introduced the idea of group testing. The basic idea is as follows: rather than screen one person using one test, we could mix samples from individuals in one pool, and use a single test for this whole pool. The task is to recover the infection status of all individuals using the pooled test results. Dorfman's original work was motivated by a biological application, namely identifying individuals with syphilis. Subsequently, group testing has found a number of related applications, including detection of HIV [51], DNA sequencing [29, 37] and protein interaction experiments [35, 49]. More recently, it has been recognised as an essential tool to moderate pandemic spread [12], where identifiying infected individuals fast and at a low cost is indispensable [32]. In particular, group testing has been identified as a testing scheme for the detection of COVID-19 [2, 17, 21]. From a mathematical perspective, group testing is a prime example of an inference problem where one wants to learn a ground truth from (possibly noisy) measurements [1, 8, 15]. Over the last decade, it has regained popularity and a significant body of research was dedicated to understand its information-theoretic and algorithmic properties [9, 13, 14, 44, 45, 46]. In this paper, we provide improved upper bounds on the number of tests that guarantee successful inference for the noisy variant of group testing.

1.2. **Related Work.**

1.2.1. *Noiseless Group Testing.* In the simplest version of group testing, we suppose that a test is positive if and only if the pool contains at least one infected individual. We refer to this as the noiseless case. In this setting, each negative test guarantees that every member of the corresponding pool is not infected, so they can be removed from further consideration. However, a positive test only tells us that at least one item in the test is defective (but not which one), and so requires further investigation. Dorfman's original work [16] proposed a simple adaptive strategy where a small pool of individuals is tested, and where each positive test is followed up by testing every individual in the corresponding pool individually. Since then it has been an important problem to find the optimal way to recover the whole population's infection status in the noiseless case (see [7] for a detailed survey). A simple counting argument (see for example [7, Section 1.4]) shows that to ensure recovery with zero error probability, since every possible defective set must give different test outcomes, the following must hold in the noiseless setting:

$$2^m \geq \binom{n}{k} \qquad \Rightarrow \qquad m \geq m^0_{\mathrm{inf}} := \frac{1}{\log 2} k \log(n/k) \tag{1.1}$$

This can be extended to the case of recovery with small error probability, for example with the bound (see [7, Eq. (1.7)]) that the success probability

$$\mathbb{P}(\mathrm{suc}) \leq \frac{2^{\mathrm{m}}}{\binom{\mathrm{n}}{\mathrm{k}}}, \tag{1.2}$$

meaning that the success probability must decay exponentially with the number of tests below m^0_{inf}. Hwang [24] provided an algorithm based on repeated binary search, which is essentially optimal in terms of the number of tests required in that it requires $m^0_{\mathrm{inf}} + O(k)$ tests, but may require many stages of testing. The question of whether non-adaptive algorithms (or even adaptive algorithms with a limited number of stages) can attain the bound (1.1) remained open until recently. [4, 14] showed that the answer depends on the prevalence of the disease, for example on the value of $\theta \in (0,1)$ in a parameterisation[1] where the number of infected individuals $k \sim n^\theta$. Non-adaptive testing schemes can be represented through a

[1]The result of [14] is two-fold. On the one hand, it provides a method to recover infected individuals w.h.p.as well as attaining (1.1) for a certain range of $\theta < \theta^*$. On the other hand they show that (1.1) cannot be attained by any testing procedure for larger $\theta > \theta^*$. One finds $\theta^* = \log(2) \cdot (1 + \log(2))^{-1}$.

binary $(m \times n)$-matrix that indicates which individual participates in which test. Significant research was dedicated to see which design attains the optimal performance, although much of the recent research analysed the performance of randomized designs. Initial research focused on the case where the matrix entries are i.i.d. [3, 5, 46], which we will refer to as Bernoulli pooling. Later work considered a constant column design where each individual is assigned to a (near-)constant number of tests [6, 13, 14, 26]. Indeed [14] showed that such a design is information-theoretically optimal in the *noiseless* setting and it is to be expected that this remains true for the noisy case. To recover the ground truth from the test results and the pooling scheme, this paper focuses on two non-adaptive algorithms, COMP and DD, which are relatively simple to perform and interpret in the noiseless case. We describe them in more detail below, but in brief COMP [10] simply builds a list of all the individuals who ever appear in a negative test and are hence certainly healthy, and assumes that the other individuals are infected. DD [5] uses COMP as a first stage and builds on it by looking for individuals who appear in a positive test that only otherwise contains individuals known to be healthy. While the noiseless case provides an interesting mathematical abstraction, it is clear that it may not be realistic in practice [40].

1.2.2. *Noisy Group Testing.* In medical applications [42] the two occurring types of noise in a testing procedure are related to sensitivity (the probability that a test containing an infected individual is indeed positive) and specificity (the probability that a test with only healthy individuals is indeed negative), and in that language we cannot assume the gold standard of tests with unit specificity and sensitivity. Thus, research attention in recent years has shifted towards the noisy version of group testing [10, 43, 44, 46, 47, 48]. On the one hand, the *adaptive* noisy case was considered in [43, 44]. On the other hand [10, 27, 28, 33, 46, 47, 48] looked at the *non-adaptive* noise case from different angles (for instance linear programming, belief propagation, and Markov Chain Monte Carlo). In [46, 47, 48] the algorithmic performance guarantees within noisy group testing under Bernoulli pooling are discussed. First of all [46] obtained a converse as well as a theoretical achievability bound, but stated the practical recovery as an direction for further research. In the following [47, 48] shed light on this question by using Bernoulli pooling.[2] In this paper we focus on the COMP and DD algorithms, since it is possible to deduce explicit performance guarantees for them. The original COMP and DD were designed for the noiseless case and do not automatically carry over to general noisy models. However, recent work of Scarlett and Johnson [48] showed that noisy versions of these algorithms can perform well under certain noise models using i.i.d. (Bernoulli pooling) test designs, particularly focusing on Z channel and reverse Z channel noise. As common medical tests have different values for sensitivity and specificity [31] the analysis of a generalized noise model beyond the Z and reverse Z channel is warranted.

1.2.3. *Model Justification.* As described for example in pandemic plans developed by the EU, US and WHO [19, 38, 39], and in COVID-specific work [36], adaptive strategies may not be suitable for pandemic prevention. For example, if a test takes one day to prepare and for the results to be known, then each stage will require an extra day to perform, meaning that adaptive group testing information can be received too late to be useful. Hence the need to perform large-scale testing to identify infected individuals fast relative to the doubling time [12, 32, 36] can make adaptive group testing unsuitable to prevent an infectious disease from spreading. Furthermore it may be difficult to preserve virus samples in a usable state for long enough to perform multi-round testing [22]. Due to its automation potential and the fact that tests can be completed in parallel (for example by the use of 96-well PCR plates [18]), the main applications of group testing such as DNA screening [11, 29, 37], HIV testing [51] and protein interaction

[2] [47] introduced an approach based on separate decoding of items for symmetric noise models. While this approach works well for small θ (in particular $\theta \to 0$), the performance drops dramatically for larger θ. For most θ this approach is worse off than the noisy DD discussed in [48]. Note there exist some noise levels with the very strong restriction assuming $p = q$ where [47] improve over our results in the θ very close to 0 regime. Due to the generality of our model we will from now on focus on [48] as benchmark for our results.

analysis [35, 49] are non-adaptive, where all tests are specified upfront and performed in parallel. For example, while group testing strategies appear to be useful to identify individuals infected with COVID-19 (see for example [17, 21]), testing for the presence of the SARS-CoV-19 virus is not perfect [52], and so we need to understand the effect of both false positive and false negative errors in this context, with non-identical error probabilities. For this reason, we consider a general $p-q$ noise model in this paper. Under this model, a truly negative test is flipped with probability p to display a positive test result, while a truly positive test is flipped to negative with probability q (Figure 1). Its formulation is sufficiently general to accommodate the recovery of the noiseless results ($p = q = 0$), Z channel ($p = 0$), reverse Z channel ($q = 0$) and the Binary Symmetric Channel ($p = q$). However, our results include the case of non-zero p and q without having to make the somewhat artificial assumption that false negative and false positive errors are equally likely. We note that it may be unrealistic to assume that the noise parameters are known exactly, and more sophisticated models may be needed to understand the real world. Nevertheless our analysis of a generalised noise model serves as a starting point towards a full understanding of the difficulties occurring while implementing group testing algorithms in laboratories.

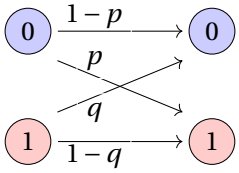

FIGURE 1. The $p-q$-noise model: the result of each standard noiseless group test is transmitted independently through the given noisy communication channel.

1.3. **Contribution.** This paper provides a simultaneous extension of [13] and [26, 48], by analysing noisy versions of COMP and DD under more general noise models for constant-column weight designs. In contrast to prior work [5, 26] assuming sampling with replacement, in this paper we use sampling without replacement, meaning that our designs have exactly the same number of tests for each item, rather than approximately the same as in those previous works. This makes little difference in practice, but may be closer to the spirit of LDPC codes for example.

We provide explicit bounds on the performance of these algorithms in a generalized noise model. We will prove that (noisy versions of) COMP as well as DD succeed with $\Theta(k\log(n/k))$ tests. Our analysis reveals the exact constants to ensure the recovery with these two inference algorithms. The main results will be stated formally in Theorems 2.1 and 2.2, but we would like to give the reader a first insight of what will follow. We analyze Algorithms 1 and 2 for the constant degree model, where there are $m = ck\log(n/k)$ tests performed and each individual chooses $\Delta = cd\log(n/k)$ tests uniformly at random. Let $p, q \geq 0, p+q < 1$ and $\epsilon > 0$.

We start with the performance of COMP (Algorithm 1), as stated in Theorem 2.1:

For any $\Delta := \Delta(c,d)$ we find a threshold $\alpha := \alpha(d,p,q)$ such that COMP succeeds in inferring the infected individuals if the number of tests

$$m \geq (1+\varepsilon)m_{COMP} = \min_{\alpha,d} \max\{b_1(\alpha,d), b_2(\alpha,d)\}\, k\log(n/k)$$

The next step on our agenda is the performance of DD (Algorithm 2), as stated in Theorem 2.1:

For any $\Delta := \Delta(c,d)$ we find thresholds $\alpha := \alpha(d,p,q)$ and $\beta := \beta(d,q)$ such that DD succeeds in inferring the infected individuals if the number of tests

$$m \geq (1+\varepsilon)m_{DD}(n,\theta,p,q) = \min_{\alpha,\beta,d} \max\big\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\alpha,\beta,d)\big\}\, k\log(n/k)$$

For all typical noise channels (Z, reverse Z and BSC) we compare the constant-column and Bernoulli design and find for all such instances that the required number of tests in the former is lower than the number needed in the latter thereby improving on results from [48], and providing the strongest performance guarantees currently proved for efficient algorithms in noisy group testing.

As group testing offers an essential tool for pandemic prevention [32] and as the the accuracy of medical testing is limited [31, 40] this paper provides the natural next step in the group testing literature.

1.4. **Test design and notation.** To formalize our notation, we write n for the number of individuals in the population, $\boldsymbol{\sigma}$ for a binary vector representing the infection status of each individual, k (the Hamming weight of $\boldsymbol{\sigma}$) for the number of infected individuals and m for the number of tests performed. We assume that k is known for the purposes of matrix design, though in practice (see [7, Remark 2.3]) it is generally enough to know k up to a constant factor to design a matrix with good properties. In this paper, in line with other work such as [5], we consider a scaling $k \sim n^\theta$ for some fixed $\theta \in (0,1)$, referred to in [7, Remark 1.1] as the sparse regime[3]. In addition to the interesting phase transitions observed using this scaling, this sparse regime is particularly relevant as it was found suitable to model the early state of a pandemic [50].

Let us next introduce the test design. With $V = (x_i)_{i\in[n]}$ denoting the set of n individuals[4] and $F = (a_i)_{i\in[m]}$ the set of m tests, the test design can be envisioned as a bipartite factor graph with n variable nodes "on the left" and m factor nodes "on the right". We draw a configuration $\boldsymbol{\sigma} \in \{0,1\}^V$, encoding the infection status of each individual, uniformly at random from vectors of Hamming weight k. The set of healthy individuals will be denoted by V_0 and the set of infected individuals by V_1. In symbols,

$$V_0 = \{x \in V : \boldsymbol{\sigma}(x) = 0\} \qquad \text{and} \qquad V_1 = V \setminus V_0 = \{x \in V : \boldsymbol{\sigma}(x) = 1\}$$

The lower bound from (1.1) suggests that in the noisy group testing setting it is natural to compare the performance of algorithms and matrix designs in terms of the prefactor of $k\log(n/k)$ in the number of tests required. To be precise, we carry out m tests, and each item is assigned to exactly Δ tests chosen uniformly at random without replacement. We parameterize m and Δ as

$$m = ck\log(n/k) \qquad \text{and} \qquad \Delta = cd\log(n/k) \tag{1.3}$$

for some suitably chosen constants $c, d \geq 0$.

Let ∂x denote the set of tests that individual x appears in and ∂a the set of individuals assigned to test a. The resulting (non-constant) collection of test degrees will be denoted by the vector $\boldsymbol{\Gamma} = (\boldsymbol{\Gamma}_a)_{a\in[m]}$. Further, let

$$\Gamma_{\min} = \min_{a\in[m]} \Gamma_a \qquad \text{and} \qquad \Gamma_{\max} = \max_{a\in[m]} \Gamma_a. \tag{1.4}$$

Throughout, $\boldsymbol{G} = \boldsymbol{G}(n,m,\Delta)$ describes the random bipartite factor graph from this construction.

Now consider the outcome of the tests. Recall from above that a standard noiseless group test a gives a positive result if and only if there is at least one defective item contained in the pool, or equivalently if $\sum_{x\in\partial a} \boldsymbol{\sigma}(x) \geq 1$. Even in the noisy case, this sum is a useful object to consider. Writing $\mathbf{1}$ for the indicator function, we define

$$\boldsymbol{\sigma}^*(a) = \mathbf{1}\left\{\sum_{x\in\partial a} \boldsymbol{\sigma}(x) \geq 1\right\} \tag{1.5}$$

to be the outcome we would observe in the noiseless case using the test matrix corresponding to $\boldsymbol{G}$. We will say that test a is *truly positive* if $\boldsymbol{\sigma}^*(a) = 1$ and truly negative otherwise.

However, we do not observe the values of $\boldsymbol{\sigma}^*(a)$ directly, but rather see what we will refer to as the *displayed* test outcomes $\hat{\boldsymbol{\sigma}}(a)$ – the outcomes of sending the true outcomes $\boldsymbol{\sigma}^*(a)$ independently through

[3] Note that the analysis directly extends to $k = \Theta(n^\theta)$ as a constant factor in front does not influence the analysis.

[4] $[n]$ will be used as an abbreviated notation for the set $\{1,\dots,n\}$.

the $p-q$ channel of Figure 1. Since in this model a truly positive test remains positive with probability $1-q$ and a truly negative test is displayed as positive with probability p we can write

$$\hat{\boldsymbol{\sigma}}(a) = \mathbf{1}\left\{\mathrm{Be}(p)=1\right\}\left(1-\boldsymbol{\sigma}^*(a)\right) + \mathbf{1}\left\{\mathrm{Be}(1-q)=1\right\}\boldsymbol{\sigma}^*(a) \tag{1.6}$$

where $\mathrm{Be}(r)$ denotes a Bernoulli random variable with parameter r independent of all other randomness in the model. For models with binary outputs, this is the most general channel satisfying the noisy defective channel property of [7, Definition 3.3], though more general models are possible under the only defects matter property [7, Definition 3.2], where the probability of a test being positive depends on the number of infected individuals it contains.

Note that if $p+q>1$, we can preprocess the outputs from (1.6) by flipping them, i.e. setting $\widetilde{p}=1-p$ and $\widetilde{q}=1-q$, where $\widetilde{p}+\widetilde{q}<1$. Hence without loss of generality we will assume throughout that $p+q<1$. In the case $p+q=1$, the test outcomes are independent of the inputs, and we cannot hope to find the infected individuals – see Corollary 2.3.

With $\boldsymbol{m}_0$ being the number of truly negative tests, let $\boldsymbol{m}_0^f$ be the number of truly negative tests that are flipped to display a positive test result and $\boldsymbol{m}_0^u$ be the number of truly negative tests that are unflipped. Similarly, define $\boldsymbol{m}_1$ as the number of truly positive tests, of which $\boldsymbol{m}_1^f$ are flipped to a negative test result and of which $\boldsymbol{m}_1^u$ are unflipped. For reference, for $t\in\{0,1\}$ we write

$$\boldsymbol{m}_t = \left|\left\{a:\boldsymbol{\sigma}^*(a)=t\right\}\right|$$
$$\boldsymbol{m}_t^f = \left|\left\{a:\boldsymbol{\sigma}^*(a)=t,\hat{\boldsymbol{\sigma}}(a)\neq t\right\}\right| \quad \text{and} \quad \boldsymbol{m}_t^u = \left|\left\{a:\boldsymbol{\sigma}^*(a)=t,\hat{\boldsymbol{\sigma}}(a)=t\right\}\right|$$

Here we use bold letters to indicate random variables. Throughout the paper, we use the standard Landau notation $o(\cdot),O(\cdot),\Theta(\cdot),\Omega(\cdot),\omega(\cdot)$ and define $0\log 0=0$. Furthermore we say that a property $\mathscr{P}$ holds *with high probability (w.h.p.)*, if $\mathbb{P}(\mathscr{P})=1$ as $n\to\infty$. In order to quantify the performance of our algorithms, for any $0<r\neq s<1$, we write

$$D_{\mathrm{KL}}(r\|s) := r\log\left(\frac{r}{s}\right)+(1-r)\log\left(\frac{1-r}{1-s}\right), \tag{1.7}$$

for the relative entropy of a Bernoulli random variable with parameter r to a Bernoulli random variable with parameter s, commonly referred to as the Kullback–Leibler divergence. Here and throughout the paper we use log to denote the natural logarithm. For r or s equal to 0 or 1 we define the value of $D_{\mathrm{KL}}(\cdot\|\cdot)$ (possibly infinite) on grounds of continuity, so for example $D_{\mathrm{KL}}(0\|s)=-\log(1-s)$.

2. Main results

With the test design and notation in place, we are now in a position to state our main results. Theorems 2.1, 2.2 are the centerpiece of this paper, featuring improved bounds for the noisy group testing problem for the general $p-q$ model. We follow up in Section 2.2 with a discussion of the combinatorics underlying both algorithms, and provide a converse bound in Section 2.3. Subsequently, in Section 2.4 we show how the bounds simplify when we consider the special cases of the Z, the reverse Z and Binary Symmetric Channel. Finally, in Section 2.5 we derive sufficient conditions under which DD requires fewer tests than the COMP algorithm and compare the bounds of our constant-column design against the Bernoulli design employed in prior literature.

2.1. **Bounds for Noisy Group Testing.** We will consider two well-known algorithms from the noiseless setting to identify infected individuals in this paper. First, we study a noisy variant of the COMP algorithm, originally introduced in [10].

1 Declare every individual that appears in $\alpha\Delta$ or more displayed negative tests as healthy.
2 Declare all remaining individuals as infected.

Algorithm 1: The noisy COMP algorithm

Note that for $\alpha\Delta = 1$ the formulation of Algorithm 1 coincides with the standard COMP algorithm where an individual is classified as healthy if it appears in at least one displayed negative test which constitutes a sufficient condition in the noiseless case. We now state the first main result of this paper.

Theorem 2.1 (Noisy COMP)**.** *Let $p, q \geq 0$, $p+q<1, d \in (0,\infty), \alpha \in (q, e^{-d}(1-p)+\left(1-e^{-d}\right)q)$. Suppose that $0<\theta<1$ and let*

$$m_{COMP} = m_{COMP}(n,\theta,p,q) = \min_{\alpha,d} \max\{b_1(\alpha,d), b_2(\alpha,d)\}\, k\log(n/k)$$
$$\text{where} \qquad b_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|q\right)}$$
$$\text{and} \qquad b_2(\alpha,d) = \frac{1}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}$$

If $m \geq (1+\varepsilon)m_{COMP}$ for some $\varepsilon > 0$, noisy COMP *will recover $\boldsymbol{\sigma}$ w.h.p. given test design $\boldsymbol{G}$ and test results $\hat{\boldsymbol{\sigma}}$.*

The noisy variant of the DD algorithm of [5] was introduced in [48] and reads as follows:

1 Declare every individual that appears in $\alpha\Delta$ or more displayed negative tests as healthy and remove such individual from every assigned test.
2 Declare every yet unclassified individual who is now the only unclassified individual in $\beta\Delta$ or more displayed positive tests as infected.
3 Declare all remaining individuals as healthy.

Algorithm 2: The noisy DD algorithm [48]

Note that the formulation of Algorithm 2 reduces to the noiseless version of DD introduced in [5] by taking $\alpha\Delta = \beta\Delta = 1$. This is because in the noiseless setting a single negative test or a single positive test with just individuals already classified as uninfected is sufficient in the noiseless case. Furthermore note that for $\beta = 0$ noisy DD and noisy COMP are the same. From now on we assume $\beta > 0$. The proof of Theorem 2.1 can be found in Appendix B. We now state the second main result of the paper.

Theorem 2.2 (Noisy DD)**.** *Let $p, q \geq 0$, $p+q<1$, $d \in (0,\infty)$, $\alpha \in (q, e^{-d}(1-p)+\left(1-e^{-d}\right)q)$ and $\beta \in (0, e^{-d}(1-q))$ and define $w = e^{-d}p + (1-e^{-d})(1-q)$. Suppose that $0<\theta<1$ and let*

$$m_{DD} = m_{DD}(n,\theta,p,q) = \min_{\alpha,\beta,d} \max\left\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\alpha,\beta,d)\right\} k\log(n/k)$$

$$\text{where} \qquad c_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|q\right)}$$

$$\text{and} \qquad c_2(\alpha,d) = \frac{1}{dD_{\mathrm{KL}}\left(\alpha\|1-w\right)}$$

$$\text{and} \qquad c_3(\beta,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\beta\|(1-q)e^{-d}\right)}$$

$$\text{and} \qquad c_4(\alpha,\beta,d) = \max_{1-\alpha\leq z\leq 1}\left\{\frac{1}{1-\theta}\frac{1}{d\left(D_{\mathrm{KL}}\left(z\|w\right)+\mathbf{1}\left\{\beta>\frac{ze^{-d}p}{w}\right\}zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right)}\right\}$$

If $m \geq (1+\varepsilon)m_{DD}$ for some $\varepsilon>0$, then noisy DD *will recover $\boldsymbol{\sigma}$ w.h.p. given test design $\boldsymbol{G}$ and test results $\hat{\boldsymbol{\sigma}}$.*

The proof of Theorem 2.2 can be found in Appendix C. While the bounds appear cumbersome at first glance, the optimization is of finite dimension and for every specific value of p and q can be efficiently solved to arbitrary precision yielding explicit values for m_{COMP} and m_{DD}. For illustration purposes, we will calculate those bounds for several values of p, q and θ.

2.2. **The combinatorics of the noisy group testing algorithms.** In the following, we outline the combinatorial structures that Algorithm 1 and 2 take advantage of.
We start with defining the three types of tests that are relevant for the classification of an individual x_i while using COMP and DD. In the first stage we find

- Type DN: Displayed negative tests
- Type DP: Displayed positive tests

Note that the only available information during the first stage of the algorithms is the test result and the pooling structure – no information about the individuals' infection status is available. We give an illustration on the left hand side of Figure 2. After this step COMP terminates by declaring all remaining individuals as infected.

The DD algorithm continues with a second step which considers just the displayed positive tests. From the first step of the algorithm one receives the estimate of the set of non-infected individuals obtained in the first round. Now distinguish the following two types, illustrated on the right hand side in Figure 2:

- Type Displayed-Positive-Single (DP-S): Displayed positive tests in which all other individuals are already declared as uninfected.
- Type Displayed-Positive-Multiple (DP-M): Displayed positive tests with at least one other individual that is not contained in the estimated set of uninfected individuals.

2.2.1. *The noisy COMP algorithm.* To get started, let us shed light on the combinatorics of noisy COMP (Algorithm 1). For the *noiseless* case, the COMP algorithm classifies each individual that appears in at least one negative test as healthy and all other individuals as infected, since the participation in a negative test is a sufficient condition for the individual to be healthy.

For the noisy case, the situation is not as straightforward, since an infected individual might appear in *displayed* negative tests that were flipped when sent through the noisy channel. Thus, a single negative test is not definitive evidence that an individual is healthy. Yet, we can use the number of negative tests to tell the infected individuals apart from the healthy individuals.

Clearly, noisy COMP (Algorithm 1) using a threshold $\alpha\Delta$ succeeds if no healthy individual appears in fewer than $\alpha\Delta$ displayed negative tests and no infected individual appears in more than $\alpha\Delta$ displayed

negative tests. To this end, we define

$$\boldsymbol{N}_x = |\{a \in \partial x : \hat{\boldsymbol{\sigma}}(a) = 0\}| \tag{2.1}$$

for the number of displayed negative tests that item x appears in. In terms of Figure 2, the algorithm determines the infection status by counting the number of tests of Type DN.

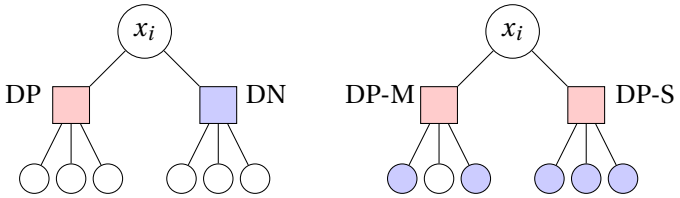

FIGURE 2. The relevant neighborhood structures for the analysis of the algorithms, on the left for the first stage and on the right for the second step. Rectangles represent tests (displayed positive in red, displayed negative in blue). Blue circles represent individuals that have been classified as healthy in the first step of DD (or by COMP). White circles represent individuals that are unclassified in the current stage. We refer to displayed negative tests as Type DN, displayed positive tests as Type DP, displayed positive with a single unclassified individual as Type DP-S and displayed positive with a multiple unclassified individual as Type DP-M

2.2.2. *The noisy* DD *algorithm.* As in the prior section, let us first consider the *noiseless* DD algorithm. The first step is identical to COMP classifying all individuals that are contained in at least one negative test as healthy. In a second step, the algorithm checks each individual to see if it is contained in a positive test as the only remaining unclassified individual after the first step of the algorithm and thus must be infected.

Again, the situation is more intricate when we add noise, since neither a single negative test gives us confidence that an individual is healthy nor does a positive test where the individual is the single remaining unclassified individual after the first step of the algorithm inform us that this individual must be infected. Instead we count and compare the number of such tests. The first step of the noisy DD algorithm is identical to noisy COMP, but we are not required to identify all healthy individuals in the first step (we are able to keep some unclassified for the second round). Thus, after the first step, we are left with all infected individuals V_1 (as the algorithm did not try to classify any individual as infected in the first step) and a set of yet unclassified healthy individuals (as some of them might exhibit a first neighbourhood that is not sufficient for a clear first round classification) which we will denote by $V_{0,\mathrm{PD}}$.These are healthy individuals who did not appear in sufficiently many displayed negative tests to be declared healthy with confidence in the first step[5]. In symbols, for some $\alpha \in (0,1)$

$$V_{0,\mathrm{PD}} = \{x \in V_0 : \boldsymbol{N}_x < \alpha\Delta\}$$

To tell V_1 and $V_{0,\mathrm{PD}}$ apart, we consider the number of displayed positive tests $\boldsymbol{P}_x$ where the individual x appears on its own after removing the individuals , which were declared healthy already, $V_0 \setminus V_{0,\mathrm{PD}}$ from the first step, i.e.

$$\boldsymbol{P}_x = \left|\left\{a \in \partial x : \hat{\boldsymbol{\sigma}}(a) = 1 \text{ and } \partial a \setminus \{x\} \subset V_0 \setminus V_{0,\mathrm{PD}}\right\}\right| \tag{2.2}$$

Referring to Figure 2, the second step of the algorithm is based on counting tests of Type DP-S. Tests of Type DP-M contain another remaining unclassified individual after the first step of the algorithm from $V_{0,PD} \cup V_1$. The noisy DD algorithm takes advantage of the fact that it is less likely for an individual $x \in V_{0,\mathrm{PD}}$ to appear as the only yet unclassified individual in a displayed positive test than it is for an

[5]Note that the bounds are taken in a way such that no infected individual is classified as uninfected in the first round.

individual in $x \in V_1$. For $x \in V_{0,\mathrm{PD}}$ such a test would be truly negative and would have been flipped (which occurs with probability p) to display a positive test result. Conversely, an individual $x \in V_1$ renders any of its tests truly positive and thus the only requirement is that the test otherwise contains only individuals which were declared healthy already, and is not flipped (which occurs with probability $1-q$). For this reason, we will see that the distribution of $\boldsymbol{P}_x$ differs between $x \in V_1$ and $x \in V_{0,\mathrm{PD}}$, and the difference $(1-q)-p>0$ helps determine the size of this difference. The second step of DD exploits this observation by counting tests of Type DP-S.

2.3. **The Channel Perspective of noisy group testing.** Motivated by (1.1), we can describe the bounds in terms of rate, in a Shannon-theoretic sense. That is, we follow the common notion to define the rate (bits learned per test) of an algorithm in this setting (for instance as in [9]) to be

$$R := \frac{\log\binom{n}{k}}{m\log 2} \sim \frac{k\log(n/k)}{m\log 2}.$$

(Recall that we take logarithms to base e throughout this paper). For example the fact that Theorems 2.1 and 2.2 show that noisy COMP and DD respectively can succeed w.h.p. ; with $m \geq (1+\epsilon)ck\log(n/k)$ tests for some c is equivalent to the fact that $R = 1/(c\log 2)$ is an achievable rate in a Shannon-theoretic sense.

We now give a counterpart to these two theorems by stating a universal converse for the $p-q$ channel below, improving on the universal counting bound from (1.1). The starting observation (see [7, Theorem 3.1]) is that no group testing algorithm can succeed w.h.p. with rate greater than C_{Chan}, the Shannon capacity of the corresponding noisy communication channel. Thus, we cannot hope to succeed w.h.p. with $m < (1-\epsilon)ck\log(n/k)$ tests where $c = 1/(C_{\mathrm{Chan}}\log 2)$. Hence as a direct consequence of the value of the channel capacity of the $p-q$ channel, we deduce the following statement.

Corollary 2.3. *Let $p, q \geq 0$, $p+q<1$ and $\epsilon>0$, write $h(\cdot)$ for the binary entropy in nats (logarithms taken to base e) and $\phi = \phi(p,q) = (h(p)-h(q))/(1-p-q)$. If we define*

$$m_{COUNT} = \left(\frac{1}{D_{\mathrm{KL}}\left(q \| 1/(1+e^{\phi})\right)}\right) k\log(n/k),$$

then for $m \leq (1-\epsilon)m_{COUNT}$ no algorithm can recover $\boldsymbol{\sigma}$ w.h.p. for any matrix design.

Remark 2.4. *This result follows from Lemma F.1 derived in Appendix F below. As discussed there, this derivation (combined with the fact that each test is negative with probability e^{-d}) suggests a choice of density for the matrix:*

$$d = d^*_{\mathrm{ch}} = \log(1-p-q) - \log\left(\frac{1}{1+e^{\phi}} - q\right).$$

While a choice of $\Delta = c \cdot d^*_{\mathrm{ch}} \cdot \log(n/k)$ is not necessarily optimal, it may be regarded as a sensible heuristic that provides good rates for a range of p and q values.

2.4. **Applying the results to standard channels.** With Theorem 2.1 and Theorem 2.2 we derived achievable rates for the generalized *p-q-model* (see Figure 1). Prior research considered the Z channel where $p=0$ and $q>0$, the Reverse Z channel where $p>0$ and $q=0$ and the Binary Symmetric Channel with $p=q>0$. These channels are common models in coding theory [41], but are also often considered in medical applications [30, 31] concerned with taking imperfect sensitivity ($q>0$), specificity ($p>0$) or both ($p>0$ and $q>0$) into account. As a consequence we also compare our results with the most recent results of Johnson and Scarlett [48]. In the following section we will demonstrate how performance guarantees on these channels can directly be obtained from our main theorems.

2.4.1. *Recovery of the noiseless model.* Note that the bounds Corollary 2.5 and Corollary 2.6 are already known [10, 26]. We would like to give the reader an idea of how one can see that our cumbersome looking bounds relate to the more accessible bounds given for the noiseless case. First, we show the noiseless bounds can be simply recovered by letting $p, q \to 0$. In the noiseless setting, it is sufficient, by definition

of the algorithm, to set both $\alpha\Delta = 1$ and $\beta\Delta = 1$. To see why, observe that in the absence of noise a single negative test is sufficient evidence that an individual is healthy. Conversely, a single positive test where the individual only appears with individuals , which were declared healthy already, implies that particular individual must surely be infected. As shown in [13] the optimal parameter choice for the density parameter d in the constant-column design in the noiseless setting is $\log(2)$. Applying these values to Theorem 2.1 we recover the noiseless bound for COMP.These bounds were first stated in [10].

Corollary 2.5 (COMP in the noiseless setting)**.** *Let $p,q \to 0$, $0 < \theta < 1$ and $\varepsilon > 0$. Further, let*

$$m_{\texttt{COMP},noiseless} = \frac{1}{(1-\theta)\log^2 2} k\log(n/k).$$

Furthermore let $m_{\texttt{COMP}}(n,\theta,p,q)$ be defined as in Theorem 2.1 Then we find

$$m_{\texttt{COMP}}(n,\theta,p,q) \underset{p,q\to 0}{\longrightarrow} m_{\texttt{COMP},noiseless}$$

Proof. We start by taking the bounds $b_1(\alpha,d)$ and $b_2(\alpha,d)$. To see how this boils down to $m_{\texttt{COMP},\text{noiseless}}$, we start with using the well-known fact that within the near constant column design $d = \log(2)$ is the optimal choice [13]. Now by taking both $p,q \to 0$ one realizes that $b_1(\alpha,\log(2))$ vanishes as $\log(p) \to -\infty$ as $p \to 0$. Turning our focus to the second bound we see that it boils down to

$$b_2(\alpha,\log(2))) = \frac{1}{(1-\theta)\log(2)} \frac{1}{\log(2) + \alpha\log(\alpha) + (1-\alpha)\log(1-\alpha)}$$

On the one hand we realize that $\alpha\log(\alpha) + (1-\alpha)\log(1-\alpha)$ is negative for all $\alpha \in (0,1)$. This leads to

$$b_2(\alpha,\log(2)) > b_2(0,\log(2))$$

On the other hand we realize that in the noiseless case a single negative test is sufficient for a classification as uninfected. Therefore we may choose $\alpha > 0$ sufficiently small. One indeed realizes that for each α we can choose $\varepsilon := \varepsilon(\alpha) > 0$ appropriately, such that the bounds given in Theorem 2.1 recover the noiseless case. □

We also recover the noiseless bounds for the DD algorithm as stated in [26].

Corollary 2.6 (DD in the noiseless setting)**.** *Let $p,q \to 0, 0 < \theta < 1$ and $\varepsilon > 0$. Further, let*

$$m_{\texttt{DD},noiseless} = \max\left\{1, \frac{\theta}{1-\theta}\right\} \frac{1}{\log^2 2} k\log(n/k).$$

Furthermore let $m_{\texttt{DD}}(n,\theta,p,q)$ be defined as in Theorem 2.2 Then we find

$$m_{\texttt{DD}}(n,\theta,p,q) \underset{p,q\to 0}{\longrightarrow} m_{\texttt{DD},noiseless}$$

Proof. We start with taking $c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d)$ and $c_4(\alpha,\beta,d)$ as defined in Theorem 2.2. First of all we take $c_4(\alpha,\beta,d)$. By assumption we find $\beta > 0$ and therefore the indicator is 1 as soon as we let $p \to 0$. Furthermore for $p \to 0$ we get $-\log(p) \to \infty$ and find $c_4 \to 0$. Second of all we take $c_1(\alpha,d)$. With a similar argument as before we see that $c_1(\alpha,d) \to 0$ for $q \to 0$ as in this case we find $-\log(q) \to \infty$. Therefore we are left with $c_2(\beta,d)$ and $c_3(\alpha,\beta,d)$. Again, we use the well known fact that in the noiseless case $d = \log(2)$ is the optimal choice. Therefore with $p,q \to 0$ the two remaining bounds read as follows:

$$c_2(\alpha,\log(2)) = \frac{1}{\log(2)\left(\log(2) + \alpha\log(\alpha) + (1-\alpha)\log(1-\alpha)\right)}$$

$$c_3(\alpha,\beta,\log(2)) = \frac{\theta}{(1-\theta)} \frac{1}{\log(2)\left(\log(2) + \beta\log(\beta) + (1-\beta)\log(1-\beta)\right)}$$

Again we see that $x\log(x)+(1-x)\log(1-x)$ is negative for $x\in(0,1)$. Therefore we find

$$c_2(\alpha,\log(2)) > c_2(0,\log(2))$$
$$c_3(\alpha,\log(2)) > c_3(0,\log(2))$$

Now as as before in this case again a single negative test as well as a single test with only already classified uninfected individuals is sufficient. Therefore we can choose $\alpha,\beta>0$ sufficiently small. One indeed realizes that for each $\alpha,\beta>0$ one can choose $\varepsilon:=\varepsilon(\alpha,\beta)$ appropriately such that the bounds of Theorem 2.2 recover the noiseless case. □

2.4.2. *The Z channel.* In the Z channel, we have $p=0$ and $q>0$, i.e. no truly negative test displays a positive test result. Thus, in this case finding one positive test with only one unclassified individual is a clear indication, therefore we again can choose $\beta>0$ sufficiently small and remain agnostic about α and d. The bounds for COMP and DD thus read as follows.

Corollary 2.7 (Noisy COMP for the Z channel)**.** *Let $p\to 0, 0<q<1, 0<\theta<1$ and $\varepsilon>0$. Further, let*

$$m_{\mathtt{COMP},Z} = \min_{\alpha,d}\max\{b_1(\alpha,d),b_2(\alpha,d)\}\,k\log(n/k)$$
$$\text{with}\quad b_1(\alpha,d)=\frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|q\right)}\quad\text{and}\quad b_2(\alpha,d)=\frac{1}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}+\left(1-e^{-d}\right)q\right)}.$$

If $m>(1+\varepsilon)m_{\mathtt{COMP},Z}$, noisy COMP *will recover $\boldsymbol{\sigma}$ w.h.p. given $\boldsymbol{G},\hat{\boldsymbol{\sigma}}$.*

Corollary 2.8 (Noisy DD for the Z channel)**.** *Let $p\to 0, 0<q<1, 0<\theta<1$ and $\varepsilon>0$. Further, let*

$$m_{\mathtt{DD},Z} = \min_{\alpha,d}\max\{c_1(\alpha,d),c_2(\alpha,d),c_3(d)\}\,k\log(n/k)$$
$$\text{with}\quad c_1(\alpha,d)=\frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|q\right)}\quad\text{and}\quad c_2(\alpha,d)=\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}+\left(1-e^{-d}\right)q\right)}$$
$$\text{and}\quad c_3(d)=\frac{\theta}{1-\theta}\frac{1}{-d\log\left(1-e^{-d}(1-q)\right)}.$$

If $m>(1+\varepsilon)m_{\mathtt{DD},Z}$, noisy DD *will recover $\boldsymbol{\sigma}$ w.h.p. given $\boldsymbol{G},\hat{\boldsymbol{\sigma}}$.*

Proof. The bounds c_1 and c_2 follow directly from Theorem 2.2 by letting $p\to 0$. An immediate consequence of $p\to 0$ is that due to the fact that $-\log(p)\to\infty$ and one finds that $c_4\to 0$, thus being trivial in this case. For c_3 we use the fact that we can choose $\beta>0$ sufficiently small we find $D_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-q)\right)=-\log\left(1-e^{-d}(1-q)\right)-\delta(\beta)$ for $\delta(\beta)>0$. Note that by definition of the noise model, we may choose an arbitrary $\beta_{\min}$ very close to zero and as a consequence $\beta=\beta_{\min}$ leading to $\delta(\beta)\to\delta_{\min}$. The assertion follows as for each β we may choose $\varepsilon:=\varepsilon(\beta)>0$ such that $(1+\varepsilon)>\left(1+\varepsilon\left(\beta_{\min}\right)\right)$. □

An illustration of the bounds from Corollary 2.7 and 2.8 for sample values of q is shown in Figure 5.

2.4.3. *Reverse Z channel.* In the reverse Z channel, we have $q=0$ and $p>0$, i.e. no truly positive test displays a negative test result. Thus, we may choose $\alpha>0$ sufficiently small and remain agnostic about β and d. The bounds for the noisy COMP and DD thus read as follows.

Corollary 2.9 (Noisy COMP for the Reverse Z channel)**.** *Let $0<p<1, q\to 0, 0<\theta<1$ and $\varepsilon>0$. Further, let*

$$m_{\mathtt{COMP},rev\,Z}=\frac{1}{1-\theta}\min_{d}\left\{\frac{1}{-d\log\left(1-e^{-d}(1-p)\right)}\right\}k\log(n/k).$$

If $m>(1+\varepsilon)m_{\mathtt{COMP},rev\,Z}$, noisy COMP *will recover $\boldsymbol{\sigma}$ w.h.p. given $\boldsymbol{G},\hat{\boldsymbol{\sigma}}$.*

Proof. The corollary follows from Theorem 2.1 and the fact that for $q \to 0$ one finds that $D_{\mathrm{KL}}(\alpha\|0)$ diverges, Thereby $b_1 \to 0$ just gives a trivial bound in this case. Furthermore for sufficiently small $\alpha > 0$ we get $D_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p)\right) \to -\log\left(1-e^{-d}(1-p)\right) - \delta(\alpha)$. Due to the noise assumption, we may choose an arbitrary $\alpha_{\min}$ very close to zero and $\alpha = \alpha_{\min}$ which leads to $\delta(\alpha) \to \delta(\alpha_{\min})$. The assertion follows by choosing $\varepsilon := \varepsilon(\alpha) > 0$ such that $(1+\varepsilon) > (1+\varepsilon(\alpha_{\min}))$. □

Note that Corollary 2.9 does not yield an immediate closed form expression for the optimal value of d.

Corollary 2.10 (Noisy DD in the Reverse Z channel)**.** *Let* $0 < p < 1, q \to 0, 0 < \theta < 1$ *and* $\varepsilon > 0$. *Further, let*

$$m_{\mathtt{DD},rev\,Z} = \min_{\beta,d} \max\left\{c_2(d), c_3(\beta,d), c_4(\beta,d)\right\} k\log(n/k)$$

$$\text{with } c_2(d) = \frac{1}{-d\log\left(1-e^{-d}(1-p)\right)} \quad \text{and} \quad c_3(\beta,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\beta\|e^{-d}\right)}$$

$$\text{and } c_4(\beta,d) = \frac{1}{1-\theta}\frac{1}{d\left(-\log\left(1-e^{-d}(1-p)\right) + D_{\mathrm{KL}}\left(\beta\|\frac{e^{-d}p}{e^{-d}p+(1-e^{-d})}\right)\right)}$$

If $m > (1+\varepsilon)m_{\mathtt{DD},rev\,Z}$, *noisy* DD *will recover* $\boldsymbol{\sigma}$ *w.h.p. given* $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$.

Proof. First of all we assume $q \to 0$. Therefore we find $c_1 \to 0$ as $-\log(q) \to \infty$. The bounds c_2, c_3 follow from Theorem 2.2 and the same manipulations as above. For c_4, we again see that by definition of the noise model we may choose $\alpha > 0$ as close to zero as we like. Therefore we get $(1-\alpha)$ close to 1, which leads to $z \to 1$. The assertion follows as for each α we can choose $\varepsilon := \varepsilon(\alpha) > 0$ such that $(1+\varepsilon) > (1+\varepsilon(\alpha_{\min}))$. □

An illustration of the bounds of Corollary 2.9 and 2.10 for sample values of p is shown in Figure 6.

2.4.4. *Binary Symmetric Channel.* In the Binary Symmetric Channel (BSC), we set $p = q > 0$. Even though information-theoretic arguments would suggest setting $d = \log 2$, we formulate the expression below with general d. We also keep the threshold parameters α and β. The bounds for the noisy DD and COMP only simplify slightly.

Corollary 2.11 (Noisy COMP in the Binary Symmetric Channel)**.** *Let* $0 < p = q < 1/2, 0 < \theta < 1$ *and* $\varepsilon > 0$. *Further, let*

$$m_{\mathtt{COMP},BSC} = \min_{\alpha,d} \max\{b_1(\alpha,d), b_2(\alpha,d)\} k\log(n/k)$$

$$\text{with} \quad b_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|p\right)} \quad \text{and} \quad b_2(\alpha,d) = \frac{1}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}+p-2e^{-d}p\right)}.$$

If $m > (1+\varepsilon)m_{\mathtt{COMP},BSC}$, *noisy* COMP *will recover* $\boldsymbol{\sigma}$ *w.h.p. given* $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$.

Corollary 2.12 (Noisy DD in the Binary Symmetric Channel)**.** *Let* $0 < p = q < 1/2, 0 < \theta < 1$ *and* $\varepsilon > 0$ *and define* $v = 1 - e^{-d} - p + 2e^{-d}p$. *Further, let*

$$m_{\mathtt{DD},BSC} = \min_{\alpha,\beta,d} \max\left\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\alpha,\beta,d)\right\} k\log(n/k)$$

$$\text{with} \quad c_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|p\right)} \quad \text{and} \quad c_2(\alpha,d) = \frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}+p-2e^{-d}p\right)}$$

$$\text{and} \quad c_3(\beta,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\beta\|(1-p)e^{-d}\right)}$$

$$\text{and} \quad c_4(\alpha,\beta,d) = \max_{1-\alpha\le z\le 1}\left\{\frac{1}{1-\theta}\frac{1}{d\left(D_{\mathrm{KL}}(z\|v) + \mathbf{1}\left\{\beta > \frac{ze^{-d}p}{v}\right\} zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{v}\right)\right)}\right\}.$$

If $m > (1+\varepsilon)m_{\mathtt{DD},BSC}$, *noisy* DD *will recover* $\boldsymbol{\sigma}$ *w.h.p. given* $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$.

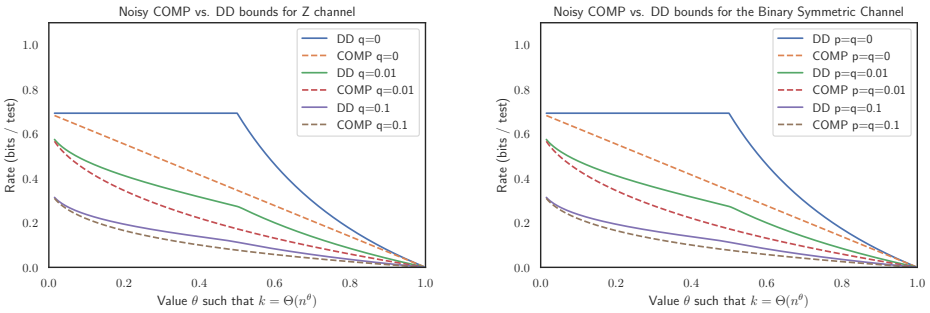

FIGURE 3. Comparison of the bound for noisy DD and noisy COMP in the Z-channel and the Binary Symmetric Channel for different noise level. (Note for black and white prints: The lines in the diagram are in the same order as given in the legend from top to bottom)

An illustration of the bounds of Corollary 2.11 and 2.12 is shown in Figure 7.

2.5. **Comparison of noisy COMP and DD.** An obvious next question is to find conditions under which the noisy DD algorithm requires fewer tests than the noisy COMP. For the noiseless setting, it can be easily shown that DD provably outperforms COMP for all $\theta \in (0,1)$. For the noisy case, matters are slightly more complicated.

Recall that noisy COMP classifies all individuals appearing in less than $\alpha\Delta$ displayed negative tests as infected while noisy DD additionally requires such individuals to appear in more than $\beta\Delta$ displayed positive tests as the only yet unclassified individual. Thus, it might well be that an infected individual is classified correctly by noisy COMP, while it is missed by the noisy DD algorithm.

That being said, our simulations indicate that noisy DD generally requires fewer tests than noisy COMP, but for the reason mentioned above we can only prove that for the reverse Z channel while remaining agnostic about the Z channel and the Binary Symmetric Channel, as the next proposition evinces.

Proposition 2.13. *For all $p,q \geq 0$ with $p+q<1$ there exists a $d^* \in (0,\infty)$ such that $m_{COMP} \geq m_{DD}$ as long as $e^{-d^*} p \geq q$.*

In terms of the common noise channels Proposition 2.13 gives the following corollary.

Corollary 2.14. *In the reverse Z channel, $m_{COMP} \geq m_{DD}$.*

The proof can be found in Appendix D. Our simulations suggest that this superior performance of noisy DD holds as well for the Z channel and Binary Symmetric Channel. Please refer to Figure 3 for an illustration.

2.6. **Relation to Bernoulli testing.** In [48] sufficient bounds for noisy group testing and a Bernoulli test design where each individual joins every test independently with some fixed probability were derived. Thus, the variable degrees fluctuate and we end up with some individuals assigned only to few tests. In contrast, we work under a model in this paper where each individual joins an equal number of tests Δ chosen uniformly at random without replacement. For the noiseless case, it is by now clear that the near-constant-column design better facilitates inference than the Bernoulli test design [13, 26]. We find that the same holds true for the noisy variant of the COMP algorithm. Let us denote by $m_{\mathrm{COMP}}^{\mathrm{Ber}}$ the number of tests required for the noisy COMP to succeed under a Bernoulli test design.

Proposition 2.15. *For all $p+q<1$, we have*

$$m_{COMP}^{Ber} \geq m_{COMP}$$

We see the same effect for the noisy variant of the DD algorithm for all simulations, but for technical reasons only prove it for the Z channel.

Proposition 2.16. *For the Z channel where $p = 0$ and $0 < q < 1$, we have*

$$m_{DD}^{Ber} > m_{DD}$$

For an illustration on the magnitude of the difference, we refer to Figure 4 and Figure 8.

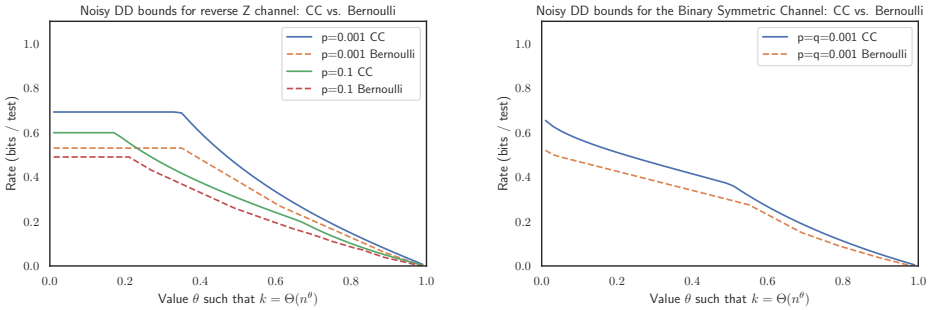

FIGURE 4. Comparison of DD bounds under a Bernoulli test design ([48]) and constant column test design (present paper) for the reverse Z and Binary Symmetric Channel. (Note for black and white prints: The solid lines as well as the dashed lines in the diagram are in the same order as given in the legend from top to bottom)

APPENDIX

The core of the technical sections is the proof of Theorems 2.1 and Theorem 2.2. Some groundwork with standard concentration bounds and group testing properties can be found in Section A. We continue with the proof of Theorems 2.1 and 2.2 in Sections B and C, respectively. The structure of the proofs follows a similar logic. First, we derive the distributions for the number of displayed positive and negative tests for infected and healthy individuals. Second, we threshold these distributions using sharp Chernoff concentration bounds to deduce the bounds stated in Theorem 2.1 and Theorem 2.2. Thereafter, we proceed to the proof of Proposition 2.13 in Section D, while the proofs of Propositions 2.15 and 2.16 follow in Section E. The proof of Corollary 2.3 can be found in Section F. Additional illustrations of our results for the different channels can be found in Section G.

APPENDIX A. GROUNDWORK

For starters, let us recall the Chernoff bound for binomial and hypergeometric distributions.

Lemma A.1 (Chernoff bound for the binomial distribution [25])**.** *Let $p < q < r \in (0,1)$ and $\boldsymbol{X} \sim \mathrm{Bin}(n,q)$ be a binomially distributed random variable. Then*

$$\mathbb{P}\left(\boldsymbol{X} \le \lceil pn \rceil\right) = \exp\left(-\left(1+n^{-\Omega(1)}\right) nD_{\mathrm{KL}}\left(p\|q\right)\right)$$
$$\mathbb{P}\left(\boldsymbol{X} \ge \lceil rn \rceil\right) = \exp\left(-\left(1+n^{-\Omega(1)}\right) nD_{\mathrm{KL}}\left(r\|q\right)\right)$$

Lemma A.2 (Chernoff bound for the hypergeometric distribution [23])**.** *Let $p < q < r \in (0,1)$ and $\boldsymbol{Y} \sim H(N,Q,n)$ be a hypergeometrically distributed random variable. Further, let $q = Q/N$. Then*

$$\mathbb{P}\left(\boldsymbol{Y} \le \lceil pn \rceil\right) = \exp\left(-\left(1+n^{-\Omega(1)}\right) nD_{\mathrm{KL}}\left(p\|q\right)\right)$$
$$\mathbb{P}\left(\boldsymbol{Y} \ge \lceil rn \rceil\right) = \exp\left(-\left(1+n^{-\Omega(1)}\right) nD_{\mathrm{KL}}\left(r\|q\right)\right)$$

The next lemma provides that the test degrees, as defined in (1.4) above, are tightly concentrated. Recall from (1.3) that the number of tests $m = ck\log(n/k)$ and each item appears in $\Delta = cd\log(n/k)$ tests.

Lemma A.3. *With probability $1-o(n^{-2})$ we have*

$$dn/k - \sqrt{dn/k}\log n \le \boldsymbol{\Gamma}_{\min} \le \boldsymbol{\Gamma}_{\max} \le dn/k + \sqrt{dn/k}\log n$$

Proof. The probability that an individual x is assigned to test a is given by

$$\mathbb{P}(x \in \partial a) = 1 - \mathbb{P}(x \notin \partial a) = 1 - \binom{m-1}{\Delta}\binom{m}{\Delta}^{-1} = \Delta/m = d/k \tag{A.1}$$

Since each individual is assigned to tests independently, the total number of individuals in a given test follows the binomial distribution $\mathrm{Bin}(n, d/k)$. The assertion now follows from applying the Chernoff bound for this binomial distribution at the expectation (Lemma A.1). □

Next, we show that the number of truly negative tests $\boldsymbol{m}_0$ (and thus the number of truly positive tests $\boldsymbol{m}_1$) is tightly concentrated.

Lemma A.4. *With probability $1-o(n^{-2})$ we have $\boldsymbol{m}_0 = e^{-d}m + O(\sqrt{m}\log^3 n)$.*

Proof. Recall from (A.1) that

$$\mathbb{P}(x \in \partial a) = d/k$$

Since infected individuals are assigned to tests mutually independently, we find for a test a that

$$\mathbb{P}(V_1 \cap \partial a = \emptyset) = \mathbb{P}(\mathrm{Bin}(k, d/k) = 0) = (1-d/k)^k = \left(1+n^{-\Omega(1)}\right)e^{-d}.$$

Consequently, $\mathbb{E}[\boldsymbol{m}_0] = \left(1 + n^{-\Omega(1)}\right) e^{-d} m$. Finally, changing the set of tests for a specific infected individual shifts the total number of negative tests by at most Δ. Therefore, the McDiarmid inequality (Lemma 1.2 in [34]) yields

$$\mathbb{P}\left(|\boldsymbol{m}_0 - \mathbb{E}[\boldsymbol{m}_0]| \geq t\right) \leq 2\exp\left(-\frac{t^2}{4k\Delta^2}\right).$$

The lemma follows from setting $t = O\left(\sqrt{m}\log^3 n\right)$. □

With the concentration of $\boldsymbol{m}_0$ and $\boldsymbol{m}_1$ at hand, we readily obtain estimates for $\boldsymbol{m}_0^f, \boldsymbol{m}_0^u, \boldsymbol{m}_1^f$ and $\boldsymbol{m}_1^u$. We remind ourselves that these are the number of flipped, unflipped negative tests and the number of flipped, unflipped positive tests as defined in Sec. 1.4.

Corollary A.5. *With probability* $1 - o(n^{-2})$ *we have*

(i) $\boldsymbol{m}_0^f = e^{-d} pm + O\left(\sqrt{m}\log^3 n\right)$
(ii) $\boldsymbol{m}_0^u = e^{-d}(1-p)m + O\left(\sqrt{m}\log^3 n\right)$
(iii) $\boldsymbol{m}_1^f = (1 - e^{-d})qm + O\left(\sqrt{m}\log^3 n\right)$
(iv) $\boldsymbol{m}_1^u = (1 - e^{-d})(1-q)m + O\left(\sqrt{m}\log^3 n\right)$

Proof. Since each test is flipped with probability p and q independently, the claims follow from Lemma A.4 and the Chernoff bound for the binomial distribution (Lemma A.1). □

In the following, let $\mathscr{E}$ be the event that the bounds from Lemma A.4 and A.5 hold. Note that $\mathscr{E}$ holds with high probability.

Appendix B. Proof of COMP bound, Theorem 2.1

Recall from (2.1) that we write $\boldsymbol{N}_x$ for the number of displayed negative tests that item x appears in (as illustrated by the right branch of Fig. 2). The proof of Theorem 2.1 is based on two pillars. First, Lemmas B.1 and B.2 provide the distribution of $\boldsymbol{N}_x$ for healthy and infected individuals, respectively. We will see that these distributions differ according to the infection status of the individual. Second, we will derive a suitable threshold $\alpha\Delta$ via Lemma B.3 and B.4 to tell healthy and infected individuals apart w.h.p. We start by analysing individuals in the infected set V_1. Throughout the section, we assume $\alpha \in (q, e^{-d}(1-p) + \left(1 - e^{-d}\right) q)$.

Lemma B.1. *Given* $x \in V_1$*, its number of displayed negative tests* $\boldsymbol{N}_x$ *is distributed as* $\mathrm{Bin}(\Delta, q)$.

Proof. Any test containing an infected individual is truly positive because of the presence of the infected individual. Since an infected individual is assigned to Δ different tests and each such test is flipped with probability q independently, the lemma follows immediately. □

Next, we consider the distribution for healthy individuals. Recall that $\mathscr{E}$ denotes the event that the bounds from Lemma A.4 and Corollary A.5 hold.

Lemma B.2. *Given* $x \in V_0$ *and conditioned on* $\mathscr{E}$*, the total variation distance of the distribution of* $\boldsymbol{N}_x$ *and* $\boldsymbol{T}_h$ *that is distributed as* $H\left(m, m\left(e^{-d}(1-p) + \left(1 - e^{-d}\right) q\right), \Delta\right)$ *tends to zero with* n*, that is*

$$d_{TV}(\boldsymbol{N}_x, \boldsymbol{T}_h) = n^{-\Omega(1)}$$

Proof. Since x is healthy, the outcome of all the tests remains the same if it is removed from consideration (if we perform group testing with $n-1$ items and the corresponding reduced matrix).

Thus, given $\mathscr{E}$, we find that with x removed the $\boldsymbol{m}_0^f, \boldsymbol{m}_0^u, \boldsymbol{m}_1^f, \boldsymbol{m}_1^u$ still satisfy the bounds from Corollary A.5. As a result the number of displayed negative tests (which consist of unflipped truly negative tests and flipped truly positive tests) is given by

$$\boldsymbol{m}_0^u + \boldsymbol{m}_1^f = \left(e^{-d}(1-p) + (1-e^{-d})q\right) m + O\left(\sqrt{m}\log^3 n\right) \tag{B.1}$$

Now, adding x back into consideration: $x \in V_0$ chooses Δ tests without replacement independently of this. Hence, given that the random quantity $\boldsymbol{m}_0^u + \boldsymbol{m}_1^f = \ell$, the $\boldsymbol{N}_x$ (the number of displayed negative tests that item x appears in) is distributed as $H(m,\ell,\Delta)$. Hence, a conditioning argument shows that the linear combination of distribution functions

$$\sum_{\ell} \mathbb{P}\left(\boldsymbol{m}_0^u + \boldsymbol{m}_1^f = \ell\right) \mathbb{P}(H(m,\ell,\Delta) \le x)$$

tends to the distribution function of $H\left(m, m\left(e^{-d}(1-p) + \left(1-e^{-d}\right)q\right), \Delta\right)$ in total variation distance, due to the concentration of $\boldsymbol{m}_0^u + \boldsymbol{m}_1^f$ as obtained in Corollary A.5. □

Moving to the second pillar of the proof, we need to demonstrate that no infected individual is assigned to more than $\alpha\Delta$ displayed negative tests as shown by the following lemma.

Lemma B.3. *If $c > (1+\eta)\frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}(\alpha\|q)}$ for some small $\eta > 0$, $\boldsymbol{N}_x < \alpha\Delta$ for all $x \in V_1$ w.h.p.*

Proof. We have to ensure that $\mathbb{P}(\exists x \in V_1 : \boldsymbol{N}_x \ge \alpha\Delta) = o(1)$. By Lemma B.1 and the union bound, we thus need to have

$$o(1) = k \cdot \mathbb{P}(\boldsymbol{N}_x \ge \alpha\Delta : x \in V_1) = k \cdot \mathbb{P}\left(\mathrm{Bin}(\Delta, q) \ge \alpha\Delta\right) = k \cdot \exp\left(-\left(1+\Delta^{-\Omega(1)}\right)\Delta D_{\mathrm{KL}}\left(\alpha\|q\right)\right),$$

by the Chernoff bound for the binomial distribution (Lemma A.1). Since $k \sim n^\theta$ and $\Delta = cd(1-\theta)\log n$ the following must hold

$$\theta - cd(1-\theta)D_{\mathrm{KL}}\left(\alpha\|q\right) < 0$$

The lemma follows from rearranging terms and the fact that if we choose the number of tests slightly above the required number of tests (larger by a factor of $1+\eta$ for $\eta > 0$), the assertion holds w.h.p. as $n \to \infty$. □

We proceed to show that no healthy individual is assigned to less than $\alpha\Delta$ displayed negative tests.

Lemma B.4. *If $c > (1+\eta)\frac{1}{1-\theta}\frac{1}{dD_{\mathrm{KL}}(\alpha\|e^{-d}(1-p)+(1-e^{-d})q)}$ for some small $\eta > 0$, $\boldsymbol{N}_x > \alpha\Delta$ for all $x \in V_0$ w.h.p.*

Proof. We need to ensure that $\mathbb{P}(\exists x \in V_0 : \boldsymbol{N}_x < \alpha\Delta) = o(1)$. Since $\mathscr{E}$ occurs w.h.p. by Lemma A.4 and Corollary A.5, we need to have by Lemma B.2 and the union bound that

$$(n-k) \cdot \mathbb{P}(\boldsymbol{N}_x \le \alpha\Delta | x \in V_0, \mathscr{E}) \le n \cdot \mathbb{P}(\boldsymbol{T}_h \le \alpha\Delta) = o(1). \tag{B.2}$$

We remind ourselves that $\boldsymbol{T}_h \sim H\left(m, m\left(e^{-d}(1-p) + \left(1-e^{-d}\right)q\right), \Delta\right)$ and together with the Chernoff bound for the hypergeometric distribution (Lemma A.2) this leads to the following condition[6]

$$1 - cd(1-\theta)D_{\mathrm{KL}}\left(\alpha\|(1-p)e^{-d} + (1-e^{-d})q\right) < 0$$

in a similar way to the proof of Lemma B.3. The lemma follows from rearranging terms and the fact that if we choose the number of tests slightly above the required number of tests (larger by a factor of $1+\eta$ for $\eta > 0$), the assertion holds w.h.p. as $n \to \infty$. □

Proof of Theorem 2.1. The theorem is now an immediate consequence of Lemma B.3 and B.4 which guarantee that w.h.p. classifying individuals according to the threshold $\alpha\Delta$ for negative displayed tests recovers $\boldsymbol{\sigma}$, and the fact that the choice of α and d is at our disposal. □

[6]Note that the additive rule of the logarithm allows us to move the error term from inside the KL-divergence to outside

APPENDIX C. PROOF OF DD BOUND, THEOREM 2.2

The proof of Theorem 2.2 follows a similar two-step approach as the proof of Theorem 2.1 by first finding the distribution of $\boldsymbol{P}_x$ (the number of displayed positive tests where individual x appears on its own after removing the individuals, which were declared healthy already, $V_0 \setminus V_{0,\mathrm{PD}}$, illustrated by DP-S in Fig. 2). We then threshold the distributions for healthy and infected individuals. To get started, we revise the second bound from Theorem 2.1 to allow $kn^{-\Omega(1)}$ healthy individuals to not be classified yet after the first step of DD. Recall that, we assume $\alpha \in (q, e^{-d}(1-p) + \left(1-e^{-d}\right)q)$ and $\beta \in (0, e^{-d}(1-q))$.

Lemma C.1. *If*

$$c > (1+\eta)\frac{1}{dD_{\mathrm{KL}}\left(\alpha \| e^{-d}(1-p) + \left(1-e^{-d}\right)q\right)}$$

for some small $\eta > 0$, we have $\left|\boldsymbol{V_{0,PD}}\right| = kn^{-\Omega(1)}$ w.h.p.

Proof. The lemma follows immediately by replacing the r.h.s. of (B.2) with $kn^{-\delta}$ for some small $\delta = \delta(\eta)$, rearranging terms and applying Markov's inequality. □

For the next lemmas, we need an auxiliary notation denoting the number of tests $\boldsymbol{m}_{0,\mathrm{nd}}$ that only contain individuals from $V_0 \setminus V_{0,\mathrm{PD}}$. In symbols,

$$\boldsymbol{m}_{0,\mathrm{nd}} = \left|\left\{a \in F : \partial a \subset V_0 \setminus V_{0,\mathrm{PD}}\right\}\right|.$$

Lemma C.2. *If*

$$c > (1+\eta)\frac{1}{dD_{\mathrm{KL}}\left(\alpha \| e^{-d}(1-p) + \left(1-e^{-d}\right)q\right)}$$

for some small $\eta > 0$, we have $\boldsymbol{m}_{0,nd} = \left(1-n^{-\Omega(1)}\right)e^{-d}m$ with probability $1-o(n^{-2})$.

Proof. As in the proof of Lemma B.2 above, we consider the graph in two rounds: in the first round we consider the tests containing infected individuals. Since each healthy individual $x \in V_0$ does not impact the number of positive and negative tests, we know by Lemma A.4 that with probability $1-o(n^{-2})$ we find that the number of truly negative tests $\boldsymbol{m}_0 = e^{-d}m + O\left(\sqrt{m}\log^4 n\right)$ after the first round. Furthermore the presence of a healthy individual has no impact on the number of displayed negative tests, as unflipped negative tests remain unflipped and flipped positive tests remain flipped. In the second round, we consider the effect of adding healthy individuals into the tests. Knowing the number of negative tests w.h.p. we can think of the participation of individuals $x \in V_{0,\mathrm{PD}}$ in these tests as a *balls into bins* experiment. Starting with the number of truly negative tests $\boldsymbol{m}_0$ (given by the first round) we conduct a worst case analysis to see how many of those tests may include one of the $x \in V_{0,\mathrm{PD}}$. Consider some particular truly negative test a. We are interested in the probability that none of the elements of $V_{0,\mathrm{PD}}$ is contained. The probability that a given individual $x \in V_{0,\mathrm{PD}}$ (knowing that it participates in $N_x \le \alpha\Delta$ displayed negative

tests, which is of lower order than m) is assigned to this test is given by[7]

(C.1)

$$\mathbb{P}\left(x \in \partial a | x \in V_{0,\mathrm{PD}}\right) = 1 - \mathbb{P}\left(x \notin \partial a | x \in V_{0,\mathrm{PD}}\right)$$

$$= 1 - \sum_{i=0}^{\alpha\Delta} \mathbb{P}\left(\boldsymbol{N}_x = i | x \in V_{0,PD}\right) \binom{m-1}{\Delta - i} \binom{m}{\Delta - i}^{-1} \tag{C.2}$$

$$\leq 1 - \left(1 + n^{-\Omega(1)}\right) \sum_{i=0}^{\alpha\Delta} \mathbb{P}\left(\boldsymbol{N}_x = i | x \in V_{0,PD}\right) \left(1 - \frac{1}{m}\right)^{\Delta - i} \tag{C.3}$$

$$\leq 1 - \left(1 + n^{-\Omega(1)}\right) \sum_{i=0}^{\alpha\Delta} \mathbb{P}\left(\boldsymbol{N}_x = i | x \in V_{0,PD}\right) \left(1 - \frac{1}{m}\right)^{\Delta} = \left(1 + n^{-\Omega(1)}\right) \left(\frac{\Delta}{m} + O(k^{-2})\right) = \frac{d}{k} + O(k^{-2}) \tag{C.4}$$

We can now calculate the probability that no individual $x \in V_{0,\mathrm{PD}}$ is assigned to a, bearing in mind that the size of $V_{0,\mathrm{PD}}$ is random, and that each such individual is assigned to tests mutually independently. Using (C.4), and decomposing the sum into two parts, this is given by (for a given V)

$$\begin{aligned}
\mathbb{P}\left(\left\{V_{0,\mathrm{PD}} \cap \partial a\right\} = \emptyset\right) &= \sum_{j=0}^{n} \mathbb{P}\left(\left|\boldsymbol{V_{0,\mathrm{PD}}}\right| = j\right) \mathbb{P}\left(\left\{V_{0,\mathrm{PD}} \cap \partial a\right\} = \emptyset \middle| \left|\boldsymbol{V_{0,\mathrm{PD}}}\right| = j\right) \\
&= \sum_{j=0}^{V} \mathbb{P}\left(\left|\boldsymbol{V_{0,\mathrm{PD}}}\right| = j\right) \left(1 - \frac{d}{k} + O\left(k^{-2}\right)\right)^{j} + \sum_{j=V+1}^{n} \mathbb{P}\left(\left|\boldsymbol{V_{0,\mathrm{PD}}}\right| = j\right) \left(1 - \frac{d}{k} + O\left(k^{-2}\right)\right)^{j} \\
&\geq \sum_{j=0}^{V} \mathbb{P}\left(\left|\boldsymbol{V_{0,\mathrm{PD}}}\right| = j\right) \left(1 - \frac{d}{k} + O\left(k^{-2}\right)\right)^{V} = \mathbb{P}\left(\left|\boldsymbol{V_{0,\mathrm{PD}}}\right| \leq V\right) \left(1 - \frac{d}{k} + O\left(k^{-2}\right)\right)^{V}
\end{aligned}$$

By Lemma C.1, we can choose $V = kn^{-\Omega(1)}$ such that $\mathbb{P}\left(\left|\boldsymbol{V_{0,\mathrm{PD}}}\right| \leq V\right)$ is arbitrarily close to 1, and knowing that $\left(1 - \frac{d}{k} + O\left(k^{-2}\right)\right)^{V} \simeq \exp(-dV/k) = \exp(-dn^{-\Omega(1)})$ we find

$$\mathbb{P}\left(\left\{V_{0,\mathrm{PD}} \cap \partial a\right\} = \emptyset\right) = 1 - n^{-\Omega(1)}.$$

By combining this with the findings of Lemma A.4 we find $\mathbb{E}\left[\boldsymbol{m}_{0,\mathrm{nd}}\right] = \left(1 - n^{-\Omega(1)}\right) e^{-d} m$. The lemma follows by a similar application of the McDiarmid inequality as used in the proof of Lemma A.4.

□

Note that, changing the set of tests for a specific individual $x \in V_1 \cup V_{0,\mathrm{PD}}$ shifts $\boldsymbol{m}_{0,\mathrm{nd}}$ by at most Δ. Thus, such an individual choosing from this set is not affecting the order of $\boldsymbol{m}_{0,\mathrm{nd}}$.
Let $\mathscr{F}$ be the event that $\boldsymbol{m}_{0,\mathrm{nd}} = \left(1 - n^{-\Omega(1)}\right) e^{-d} m$. By Lemma C.2, $\mathbb{P}(\mathscr{F}) = 1 - o(n^{-2})$ if

$$c > (1+\eta) \frac{1}{dD_{\mathrm{KL}}\left(\alpha \| e^{-d}(1-p) + \left(1 - e^{-d}\right) q\right)}$$

for some small $\eta > 0$. With Lemma C.2 at hand, we are in a position to describe the distribution of $\boldsymbol{P}_x$ for healthy and infected individuals (recall the definition of $\boldsymbol{P}_x$ in (2.2)). Let us start with infected individuals.

Lemma C.3. *Given $x \in V_1$ and conditioned on $\mathscr{F}$, the total variation distance between $\boldsymbol{P}_x$ and $\boldsymbol{Q}_H$, a random variable with hypergeometric distribution $H\left(m, me^{-d}(1-q), \Delta\right)$, tends to zero with n, that is*

$$d_{TV}\left(\boldsymbol{P}_x, \boldsymbol{Q}_H\right) = n^{-\Omega(1)}.$$

[7] We refer the reader to [20] for two results we use while obtaining (C.3) (apply Claim 7.3 to the binomial coefficients) as well as (C.4) (apply Claim 7.4 as error corrected version of Bernoulli's inequality). Please note that these bounds in particular hold for $\Delta = \Theta(\log(n))$ and $k \sim n^{\theta}$.

Proof. We are interested in the neighborhood structure of one given infected individual $x \in V_1$, and we check how the remaining individuals influence the test types. In particular we are interested in the number of tests $a \in F$ such that $\partial a \subset V_0 \setminus V_{0,\mathrm{PD}}$ are contained in the neighborhood of an infected individual x. Knowing the total number of tests m and fixed degree Δ, for a given value of the random quantity $\boldsymbol{m}_{0,\mathrm{nd}} = \ell$, we find that this quantity of interest follows a $H(m,\ell,\Delta)$-distribution. Given $\mathscr{F}$, Lemma C.2 gives that $\boldsymbol{m}_{0,\mathrm{nd}}$ is highly concentrated,

$$\boldsymbol{m}_{0,\mathrm{nd}} = \left(1 - n^{-\Omega(1)}\right) e^{-d} m$$

with high probability. Hence a conditioning argument, similar to Lemma B.2, shows that the linear combination of distribution functions

$$\sum_{\ell} \mathbb{P}(\boldsymbol{m}_{0,\mathrm{nd}} = \ell)\mathbb{P}(H(m,\ell,\Delta) \leq x)$$

tends to the distribution function of $H\left(m, me^{-d}, \Delta\right)$ in total variation distance, due to the concentration result obtained in Lemma C.2. Since each test featuring x will truly be positive (as we assume x to be infected) and will be displayed positive with probability $1-q$ independently, the lemma follows immediately. □

To describe the distribution of $\boldsymbol{P}_x$ for healthy individuals, let us introduce the random variable $\boldsymbol{P}_x(P)$, which is $\boldsymbol{P}_x$ conditioned on the individual appearing in P displayed positive tests, as follows:

$$\mathbb{P}\left(\boldsymbol{P}_x(P) = t\right) = \mathbb{P}\left(\boldsymbol{P}_x = t | \boldsymbol{N}_x = \Delta - P\right)$$

Then, we find for healthy individuals the following conditional distribution.

Lemma C.4. *Given $x \in V_0$,conditioned on $\mathscr{E}$ and $\mathscr{F}$, the total variation distance between $\boldsymbol{P}_x(P)$ and $\boldsymbol{B}_h \sim H\left(m\left(e^{-d}p + (1-e^{-d})(1-q)\right), m\left(e^{-d}p\right), P\right)$ tends to zero with n. That is*

$$d_{TV}(\boldsymbol{P}_x(P), \boldsymbol{B}_h) = n^{-\Omega(1)}.$$

Proof. We proceed with the same exposition and reasoning as in the proof of Lemma C.3. Due to the fact that x is healthy we can remove it without affecting the test result. Therefore we can analyse its neighborhood structure induced by the pooling graph while excluding it. Since by assumption individual $x \in V_0$ is assigned to exactly P displayed positive and the total number of displayed positive test is given by $\boldsymbol{m}_0^f + \boldsymbol{m}_1^u$, we see that $\boldsymbol{P}_x(P)$ is $H\left(\boldsymbol{m}_0^f + \boldsymbol{m}_1^u, \boldsymbol{m}_{0,\mathrm{nd}}, P\right)$-distributed. Due to the fact that the event $\mathscr{E}$ pinpoints the amount of displayed positive and negative tests we can derive the distribution of neighbors the individual may choose from. Recalling the results of Corollary A.5, we see that w.h.p.

$$\begin{aligned} \boldsymbol{m}_0^f &= e^{-d} pm + O\left(\sqrt{m}\log^3 n\right), \\ \text{and } \boldsymbol{m}_1^u &= (1-e^{-d})(1-q)m + O\left(\sqrt{m}\log^3 n\right). \end{aligned}$$

Furthermore we get from Lemma C.2 that w.h.p.

$$\boldsymbol{m}_{0,\mathrm{nd}} = \left(1 - n^{-\Omega(1)}\right) e^{-d} m.$$

Now we apply the concentration results obtained in Corollary A.5 and Lemma C.2 to obtain a linear combination of distribution functions

$$\sum_{\ell,\nu} \mathbb{P}(\boldsymbol{m}_{0,\mathrm{nd}} = \ell, \boldsymbol{m}_0^f + \boldsymbol{m}_1^u = \nu) \cdot \mathbb{P}(H(\nu,\ell,\Delta) \leq x)$$

that tends to $H\left(m\left(e^{-d}p + (1-e^{-d})(1-q)\right), me^{-d}, P\right)$. The lemma follows since truly negative tests get flipped independently with probability p. □

Having derived the distributions for $\boldsymbol{P}_x$ for $x \in V_1$ and $\boldsymbol{P}_x(P)$ for $x \in V_0$ we can now determine a threshold $\beta\Delta$ of displayed positive tests where the individual appears only with individuals from the set $V_0 \setminus V_{0,\mathrm{PD}}$ such that we can tell V_1 and $V_{0,\mathrm{PD}}$ apart and thus recover $\boldsymbol{\sigma}$. Let us start with infected individuals.

Lemma C.5. *As long as*

$$c > (1+\eta)\max\left\{\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}, \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\beta\|(1-q)e^{-d}\right)}\right\}$$

for some small $\eta > 0$, we have $\boldsymbol{P}_x > \beta\Delta$ for all $x \in V_1$ w.h.p.

Proof. We need to ensure that $\mathbb{P}(\exists x \in V_1 : \boldsymbol{P}_x < \beta\Delta) = o(1)$. For the bound on c from the lemma, we know that $\mathscr{F}$ occurs w.h.p. by Lemma C.2. In combination with Lemma C.3 and the union bound we need to ensure that

$$k \cdot \mathbb{P}\left(\boldsymbol{P}_x \le \beta\Delta | x \in V_1, \mathscr{F}\right) = k \cdot \mathbb{P}(\boldsymbol{Q}_H \le \beta\Delta) + kn^{-\Omega(1)} = o(1), \tag{C.5}$$

where as before $\boldsymbol{Q}_H$ is a random variable with hypergeometric distribution $H\left(m, me^{-d}(1-q), \Delta\right)$. Using the Chernoff bound for the hypergeometric distribution (Lemma A.2), the following condition for (C.5) to hold arises

$$\theta - cd(1-\theta)D_{\mathrm{KL}}\left(\beta\|(1-q)e^{-d}\right) < 0 \tag{C.6}$$

The lemma follows from rearranging terms in (C.6) and the fact that if we choose the number of tests slightly above the required number of tests (larger by a factor of $1+\eta$ for $\eta > 0$), the assertion holds w.h.p. as $n \to \infty$. □

We proceed with the set of individuals $V_{0,\mathrm{PD}}$.

Lemma C.6. *As long as*

$$c > (1+\eta)\max\left\{\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}, \max_{1-\alpha\le z\le 1}\left\{\frac{1}{1-\theta}\frac{1}{d\left(D_{\mathrm{KL}}\left(z\|e^{-d}p+(1-e^{-d})(1-q)\right)+zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{e^{-d}p+(1-e^{-d})(1-q)}\right)\right)}\right\}\right\}$$

for some small $\eta > 0$, we have $\boldsymbol{P}_x < \beta\Delta$ for all $x \in V_{0,PD}$ w.h.p.

Proof. We need to ensure that $\mathbb{P}(\exists x \in V_{0,PD} : \boldsymbol{P}_x > \beta\Delta) = o(1)$. For the bound on c from the lemma, we know that $\mathscr{F}$ occurs w.h.p. by Lemma C.2. Moreover, $\mathscr{E}$ occurs w.h.p. by Lemma A.4 and Corollary A.5. We write $w = e^{-d}p + \left(1-e^{-d}(1-q)\right)$ for brevity. Combining this fact with Lemma B.2 and C.4 we need to ensure

$$(n-k)\sum_{P=(1-\alpha)\Delta}^{\Delta}\mathbb{P}\left(\boldsymbol{N}_x = \Delta - P | x \in V_0, \mathscr{E}\right)\mathbb{P}\left(\boldsymbol{P}_x(P) \ge \beta\Delta | x \in V_0, \mathscr{F}\right) \tag{C.7}$$

$$= \left(1-n^{-\Omega(1)}\right)n\sum_{P=(1-\alpha)\Delta}^{\Delta}\mathbb{P}\left(\boldsymbol{T}_{\boldsymbol{h}} = P\right)\cdot\mathbb{P}\left(\boldsymbol{B}_{\boldsymbol{h}} \ge \beta\Delta\right) = o(1) \tag{C.8}$$

We remind ourselves that

$$\boldsymbol{T}_{\boldsymbol{h}} \sim H\left(m, m\left(e^{-d}(1-p)+\left(1-e^{-d}\right)q\right), \Delta\right)$$
$$\text{and} \quad \boldsymbol{B}_{\boldsymbol{h}} \sim H\left(m\left(e^{-d}p+(1-e^{-d})(1-q)\right), m\left(e^{-d}p\right), P\right).$$

Now by the Chernoff bound for the hypergeometric distribution (Lemma A.2) and setting $z = P/\Delta$, we establish the following two bounds for the probability terms:

$$\mathbb{P}\left(H\left(m, m\left(w+n^{-\Omega(1)}\right), \Delta\right) = P\right) = \exp\left(-(1+n^{-\Omega(1)})\Delta\left(D_{\mathrm{KL}}(z\|w)\right)\right) \tag{C.9}$$

$$\begin{aligned}&\mathbb{P}\left(H\left(m\left(w+n^{-\Omega(1)}\right), m\left(e^{-d}p+n^{-\Omega(1)}\right), P\right) \geq \beta\Delta\right)\\ &= \exp\left(-\left(1+n^{-\Omega}\right) z\Delta \mathbf{1}\left\{\beta > \frac{ze^{-d}p}{w}\right\} zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right)\end{aligned} \tag{C.10}$$

(Note that the indicator in (C.10) appears due to the condition given by Lemma A.2) We reformulate the left-hand-side of (C.8) to

$$\begin{aligned}&n\sum_{P=(1-\alpha)\Delta}^{\Delta} \exp\left(-(1+o(1))\Delta\left(D_{\mathrm{KL}}(z\|w) + \mathbf{1}\left\{\beta > \frac{ze^{-d}p}{w}\right\} zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right)\right)\\ &\quad = \left(1+n^{-\Omega(1)}\right) n \max_{1-\alpha\leq z\leq 1}\left\{\exp\left(-(1+o(1))\Delta\left(D_{\mathrm{KL}}(z\|w) + \mathbf{1}\left\{\beta > \frac{ze^{-d}p}{w}\right\} zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right)\right)\right\}\end{aligned}$$

where the second equality follows since the sum consists of $\Theta(\Delta) = \Theta(\log n)$ many summands. Since $\mathbb{P}(\mathscr{F}) = 1 - n^{-\Omega(1)}$ for our choice of c by Lemma C.2 rearranging terms readily yields that the expression in (C.7) is indeed of order $o(1)$.

To see this, we remind ourselves that by definition $\Delta = cd\log\left(\frac{n}{k}\right) = (1-\theta)cd\log(n)$. Furthermore we plug in the definition for $w = e^{-d}p + \left(1-e^{-d}(1-q)\right)$. In the end we have to ensure that

$$1 < (1-\theta)cd\left(D_{\mathrm{KL}}(z\|w) + \mathbf{1}\left\{\beta > \frac{ze^{-d}p}{e^{-d}p+\left(1-e^{-d}(1-q)\right)}\right\} zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{e^{-d}p+\left(1-e^{-d}(1-q)\right)}\right)\right)$$

We solve this inequality for c. As we are only interested in a worst case bound, the assertion follows from the non-negativity of $D_{\mathrm{KL}}(*\|*)$.

□

Proof of Theorem 2.2. The theorem is now immediate from Lemma B.3, C.1, C.5 and C.6 and the fact that the choice of α, β and d is at our disposal. □

Appendix D. Comparison of the noisy DD and COMP bounds

The following section is intented to provide sufficient conditions under which the DD algorithm attains reliable performance requiring fewer tests than the COMP. However, these conditions are not necessary and DD might (and for all performed simulations does) require fewer tests than COMP for even wider settings.

Proof of Proposition 2.13. In order to prove the proposition, we need to find conditions under which

$$\min_{\alpha,d}\max\{b_1(\alpha,d), b_2(\alpha,d)\} \geq \min_{\alpha,\beta,d}\max\left\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\alpha,\beta,d)\right\}$$

We write α^* and d^* for the values that minimise the maximum of the two terms at the LHS, at which point we know that $b_1(\alpha^*,d^*) = b_2(\alpha^*,d^*)$. Then it is sufficient to show that there exists β^* such that

$$b_1(\alpha^*,d^*) = b_2(\alpha^*,d^*) \geq \max\left\{c_1(\alpha^*,d^*), c_2(\alpha^*,d^*), c_3(\beta^*,d^*), c_4(\alpha^*,\beta^*,d^*)\right\}$$

By inspection for any α and d $b_1(\alpha,d) = c_1(\alpha,d)$ and $b_2(\alpha,d) \geq c_2(\alpha,d)$ since $\theta \in (0,1)$.

Next, we will show that $b_2(\alpha,d) \geq c_4(\alpha,\beta,d)$ for any α,β in the respective bounds and $d \in (0,\infty)$. Writing $w = e^{-d}p + (1-e^{-d})(1-q)$, and recalling that by assumption that $\alpha \leq 1-w$ (or $w \leq 1-\alpha$) we readily find that

$$D_{\mathrm{KL}}(\alpha\|1-w) = \min_{1-\alpha\leq z\leq 1}(D_{\mathrm{KL}}(z\|w)) \leq \min_{1-\alpha\leq z\leq 1}\left(D_{\mathrm{KL}}(z\|w) + z\mathbf{1}\left\{\beta > \frac{ze^{-d}p}{w}\right\}D_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right) \tag{D.1}$$

where the first equality follows since $D_{\mathrm{KL}}(\alpha\|1-w) = D_{\mathrm{KL}}(1-\alpha\|w)$ and $D_{\mathrm{KL}}(z\|w) > D_{\mathrm{KL}}(1-a\|w)$ for any $z > 1-\alpha$. The bound follows. Note that (D.1) indeed holds for any choice of α,β and d in the respective bounds stated in the theorem.

Finally, we need to demonstrate that $c_3(\beta^*,d^*) \leq b_2(\alpha^*,d^*)$. Since β is not an optimisation parameter in $b_2(\alpha^*,d^*)$ and the bound in (D.1) holds for any value of β, we can simply set it to the value that minimizes $c_3(\beta^*,d^*)$ which is $\beta = 1/\Delta$ and for which we find

$$c_3(\beta^*,d^*) = \frac{\theta}{1-\theta}\frac{1}{d^*\log\left(1-e^{-d^*}(1-q)\right.}.$$

Thus, to obtain the desired inequality we need to ensure that for the optimal choice α^* from COMP

$$\theta D_{\mathrm{KL}}\left(\alpha^*\|e^{-d^*}(1-p) + \left(1-e^{-d^*}\right)q\right) \leq -\log\left(1-e^{-d^*}(1-q)\right)$$

Using the bound

$$\begin{aligned}\theta D_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p) + \left(1-e^{-d}\right)q\right) &\leq -\theta\log\left(1-\left(e^{-d}(1-p) + \left(1-e^{-d}\right)q\right)\right)\\ &\leq -\log\left(1-\left(e^{-d}(1-p) + \left(1-e^{-d}\right)q\right)\right)\end{aligned}$$

which is obtained by setting $\alpha = 1/\Delta$, we find that $c_3(\beta^*,d^*) \leq b_2(\alpha^*,d^*)$ if

$$-\log\left(1-e^{-d^*}(1-q)\right) \geq -\log\left(1-e^{-d^*}(1-p) + \left(1-e^{-d^*}\right)q\right) \Leftrightarrow e^{-d^*}p \geq q$$

□

As mentioned before, due to bounding $b_2(\alpha^*,d^*)$ the result is not sharp. However, one immediate consequence of Proposition 2.13 is that DD is guaranteed to require fewer tests than COMP for the reverse Z channel.

Appendix E. Relation to Bernoulli testing

In the noiseless case [26] shows that the constant column weight design (where each individual joins exactly Δ different tests) requires fewer tests to recover $\boldsymbol{\sigma}$ than the i.i.d. (Bernoulli pooling) design (where each individual is included in each test with a certain probability independently). In this section we show that in the noisy case, the COMP algorithm requires fewer tests for the constant column weight design than for the i.i.d. design, and derive sufficient conditions under which the same is true for the noisy DD algorithm.

To get started, let us state the relevant bounds for the Bernoulli design, taken from [48, Theorem 5] and rephrased in our notation.

Proposition E.1 (Noisy COMP under Bernoulli)**.** *Let $p,q \geq 0$, $p+q<1$, $d \in (0,\infty)$, $\alpha \in (q, e^{-d}(1-p) + \left(1-e^{-d}\right)q)$. Suppose that $0<\theta<1$ and $\varepsilon>0$ and let*

$$m^{Ber}_{COMP} = m^{Ber}_{COMP}(n,\theta,p,q) = \min_{\alpha,d}\max\{b_1(\alpha,d), b_2(\alpha,d)\}\,k\log(n/k)$$

$$\text{where} \qquad b_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{kD_{\mathrm{KL}}\left(\alpha d/k\|qd/k\right)}$$

$$\text{and} \qquad b_2(\alpha,d) = \frac{1}{1-\theta}\frac{1}{kD_{\mathrm{KL}}\left(\alpha d/k\|(e^{-d}(1-p)+(1-e^{-d})q)d/k\right)}$$

If $m > (1+\varepsilon)m^{Ber}_{COMP}$, COMP *will recover* $\boldsymbol{\sigma}$ *under the Bernoulli test design w.h.p. given* $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$.

Proposition E.2 (Noisy DD under Bernoulli)**.** *Let* $p, q \geq 0$, $p+q<1$, $d \in (0,\infty)$, $\alpha \in (q, e^{-d}(1-p)+\left(1-e^{-d}\right)q)$ *and* $\beta \in (e^{-d}p, e^{-d}(1-q))$. *Suppose that* $0<\theta<1, \zeta \in (0,\theta)$ *and* $\varepsilon > 0$ *and let*

$$m^{Ber}_{DD} = m^{Ber}_{DD}(n,\theta,p,q) = \min_{\alpha,\beta,d} \max\left\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\beta,d)\right\} k\log(n/k)$$

$$\begin{aligned}
&\textit{where} \quad c_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{kD_{\mathrm{KL}}\left(\alpha d/k \| q d/k\right)}\\
&\textit{and} \quad c_2(\alpha,d) = \frac{1-\zeta}{1-\theta}\frac{1}{kD_{\mathrm{KL}}\left(\alpha d/k \| (e^{-d}(1-p)+(1-e^{-d})q) d/k\right)}\\
&\textit{and} \quad c_3(\beta,d) = \frac{\theta}{1-\theta}\frac{1}{k\cdot D_{\mathrm{KL}}\left(\beta d/k \| e^{-d}(1-q) d/k\right)}\\
&\textit{and} \quad c_4(\beta,d) = \frac{\zeta}{1-\theta}\frac{1}{k\cdot D_{\mathrm{KL}}\left(\beta d/k \| e^{-d} p d/k\right)}
\end{aligned}$$

If $m > (1+\varepsilon)m^{Ber}_{DD}$, DD *will recover* $\boldsymbol{\sigma}$ *under the Bernoulli test design w.h.p. given* $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$.

To compare the bounds of the Bernoulli and constant-column test design we employ the following handy observation.

Lemma E.3. *Let* $0<x,y<1$ *and* $d>0$ *be constants independent of* k. *As* $k\to\infty$

$$kD_{\mathrm{KL}}\left(\frac{xd}{k}\|\frac{yd}{k}\right) = d\left(D_{\mathrm{KL}}\left(x\|y\right)+v(x,y)\right)+o(1/k)$$

with

$$v(x,y) = y-x+(1-x)\log\left(\frac{1-y}{1-x}\right) \leq 0 \tag{E.1}$$

Proof. Applying the definition of the Kullback-Leibler divergence and Taylor expanding the logarithm we obtain

$$\begin{aligned}
k\cdot D_{\mathrm{KL}}\left(\frac{xd}{k}\|\frac{yd}{k}\right) &= xd\cdot\log\left(\frac{x}{y}\right)+(k-xd)\left(\log\left(1-\frac{xd}{k}\right)-\log\left(1-\frac{yd}{k}\right)\right)\\
&= xd\cdot\log\left(\frac{x}{y}\right)+(k-xd)\left(-\frac{xd}{k}+\frac{yd}{k}+o\left(\frac{1}{k^2}\right)\right)\\
&= d\left(x\cdot\log\left(\frac{x}{y}\right)-x+y\right)+o(1/k)\\
&= d\left(D_{\mathrm{KL}}\left(x\|y\right)+y-x-(1-x)\log\left(\frac{1-x}{1-y}\right)\right)+o(1/k).
\end{aligned}$$

We can bound $v(x,y)$ from above by writing the final term as $(1-x)\log\left(1+\frac{x-y}{1-x}\right) \leq (1-x)\frac{x-y}{1-x} = x-y$, using the standard linearisation of the logarithm. □

We are now in a position to prove Proposition 2.15 and 2.16.

Proof of Proposition 2.15. The lemma follows by comparing the bounds from Theorem 2.1 and Proposition E.1 and applying Lemma E.3. □

Proof of Proposition 2.16. As evident from Corollary 2.8, the fourth bound $c_4(\alpha,\beta,d)$ vanishes under the Z channel. Now comparing the bounds from Theorem 2.2 and Proposition E.2, observing that $(1-\zeta)/(1-\theta)>1$ for $\zeta<\theta$ and applying Lemma E.3 immediately implies the lemma. □

APPENDIX F. NOTES ON COROLLARY 2.3

Lemma F.1. *If $p+q<1$ the Shannon capacity of the $p-q$ channel of Figure 1 measured in nats is*

$$C_{Chan} = D_{\mathrm{KL}}\left(q\|\frac{1}{1+e^{\phi}}\right) = D_{\mathrm{KL}}\left(p\|\frac{1}{1+e^{-\phi}}\right), \tag{F.1}$$

where $\phi = (h(p)-h(q))/(1-p-q)$. This is achieved by taking

$$\mathbb{P}(X=0) = \frac{1}{1-p-q}\left(\frac{1}{1+e^{\phi}} - q\right). \tag{F.2}$$

Please note that the proof might be a standard result for readers from some research communities, but for others it might be less standard. Therefore we state it here to prevent the interested (but unfamiliar) reader from a long textbook search.

Proof. Write $\mathbb{P}(X=0)=\gamma$ and $\mathbb{P}(Y=0)=T(\gamma) := (1-p)\gamma + q(1-\gamma)$. Then since the mutual information

$$I(X;Y) = h(Y) - h(Y|X) = h\big(T(\gamma)\big) - \big(\gamma h(p) + (1-\gamma)h(q)\big), \tag{F.3}$$

we can find the optimal T by solving

$$0 = \frac{\partial}{\partial\gamma} I(X;Y) = (1-p-q)\log\left(\frac{1-T(\gamma)}{T(\gamma)}\right) - \big(h(p)-h(q)\big),$$

which implies that the optimal $T^* = 1/(1+e^{\phi})$. We can solve for this for $\gamma^* = (T^*-q)/(1-p-q)$ to find the expression above. As $\frac{\partial}{\partial^2\gamma} I(X;Y) < 0$ it is indeed a maximum. Substituting this in (F.3) we obtain that the capacity is given by

$$\begin{aligned} h(T^*) - \big(\gamma^* h(p) + (1-\gamma^*)h(q)\big) &= h\left(\frac{1}{1+e^{\phi}}\right) - \big((T^*-q)\phi + h(q)\big) \\ &= \log(1+e^{\phi}) - \phi(1-q) - h(q) \\ &= D_{\mathrm{KL}}\big(q\|1/(1+e^{\phi})\big) \end{aligned} \tag{F.4}$$

as claimed in the first expression in (F.1) above. We can see that the second expression in (F.1) matches the first by writing the corresponding expression as $D_{\mathrm{KL}}\big(1-p\|1/(1+e^{\phi})\big) = \log(1+e^{\phi}) - \phi p - h(p)$, which is equal to (F.4) by the definition of ϕ. □

Note that this result suggests a choice of density for the matrix: since each test is negative with probability e^{-d}, equating this with (F.2) suggests that we take

$$d = d^*_{\mathrm{ch}} = \log(1-p-q) - \log\left(\frac{1}{1+e^{\phi}} - q\right).$$

This is unlikely to be optimal in a group testing sense, since we make different inferences from positive and negative tests, but gives a closed form expression that may perform well in practice. For the noiseless and BSC case observe that $\phi = 0$, and we obtain $d^*_{\mathrm{ch}} = \log 2$.

APPENDIX G. ILLUSTRATION OF BOUNDS FOR Z, REVERSE Z CHANNEL AND THE BSC

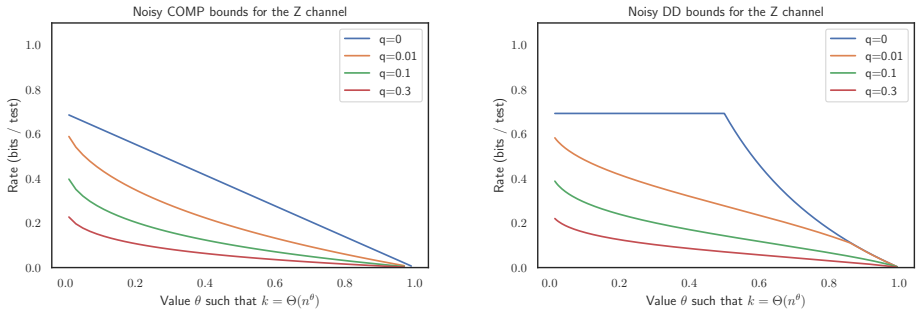

FIGURE 5. Illustration of achievability bounds for noisy COMP and DD under the Z channel. (Note for black and white prints: The solid lines as well as the dashed lines in the diagram are in the same order as given in the legend from top to bottom)

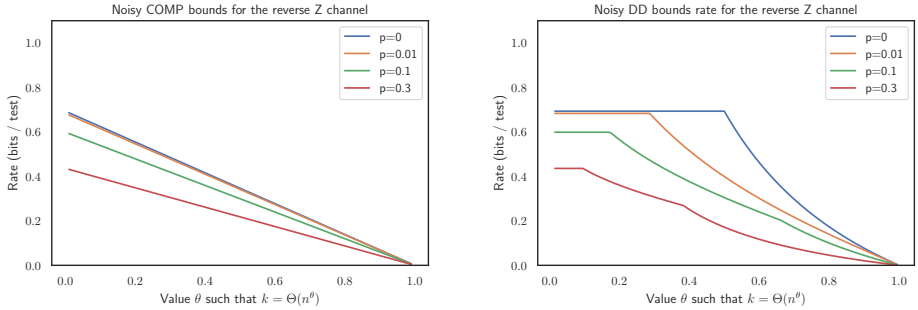

FIGURE 6. Illustration of achievability bounds for noisy COMP and DD under the reverse Z channel. (Note for black and white prints: The solid lines as well as the dashed lines in the diagram are in the same order as given in the legend from top to bottom)

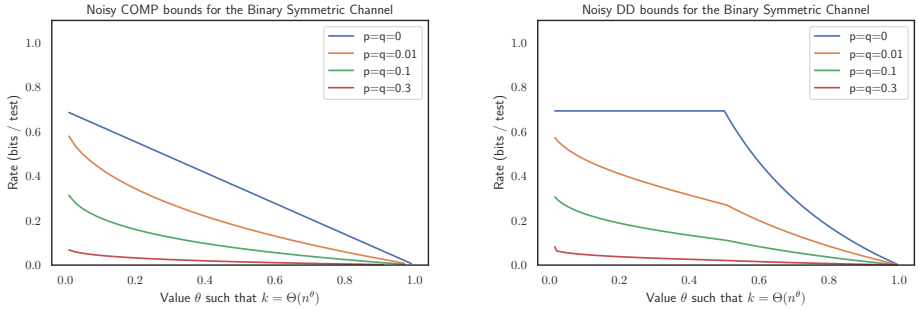

FIGURE 7. Illustration of achievability bounds for noisy COMP and DD under the Binary Symmetric Channel. (Note for black and white prints: The solid lines as well as the dashed lines in the diagram are in the same order as given in the legend from top to bottom)

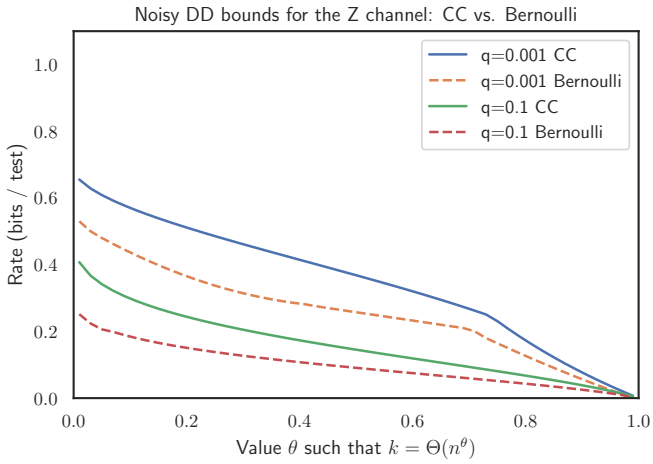

FIGURE 8. Comparison of the noisy DD rates under Bernoulli pooling ([48]) with the DD bounds with constant-column design as provided in the paper at hand within the Z-Channel.(Note for black and white prints: The solid lines as well as the dashed lines in the diagram are in the same order as given in the legend from top to bottom).

ACKNOWLEDGMENT

The authors would like to thank two anonymous referees for their detailed reading of this paper and for the suggestions they made to improve its presentation. Oliver Gebhard and Philipp Loick are supported by DFG CO 646/3.

REFERENCES

[1] E. Abbe, A. Bandeira, and G. Hall. Exact recovery in the stochastic block model. *IEEE Transactions on Information Theory*, 62:471–487, 2016.

[2] B. Abdalhamid, C. Bilder, E. McCutchen, S. Hinrichs, S. Koepsell, and P. Iwen. Assessment of specimen pooling to conserve SARS-CoV-2 testing resources. *American Journal of Clinical Pathology*, 153:715–718, 2020.

[3] M. Aldridge. The capacity of Bernoulli nonadaptive group testing. *IEEE Transactions on Information Theory*, 63:7142–7148, 2017.

[4] M. Aldridge. Individual testing is optimal for nonadaptive group testing in the linear regime. *IEEE Transactions on Information Theory*, 65:2058–2061, 2019.

[5] M. Aldridge, L. Baldassini, and O. Johnson. Group testing algorithms: bounds and simulations. *IEEE Transactions on Information Theory*, 60:3671–3687, 2014.

[6] M. Aldridge, O. Johnson, and J. Scarlett. Improved group testing rates with constant column weight designs. *Proceedings of 2016 IEEE International Symposium on Information Theory (ISIT'16)*, pages 1381–1385, 2016.

[7] M. Aldridge, O. Johnson, and J. Scarlett. Group testing: an information theory perspective. *Foundations and Trends in Communications and Information Theory*, 15(3–4):196–392, 2019.

[8] E. ArÄ±kan. Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memory-less channels. *IEEE Transactions on Information Theory*, 55:3051â€“3073, 2009.

[9] L. Baldassini, O. Johnson, and M. Aldridge. The capacity of adaptive group testing. *Proceedings of 2013 IEEE International Symposium on Information Theory (ISIT'13)*, 1:2676–2680, 2013.

[10] C. Chan, P. Che, S. Jaggi, and V. Saligrama. Non-adaptive probabilistic group testing with noisy measurements: near-optimal bounds with efficient algorithms. *Proceedings of 49th Annual Allerton Conference on Communication, Control, and Computing*, 1:1832–1839, 2011.

[11] H. Chen and F. Hwang. A survey on nonadaptive group testing algorithms through the angle of decoding. *Journal of Combinatorial Optimization*, 15:49–59, 2008.

[12] I. Cheong. The experience of South Korea with COVID-19. *Mitigating the COVID Economic Crisis: Act Fast and Do Whatever It Takes (CEPR Press)*, pages 113–120, 2020.

[13] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick. Information-theoretic and algorithmic thresholds for group testing. *IEEE Transactions on Information Theory*, DOI: 10.1109/TIT.2020.3023377, 2020.
[14] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick. Optimal group testing. *Proceedings of 33rd Conference on Learning Theory (COLT'20)*, 2020.
[15] D. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52:1289–1306, 2006.
[16] R. Dorfman. The detection of defective members of large populations. *Annals of Mathematical Statistics*, 14:436–440, 1943.
[17] S. Ciesek E. Seifried. Pool testing of SARS-CoV-02 samples increases worldwide test capacities many times over, 2020. https://www.bionity.com/en/news/1165636/pool-testing-of-sars-cov-02-samples-increases-worldwide-test-capacities-many-times-over.html, last accessed on 2020-11-16.
[18] Y. Erlich, A. Gilbert, H. Ngo, A. Rudra, N. Thierry-Mieg, M. Wootters, D. Zielinski, and O. Zuk. Biological screens from linear codes: theory and tools. *bioRxiv*, page 035352, 2015.
[19] European Centre for Disease Prevention and Control. Surveillance and studies in a pandemic in Europe, 2009. https://www.ecdc.europa.eu/en/publications-data/surveillance-and-studies-pandemic-europe (last accessed on 2020-11-16).
[20] Oliver Gebhard, Max Hahn-Klimroth, Olaf Parczyk, Manuel Penschuck, Maurice Rolvien, Jonathan Scarlett, and Nelvin Tan. Near optimal sparsity-constrained group testing: improved bounds and algorithms. *Arxiv-Preprint*, 2021.
[21] Y. Gefen, M. Szwarcwort-Cohen, and R. Kishony. Pooling method for accelerated testing of COVID-19, 2020. https://www.technion.ac.il/en/2020/03/pooling-method-for-accelerated-testing-of-covid-19/ (last accessed on 2020-11-16).
[22] E. Gould. Methods for long-term virus preservation. *Mol Biotechnol*, 13:57–66, 1999.
[23] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:301:13–30, 1963.
[24] F. Hwang. A method for detecting all defective members in a population by group testing. *Journal of the American Statistical Association*, 67:605–608, 1972.
[25] S. Janson, T. Luczak, and A. Rucinski. *Random Graphs*. John Wiley and Sons, 2011.
[26] O. Johnson, M. Aldridge, and J. Scarlett. Performance of group testing algorithms with near-constant tests per item. *IEEE Transactions on Information Theory*, 65:707–723, 2018.
[27] O. Johnson and D. Sejdinovic. Note on noisy group testing: Asymptotic bounds and belief propagation reconstruction. *Proceedings of 48th Allerton Conference on Communication, Control, and Computing*, 2010.
[28] E. Knill, A. Schliep, and D. Torney. Interpretation of pooling experiments using the Markov chain Monte Carlo method. *Journal of Computational Biology*, 3:395â€“406, 1996.
[29] H. Kwang-Ming and D. Ding-Zhu. Pooling designs and nonadaptive group testing: important tools for dna sequencing. *World Scientific*, 2006.
[30] A. Lalkhen. Clinical tests: sensitivity and specificity. *Continuing Education in Anaesthesia Critical Care & Pain*, 8, 2008.
[31] S. Long, C. Prober, and M. Fischer. Principles and practice of pediatric infectious diseases. *Elsevier*, 2018.
[32] N. Madhav, B. Oppenheim, M. Gallivan, P. Mulembakani, E. Rubin, and N. Wolfe. Pandemics: Risks, impacts and mitigation. *The World Bank:Disease control priorities*, 9:315–345, 2017.
[33] D. M. Malioutov and M. Malyutov. Boolean compressed sensing: Lp relaxation for group testing. *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, 2012.
[34] C. McDiarmid. On the method of bounded differences. *Surveys in Combinatorics, 1989: Invited Papers at the 12th British Combinatorial Conference*, page 148â€“188, 1989.
[35] R. Mourad, Z. Dawy, and F. Morcos. Designing pooling systems for noisy high-throughput protein-protein interaction experiments using Boolean compressed sensing. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 10:1478–1490, 2013.
[36] L. Mutesa, P. Ndishimye, Y. Butera, J. Souopgui, A. Uwineza, R. Rutayisire, E. Musoni, N. Rujeni, T. Nyatanyi, E. Ntagwabira, M. Semakula, C. Musanabaganwa, D. Nyamwasa, M. Ndashimye, E. Ujeneza, I. Mwikarago, C. Muvunyi, J. Mazarati, S. Nsanzimana, N. Turok, and W. Ndifon. A strategy for finding people infected with SARS-CoV-2: optimizing pooled testing at low prevalence. *Nature*, 589:276â€“280, 2021. doi:10.1038/s41586-020-2885-5.
[37] H. Ngo and D. Du. A survey on combinatorial group testing algorithms with applications to DNA library screening. *Discrete Mathematical Problems with Medical Applications*, 7:171–182, 2000.
[38] U.S. Department of Health and Human Services. Pandemic influenza plan, 2017. https://www.cdc.gov/flu/pandemic-resources/pdf/pandemic-influenza-implementation.pdf (last accessed on 2020-11-16).
[39] World Health Origanisation. Global surveillance during an influenza pandemic, 2009. www.who.int/csr/resources/publications/swineflu (last accessed on 2020-11-16).
[40] M. Plebani. Diagnostic errors and laboratory medicine â€“ causes and strategies. *Electronic Journal of the International Federation of Clinical Chemistry and Laboratory Medicine*, 26:7–14, 2015.
[41] T. Richardson and R. Urbanke. Modern coding theory. *Cambridge University Press*, 2007.
[42] C. Sammut and G. Webb. Encyclopedia of machine learning. *Springer*, 2011.

[43] J. Scarlett. Noisy adaptive group testing: Bounds and algorithms. *IEEE Transactions on Information Theory*, 65:3646–3661, 2018.

[44] J. Scarlett. An efficient algorithm for capacity-approaching noisy adaptive group testing. *Proceedings of 2019 IEEE International Symposium on Information Theory (ISIT'19)*, pages 2679–2683, 2019.

[45] J. Scarlett and V. Cevher. Converse bounds for noisy group testing with arbitrary measurement matrices. *Proceedings of 2016 IEEE International Symposium on Information Theory (ISIT'16)*, pages 2868–2872, 2016.

[46] J. Scarlett and V. Cevher. Phase transitions in group testing. *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms(SODA'16)*, 1:40–53, 2016.

[47] J. Scarlett and V. Cevher. Near-optimal noisy group testing via separate decoding of items. *IEEE Journal of Selected Topics in Signal Processing*, 2017.

[48] J. Scarlett and O. Johnson. Noisy non-adaptive group testing: A (near-)definite defectives, approach. *IEEE Transactions on Information Theory*, 66(6):3775–3797, 2020.

[49] N. Thierry-Mieg. A new pooling strategy for high-throughput screening: the shifted transversal design. *BMC Bioinformatics*, 7:28, 2006.

[50] L. Wang, X. Li, Y. Zhang, and K. Zhang. Evolution of scaling emergence in large-scale spatial epidemic spreading. *Public Library of Science ONE*, 6, 2011.

[51] L. Wein and S. Zenios. Pooled testing for HIV screening: Capturing the dilution effect. *Operations Research*, 44:543â€“569, 1996.

[52] S. Woloshin, N. Patel, and A. Kesselheim. False negative tests for SARS-CoV-2 infection â€" challenges and implications. *New England Journal of Medicine*, 2020.

APPENDIX E. ON THE PARALLEL RECONSTRUCTION FROM POOLED DATA

On the Parallel Reconstruction from Pooled Data

Oliver Gebhard
TU Dortmund University
Dortmund, Germany
oliver.gebhard@tu-dortmund.de

Max Hahn-Klimroth
TU Dortmund University
Dortmund, Germany
maximilian.hahnklimroth@tu-dortmund.de

Dominik Kaaser
Universität Hamburg
Hamburg, Germany
dominik.kaaser@uni-hamburg.de

Philipp Loick
Goethe University Frankfurt
Frankfurt, Germany
loick@math.uni-frankfurt.de

***Abstract*—In the pooled data problem the goal is to efficiently reconstruct a binary signal from additive measurements. Given a signal $\boldsymbol{\sigma} \in \{0,1\}^n$, we can query multiple entries at once and get the total number of non-zero entries in the query as a result. We assume that queries are time-consuming and therefore focus on the setting where all queries are executed in parallel. For the regime where the signal is sparse such that $||\boldsymbol{\sigma}||_1 = o(n)$ our results are twofold: First, we propose and analyze a simple and efficient greedy reconstruction algorithm. Secondly, we derive a sharp information-theoretic threshold for the minimum number of queries required to reconstruct $\boldsymbol{\sigma}$ with high probability. Our first result matches the performance guarantees of much more involved constructions (Karimi et al. 2019). Our second result extends a result of Alaoui et al. (2014) and Scarlett & Cevher (2017) who studied the pooled data problem for dense signals. Finally, our theoretical findings are complemented with empirical simulations. Our data not only confirm the information-theoretic thresholds but also hint at the practical applicability of our pooling scheme and the simple greedy reconstruction algorithm.**

***Index Terms*—Reconstruction, Sparse Signal, Pooled Data, Information Theory, Phase Transitions**

I. INTRODUCTION

We consider the *binary pooled data problem with additive queries* which is defined as follows. We are given a *signal* of length n, a large vector $\boldsymbol{\sigma} \in \{0,1\}^n$ of Hamming weight k and a querying method. Each query pools multiple entries of $\boldsymbol{\sigma}$ together and returns the exact number of non-zero entries contained in the pool (see Fig. 1 for an example). The goal is to reconstruct $\boldsymbol{\sigma}$ using as few queries as possible.

In many real-world scenarios the time to compute a reconstruction of $\boldsymbol{\sigma}$ is dominated by the time to perform a single query. The evaluation of such a query may require, e.g., computations using a deep neural network on a GPU [20], biological processes such as DNA screening [7], [26], or PCR tests in a bio-medical context [4]. To obtain a substantial speed-up, we therefore focus on *parallel* schemes where all queries are specified a priori and executed simultaneously. This assumption makes sense in the context of a life sciences laboratory: queries can be envisioned as measurements conducted by a liquid handling robot. The time to perform all (parallel) queries then clearly dominates the time to run an efficient (sequential) reconstruction algorithm (for practical input sizes).

OG and PL were supported by DFG CO 646/3. MHK was supported by DFG FOR 2975 and Stiftung Polytechnische Gesellschaft.

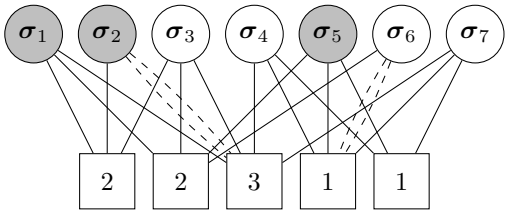

Fig. 1. A small example with signal $\boldsymbol{\sigma} = (1,1,0,0,1,0,0) \in \{0,1\}^7$ at the top and queries $a_1, \dots, a_5$ at the bottom. The edges of the bipartite (multi-) graph $\boldsymbol{G}$ show which entries are contained in a specific query. The dashed lines highlight the occurrence of multi-edges. The goal is to reconstruct $\boldsymbol{\sigma}$ given only $\boldsymbol{G}$ and the query results $(2,2,3,1,1)$.

In this paper we focus on the *sublinear regime* where the number of non-zero entries k scales sub-linearly in the signal's length n such that $k = n^\theta$ for some $\theta < 1$. In this setting, our main task is to specify a suitable parallel pooling design and an efficient reconstruction algorithm that allows us to compute $\boldsymbol{\sigma}$ efficiently from the queried data. We are interested in two different types of *phase-transitions* that commonly arise in the analysis of reconstruction and statistical inference problems:

1) What is the minimum number of queries that allows us to infer $\boldsymbol{\sigma}$ from the query results given unlimited computational power?
2) How many queries are required such that an efficient algorithm can compute $\boldsymbol{\sigma}$ from the query results?

We will refer to the first phase-transition as the *information-theoretic threshold* and to the second phase-transition as the *algorithmic threshold*.

A. The Teacher-Student Model

As in many related reconstruction problems, the teacher-student model provides the fundamental means towards analyzing information-theoretic questions. The challenge in such reconstruction problems lies in deriving probability distributions that are dependent on a variety of random variables

and hard to express per se. However, deriving probability distributions conditioned on certain high-probability events is feasible. For an introduction and mathematical justification of the model, we refer the reader to [10]. The setup is the following: a teacher aims to convey some *ground truth* to a student. Rather than directly providing the ground truth to the student, the teacher generates observable data from the ground truth via some statistical model and passes both the data and the model to the student. The student now aims to infer the ground truth from the observed data and the model.

In terms of this paper we see $\boldsymbol{\sigma}$ as the ground truth. Its distribution is inherited from all vectors in $\{0,1\}^n$ of Hamming weight k. The observable data $\boldsymbol{y}$, together with the conducted queries (expressed as a graph $\boldsymbol{G}$) are passed to the student in order to infer $\boldsymbol{\sigma}$. In the following, we analyze the chances of the student to infer the ground truth from the observable data. First, we derive the model distribution from the provided information $\boldsymbol{G}$ and the query results $\boldsymbol{y}$. Afterwards, we use the gained knowledge to analyze the chances of the student to recover the ground truth by estimating the number of possible input vectors that are consistent with the observed query results. As our goal is to recover $\boldsymbol{\sigma}$ with high probability, we condition on the event that the underlying bipartite multi-graph $\boldsymbol{G}$, which will be defined properly in due course, behaves almost *as expected*. We exploit the knowledge about $\boldsymbol{G}$ to derive high-probability events which we can condition on. Eventually, our analysis conveys the information whether there is a unique input vector or multiple possible input vectors out of which the student has to guess the correct one.

B. Related Work

The binary pooled data problem, sometimes called *quantitative group testing*, finds its roots in early works of Dorfman [13], Djackov [11], and Shapiro [27]. It has recently gained a lot of interest in the literature [1], [6], [14], [18], [25], with applications in a multitude of disciplines such as DNA screening [26], identifying genetic carriers [7] and machine learning [20], [23], [33]. Variants of the problem include binary group testing [2], [9] or threshold group testing [8], [22]. We start our discussion with an overview of related work from information theory.

Information-Theoretic Aspects. A simple information-theoretic lower bound can be obtained by a folklore counting argument: each query returns a number from 0 to k, thus a pooling design with m queries can produce at most $(k+1)^m$ different outcomes. This number must be larger than $\binom{n}{k}$ in order to distinguish all possible input vectors of length n with Hamming weight k. By standard asymptotic bounds, we obtain

$$m_{\text{seq}}^{\text{BPD}} \geq (1-o(1))\,\frac{\ln \frac{n}{k}}{\ln k}k. \tag{1}$$

The universal lower bound on $m_{\text{seq}}^{\text{BPD}}$ holds in any case, even if the queries do not need to be conducted in parallel. Restricted to the important special case in which all queries are conducted in parallel, [11] shows that reconstruction of $\boldsymbol{\sigma}$ requires at least

$$m_{\text{para}}^{\text{BPD}} = (2-o(1))\frac{\ln \frac{n}{k}}{\ln k}k = 2\cdot m_{\text{seq}}^{\text{BPD}} \tag{2}$$

queries, even with unlimited computational power. On the positive side, Bshouty [6] proves that reconstruction of $\boldsymbol{\sigma}$ is efficiently possible with $(2+\varepsilon)m_{\text{seq}}^{\text{BPD}}$ queries if they are conducted sequentially and Grebinski and Kucherov [17] provide a parallelizable design with an exponential-time reconstruction decoding algorithm which guarantees inference with $(2+\varepsilon)m_{\text{para}}^{\text{BPD}}$ queries using *separating matrices*. The latter positive result was extended to the so-called *Subset Select problem* [21], a relaxation of the pooled data problem that asks to identify only a subset of positive entries correctly. Recently, [14] improved the result for this relaxation by a factor of 2. So far, these results hold independently of k. For the linear regime where $k=\Theta(n)$, much stricter results are already known: Alaoui et al. [1] and Scarlett and Cevher [25] show that there is an exponential-time construction that achieves reconstruction with $(1+\varepsilon)m_{\text{para}}^{\text{BPD}}$ parallel queries – a result that is dependent on k scaling linearly in n.

Algorithmic Aspects. If allowed for sequential queries, Bshouty [6] presents an efficient reconstruction algorithm that succeeds at recovery of $\boldsymbol{\sigma}$ with no more than $(2+o(1))m_{\text{seq}}^{\text{BPD}}$ queries. However, for parallel schemes, there are significant gaps between the information-theoretic lower bound and the currently best known efficient algorithms [1], [12], [14], [15], [19], [24]. For instance, Alaoui et al. [1] present an *Approximate Message Passing* algorithm for dense signals ($k=\Theta(n)$). Furthermore, Donoho and Tanner [12] give a decoding strategy based on ℓ_1-minimization, and Foucart and Rauhut [15] introduce the *Basis Pursuit*-algorithm. They can be used to recover $\boldsymbol{\sigma}$ with

$$(2+o(1))k\ln\frac{n}{k} \quad \text{and} \quad (2+o(1))k\ln n \sim \frac{2}{1-\theta}k\ln\frac{n}{k}$$

queries, respectively, if the signal is sparse ($k \ll n$). Note that these algorithms solve the more general compressed sensing problem. Various improvements over the Basis Pursuit algorithm are known (e.g., the Orthogonal Matching Pursuit [24] and its improved version for discrete signals [29]) but as Wang and Yin [32] discuss, they do not perform asymptotically better in the setting discussed in this paper. More recent algorithms explicitly designed for recovery of $\boldsymbol{\sigma}$ from additive queries in the sparse regime are due to Karimi et al. [18], [19]. They provide two algorithms based on graph codes that require

$$(1.72+o(1))k\ln\frac{n}{k} \quad \text{and} \quad (1.515+o(1))k\ln\frac{n}{k}$$

queries, respectively. Furthermore, in a yet unpublished draft that appeared subsequently to our work on arXiv, Feige and Lellouche [14] analyze the Subset Select problem. They prove that, under mild assumptions, an algorithm succeeding at this relaxation can be turned into an algorithm for recovery of $\boldsymbol{\sigma}$ without significantly increasing the required number of queries.

C. Our Contributions

We study the pooled data problem under the random regular model $\boldsymbol{G}$ which is known to be information-theoretically optimal in the linear regime as well as in similar inference problems [9]. More precisely, we let $\boldsymbol{G} = (V \cup F, E)$ be a random bipartite multi-graph with *query-nodes* $F = \{a_1, \dots, a_m\}$ representing the queries, *entry-nodes* $V = \{x_1, \dots, x_n\}$ representing the coordinates of $\boldsymbol{\sigma}$, and edges E indicating how often a specific entry is contained in a given query. Hereby, each query $a_i \in F$ contains exactly $\Gamma = n/2$ entries chosen uniformly at random with replacement.

Algorithmic Results. For the aforementioned pooling design we present a fairly intuitive greedy algorithm called *Maximum Neighborhood (MN) Algorithm* that allows reconstruction of $\boldsymbol{\sigma}$ w.h.p.[1] It follows a thresholding approach that is much simpler than the known algorithms by Karimi et al. [18], [19], which are technically highly challenging. A formal definition of the MN-Algorithm is given in Algorithm 1.

Algorithm 1: The Maximum Neighborhood Algorithm

Input: m, k, querying method `query`
Output: estimation $\tilde{\boldsymbol{\sigma}}$ for $\boldsymbol{\sigma}$.

1 **for** $i = 1$ **to** m **do in parallel**
2 sample a multiset a_i of size Γ from $[n]$
3 compute $\boldsymbol{y}_i \leftarrow$ `query` (a_i)
   // The query method guarantees that $\boldsymbol{y}_i = \sum_{j \in a_i} \boldsymbol{\sigma}(j)$.
4 **for** $i = 1$ **to** n **do**
5 calculate $\Psi_i \leftarrow \sum_{j=1}^{m} \mathbf{1}\{i \in a_j\} \cdot \boldsymbol{y}_j$
6 calculate $\boldsymbol{\Delta}_i^\star \leftarrow \sum_{j=1}^{m} \mathbf{1}\{i \in a_j\}$
7 sort coordinates of $\tilde{\boldsymbol{\sigma}}$ in decreasing order by $\Psi_i - \boldsymbol{\Delta}_i^\star \frac{k}{2}$
8 set $\tilde{\boldsymbol{\sigma}}$ to 1 for the first k (sorted) coordinates
9 set $\tilde{\boldsymbol{\sigma}}$ to 0 for the remaining $n - k$ (sorted) coordinates

On an intuitive level, the MN-Algorithm works as follows. First, we query m times exactly Γ randomly chosen entries of the signal in parallel, which yields the graph representation $\boldsymbol{G}$. Secondly, we sum up the query results $(\boldsymbol{y}_a)_{a \in \partial x}$ in the neighborhood induced by $\boldsymbol{G}$ of each coordinate, counting multi-edges only once. The sum is then centralized by its expected value. Finally, those coordinates with a large *score* are very likely to have the value 1 under $\boldsymbol{\sigma}$. Our first main theorem states how many parallel queries are required for the MN-Algorithm to recover the correct $\boldsymbol{\sigma}$ w.h.p.

Theorem 1. *Suppose that $0 < \theta < 1$, $k = n^\theta$, and $\varepsilon > 0$ and let*

$$m_{\text{MN}}(n, \theta) = 4\left(1 - \frac{1}{\sqrt{e}}\right)\frac{1 + \sqrt{\theta}}{1 - \sqrt{\theta}} k \ln(n/k).$$

If $m > (1 + \varepsilon) m_{\text{MN}}(n, \theta)$, then Algorithm 1 outputs $\boldsymbol{\sigma}$ w.h.p. on input m and k and an additive querying method `query` that returns the total number of one-entries in a query.

[1]The expression *with high probability (w.h.p.)* refers to a probability that tends to 1 as $n \to \infty$.

While the MN-algorithm takes k as an input, the proof reveals that prior knowledge of k is not required in detail. More precisely, a lower bound on k suffices, as in this case enough queries are conducted and the design of $\boldsymbol{G}$ is independent from k. Observe that one additional parallel query on all entries reveals the exact value of k immediately without increasing m asymptotically and therefore the only dependence on k in Algorithm 1 (Line 7) can be easily removed by this one additional query. Beside not being strictly dependent on k, a main novelty of the MN-algorithm is its greedy fashion, providing a straightforward approach compared to the technically challenging algorithms presented in [18], [19].

Parallelized Reconstruction. Observe that our reconstruction algorithm, apart from sampling the test design and performing all queries in parallel, is specified in a sequential fashion. This emphasizes the local structure of the reconstruction algorithm. In the context of a parallel computation we observe that our algorithm can be readily parallelized. When individual queries can be conducted much faster, this further reduces the overall running time of our approach. Such improved reconstruction algorithms can be used in the context of machine learning, see, e.g., [33] for an application.

Recall that our test design is described by a random bipartite graph $\boldsymbol{G}$ and let $M = M(G) = (m_{ij}) \in \{0,1\}^{n \times m}$ be the unweighted biadjacency matrix of G. Intuitively, the entries of M are those values that are summed up in Line 6 of Algorithm 1. It follows that the Ψ_i and $\boldsymbol{\Delta}_i^\star$ vectors are matrix-vector products $\boldsymbol{\Delta}^* = M1$ and $\Psi = M\boldsymbol{y}$ where $1 = (1, \dots, 1)$ is the all-one-vector and $\boldsymbol{y}$ is the query result vector. The sums computed in Lines 4 to 6 of Algorithm 1 can therefore be expressed in terms of two matrix-vector products for which efficient parallelizations are known. Finally, in Lines 7 to 9 of Algorithm 1 the (coordinates of) the resulting vector are sorted. See [28] for a rather recent survey (with a focus on but not limited to GPUs) on parallel sorting algorithms.

Information-Theoretic Results. We prove that in the sublinear regime where $k = n^\theta$ for some $\theta \in (0,1)$ it is possible to reconstruct $\boldsymbol{\sigma}$ from $(\boldsymbol{G}, \boldsymbol{y})$ with high probability with no more than $(1 + \varepsilon) m_{\text{para}}^{\text{BPD}}$ parallel queries for some arbitrarily small $\varepsilon > 0$. More precisely, we show that there is, with high probability, no second input vector $\tau \in \{0,1\}^n$ leading to the same sequence of query results.

Theorem 2. *Suppose that $0 < \theta < 1$, $k = n^\theta$, and $\varepsilon > 0$ and let*

$$m_{\text{para}}^{\text{BPD}} = 2\frac{k \ln(n/k)}{\ln k} = 2\frac{1 - \theta}{\theta} k.$$

If $m > (1 + \varepsilon) m_{\text{para}}^{\text{BPD}}$, $\boldsymbol{\sigma}$ can be computed from $\boldsymbol{G}$ and $\boldsymbol{y}$ w.h.p.

Our result reduces the previously known upper bound of Grebinski and Kucherov [17] by a factor of two and we provide the missing counter part of (2) which establishes the existence of a phase-transition at $m_{\text{para}}^{\text{BPD}}$ for parallel designs.

D. Discussion

Our results extend information-theoretic results of Alaoui et al. [1] from the linear regime to the sublinear regime. For $\theta \to 1$, our threshold of Theorem 1 turns out to converge towards the threshold of [1]. The study of the sublinear regime is inspired by studies of the compressed sensing problem with a sparse underlying signal [3]. In the special case of the binary pooled data problem, those studies were initiated by [19]. The sparse regime is indeed interesting in real-world applications, with examples including epidemiology where Heaps law models the early spread of pandemics [5], [31] or the detection of rare features in image classification in machine learning [20]. The relevance of the sublinear regime can be seen in the following example. Suppose a screening for HIV is conducted. Out of about 67,220,000 residents of the UK, 105,200 are known to be infected with the HI virus. Hence, by screening n = 10.000 random probes, we expect 16 positive entries in the signal corresponding to the infection status. Thus, the choice $\theta = 0.3$ describes the situation quite well.

It is not surprising that also similar problems have been recently analyzed in the sublinear regime. By now, a vast body of related literature exists (see, e.g., the survey by Aldridge et al. [2]). Interestingly, for the (presumably more difficult) variant in which a query only returns the information whether at least one non-zero entry was found, a very sophisticated efficient algorithm is known for $\theta \leq \ln 2/(1+\ln 2) \approx 0.409$ which requires $m_{GT} \sim \ln^{-1}(2) k \ln \frac{n}{k}$ parallel queries [9]. Thus, dropping most of the available information and using this approach outperforms not only the simple greedy approach discussed in this paper for small values of θ, but also the quite involved algorithms by Karimi et al. [18], [19]. This result is of fundamental theoretical interest, since it solves an open complexity theoretical question. Nevertheless, their proposed algorithm appears to be of rather limited interest for practical applications, as it requires, e.g., that $\sqrt{\ln \ln n}$ is large. This is in contrast to our simple greedy scheme, which our simulations have shown to work well for real-world input sizes.

As in state-of-the art designs for similar reconstruction problems [2], [9], we allow a specific entry to be included multiple times in one query. While this seems counter-intuitive in the first place, it does not affect practicability of the proposed design.

II. Model and Notation

In this section we formally introduce the pooling design. As before, $\boldsymbol{\sigma} \in \{0,1\}^n$ is the ground truth chosen uniformly at random from all $0-1$ vectors of length n with exactly k non-zero entries, where $k = n^\theta$ for some $\theta \in (0,1)$. We use $\boldsymbol{G} = \boldsymbol{G}(n, m, \Delta)$ to denote the random bipartite multi-graph that models the pooling design, where m denotes the total number of queries and $\boldsymbol{\Delta} = \{\boldsymbol{\Delta}_1, \dots, \boldsymbol{\Delta}_n\}$ describes the number of queries each individual participates in. Observe that $\boldsymbol{\Delta}_i \sim \mathrm{Bin}(mn/2, 1/n)$. Similarly, we let $\boldsymbol{\Delta}^\star = \{\boldsymbol{\Delta}_1^\star, \dots, \boldsymbol{\Delta}_n^\star\}$ denote the number of *distinct* queries with expected value $\mathbb{E}[\boldsymbol{\Delta}_i^\star] = (1-\exp(-1/2))m$. We let the vector $\boldsymbol{y} \in \{0, \dots, \Gamma\}^m$ denote the sequence of query results. When we refer to any other input vector than $\boldsymbol{\sigma}$, we simply write σ for the input vector and $y = y(\boldsymbol{G}, \sigma)$ for the corresponding results' vector. Additionally, we write $V = \{x_1, \dots, x_n\}$ for the set of the n entries of $\boldsymbol{\sigma}$ and let $V_0 = \{x_i \in V : \boldsymbol{\sigma}(i) = 0\}$ and $V_1 = V \setminus V_0$ be the set of entries with value 0 and 1, respectively. For $x_i \in V$, we write ∂x_i for the multiset of queries a_j in which x_i is contained. Similarly, we write $\partial^\star x_i$ for the set of *distinct* such queries. Analogously, for a query a_i, we denote by ∂a_i the multiset of entries that are contained.

Recall that in our model every query contains exactly $\Gamma = n/2$ entries, and those entries are assigned uniformly at random with replacement. If a one-entry x_i participates in a query a_j more than once, it increases $\boldsymbol{y}_j$ multiple times. For each $x_i \in V$, we let Ψ_i be the sum of its query results for *distinct* queries it belongs to. That is, even if the entry appears more than once in a query and thus contributes to the result multiple times, this query's result contributes to Ψ_i only once. Of course, the value of x_i under $\boldsymbol{\sigma}$ has a significant impact on this sum, increasing it by $\boldsymbol{\Delta}_i$, if x_i is non-zero. To account for this effect in our analysis, we introduce a second variable Φ_i that sums all the query results in which x_i is contained and excludes the impact of x_i. Formally, for any configuration $\sigma \in \{0,1\}^n$ we define

$$\Psi_i(\sigma) = \sum_{j \in \partial^\star x_i} y_{a_j} \quad \text{and} \quad \Phi_i(\sigma) = \Psi_i(\sigma) - \mathbf{1}\{\sigma(i) = 1\}\Delta_i$$

and let $\Psi = (\Psi_1, \dots, \Psi_n)$ and $\Phi = (\Phi_1, \dots, \Phi_n)$. When we consider a specific instance $(\boldsymbol{G}, \boldsymbol{y})$, we will write $\boldsymbol{\Psi}_i = \Psi_i(\boldsymbol{\sigma})$ and $\boldsymbol{\Phi}_i = \Phi_i(\boldsymbol{\sigma})$ for the sake of brevity. Notably, while $\boldsymbol{\Psi}_i$ is known to the observer or an algorithm instantly from the queries, $\boldsymbol{\Phi}_i$ is not, since the ground truth $\boldsymbol{\sigma}$ itself is unknown.

To express the number m of queries conducted, we let $c(n) > 0$ denote a positive function from $\mathbb{N}$ to $\mathbb{R}^+$ such that

$$m = c(n) k \frac{\ln(n/k)}{\ln k}.$$

While it turns out that $c(n) = \Theta(1)$ suffices in the analysis of the information-theoretic bound, we will see that the performance guarantee of the MN-algorithm requires $c(n)$ to scale as $\Theta(\ln n)$. Finally, we define a high probability event $\mathcal{R}$ that we will condition on as explained in the teacher-student model. Let $\mathcal{R}$ be the event that, for all $i \in [n]$, we have

$$\begin{aligned} \boldsymbol{\Delta}_i &= \frac{m}{2} + O\left(\sqrt{m}\ln n\right) \\ \text{and} \quad \boldsymbol{\Delta}_i^\star &= (1-\exp(-1/2))\, m + O\left(\sqrt{m}\ln n\right), \end{aligned} \tag{3}$$

meaning that the underlying random graph satisfies concentration properties. The following lemma states that $\mathcal{R}$ is indeed a high probability event.

Lemma 3. *If $\boldsymbol{G}$ is constructed according to our pooling scheme, then $\mathbb{P}(\mathcal{R}) = 1 - o(1)$.*

The proof follows from standard concentration results, see the appendix for the technical details. Since Theorems 1 and 2 only contain w.h.p.-assertions, we can safely condition on $\mathcal{R}$ for the remainder of our analysis.

III. MN-Algorithm

Outline. Recall that Ψ_i is the sum over all query results in which the entry x_i is contained (multi-edges counted only once) and $\mathbf{\Delta}_i^\star$ is the (random) number of disjoint such queries. Furthermore, let $\mathcal{E}_j$ be the $\sigma-$algebra generated by the edges connected with x_j. As already discussed, we get

$$\mathbf{\Delta}_i^\star = (1+o(1))\,(1-\exp(-1/2))\,m$$

w.h.p. Therefore, intuitively spoken, a non-zero entry x_i increases the value of $\mathbf{\Psi}_i$ by $\mathbf{\Delta}_i = (1+o(1))m/2$, other than zero-entries. Moreover, by construction of the random bipartite (multi-)graph $\boldsymbol{G}$, we get that the second neighborhood of x_i contains $\mathrm{Bin}\,(\Gamma\mathbf{\Delta}_i^\star, k/n)$ non-zero entries. Thus we expect

$$\mathbb{E}\left[\Psi_i - \mathbf{\Delta}_i^\star\frac{k}{2}\middle|\mathcal{E}_i\right] = \mathbf{1}\{\boldsymbol{\sigma}(i)=1\}\mathbf{\Delta}_i.$$

Therefore, if $\Psi_i - \mathbf{\Delta}_i^\star\frac{k}{2}$ is called the *score* of entry x_i, we observe that the scores differ between zero entries and non-zero entries. The whole proof of the algorithmic performance boils down to identify a threshold value $T(\alpha) = T(n,k,\alpha)$ such that, if sufficiently many queries are conducted, all scores of zero entries are below $T(\alpha)$ while the scores of all non-zero entries exceed this threshold w.h.p. If we conduct $m = dk\ln\frac{n}{k}$ queries, with $d = c(n)\ln(k)^{-1}$, we get by a standard application of a Chernoff bound and a union bound over all $k = n^\theta$ non-zero entries $x_i \in V_1$ and, respectively, $n-k = \Theta(n)$ zero-entries $x_i \in V_0$ that $T(\alpha)$ is a valid threshold whenever

$$\begin{aligned}&\frac{-(1-\theta)\alpha^2 d}{4\,(1-\exp(-1/2))\,(1+o(1))} + \theta < 0\\ \text{and}\quad &\frac{-(1-\theta)(1-\alpha)^2 d}{4\,(1-\exp(-1/2))\,(1+o(1))} + 1 < 0,\end{aligned} \tag{4}$$

which will become clear in a second. Optimizing (4) with respect to $\alpha \in (0,1)$ and plugging d into $m = dk\ln(n/k)$ yields for any $\varepsilon > 0$ the sufficient condition

$$m \geq (4+\varepsilon)(1+o(1))\,(1-\exp(-1/2))\,\frac{1+\sqrt{\theta}}{1-\sqrt{\theta}}k\ln(n/k).$$

Formal Analysis. Let $\boldsymbol{A}_{ij} \in \mathbb{N}_0$ denote how often entry x_i appears in query a_j and let $\boldsymbol{A} = (\boldsymbol{A}_{ij})_{i\in[n],j\in[m]}$ be the adjacency matrix of $\boldsymbol{G}$. Then the following holds.

Corollary 4. *Let $1 \leq j \leq n$. Given $\mathcal{E}_j$, the random variable*

$$\boldsymbol{S}_j = \boldsymbol{\psi}_j - \mathbf{\Delta}_j = \sum_{i=1}^{m}\mathbf{1}\{\boldsymbol{A}_{ij}>0\}\left(\boldsymbol{y}_j - \boldsymbol{A}_{ij}\right)$$

has distribution $\mathrm{Bin}\left(\mathbf{\Delta}_j^\star\Gamma - \mathbf{\Delta}_j, \dfrac{k-\mathbf{1}\{\boldsymbol{\sigma}(j)=1\}}{n-1}\right)$.

Proof. This is an immediate consequence of the model definition. There are $\Gamma\mathbf{\Delta}_j^\star - \mathbf{\Delta}_j$ half-edges connected to query-nodes in the neighborhood of x_j that are connected to entry-nodes $x_i \neq x_j$. Each of these half-edges is connected to one of $k - \mathbf{1}\{\boldsymbol{\sigma}(j)=1\}$ entry-nodes belonging to an entry of value 1, independently, from the $n-1$ remaining entry-nodes. □

Now it is possible to immediately infer the expectation of $\boldsymbol{S}_j$ conditioned on the event $\mathcal{R}$ (as defined in (3)). For the sake of brevity let $\gamma = 1-\exp(-1/2)$. Given the event $\mathcal{R}$ which guarantees concentration properties of the underlying graph, we get w.h.p.

$$\mathbb{E}\left[\boldsymbol{S}_j \mid \mathcal{E}_j, \mathcal{R}\right] = (1\pm\delta)\,\frac{\gamma km}{2} \tag{5}$$

where $$\delta := \frac{\sqrt{2}\ln n}{\sqrt{\gamma mk}} = o(1).$$

The Chernoff bound allows us to bound $\boldsymbol{S}_j$ as follows.

Lemma 5. *Let $\alpha \in (0,1)$ be a constant and $m = dk\ln\frac{n}{k}$. Then*

$$\begin{aligned}&\mathbb{P}\left(|\boldsymbol{S}_j - \mathbb{E}\left[\boldsymbol{S}_j \mid \mathcal{E}_j,\mathcal{R}\right]| \geq (1-\alpha)m/2 \mid \mathcal{E}_j,\mathcal{R}\right)\\ &\qquad\leq \exp\left(-(1+o(1))\,\frac{(1-\alpha)^2 d}{4\gamma(1+o(1))}\ln\frac{n}{k}\right).\end{aligned}$$

Proof. The Chernoff bound (Lemma 12) directly implies

$$\begin{aligned}&\mathbb{P}\left(|\boldsymbol{S}_j - \mathbb{E}\left[\boldsymbol{S}_j \mid \mathcal{E}_j\right] \mid \mathcal{E}_j,\mathcal{R}| \geq (1-\alpha)m/2 \mid \mathcal{R}\right)\\ &\qquad\leq \exp\left(-(1+o(1))\,\frac{(1-\alpha)^2 m}{8\mathbb{E}\left[\boldsymbol{S}_j \mid \mathcal{E}_j,\mathcal{R}\right]}\right)\\ &\qquad= \exp\left(-(1+o(1))\,\frac{(1-\alpha)^2 d}{4\gamma(1+o(1))}\ln(n/k)\right).\end{aligned}$$ □

Next we show that, with a suitable choice of a threshold, the scores of zero- and one-entries are well separated.

Corollary 6. *Let $\varepsilon > 0$ be an arbitrary constant. If $m \geq (4+\varepsilon)\,(1-\exp(-1/2))\,\frac{1+\sqrt{\theta}}{1-\sqrt{\theta}}k\ln\frac{n}{k}$ then there exists an $\alpha \in (0,1)$ such that, w.h.p., we have*

$$\boldsymbol{S}_j + \mathbf{\Delta}_j \geq \mathbb{E}\left[\boldsymbol{S}_j \mid \mathcal{E}_j,\mathcal{R}\right] + (1-\alpha)m/2$$

for all x_j where $\boldsymbol{\sigma}(j)=1$, and

$$\boldsymbol{S}_j < \mathbb{E}\left[\boldsymbol{S}_j \mid \mathcal{E}_j,\mathcal{R}\right] + (1-\alpha)m/2$$

for all x_j where $\boldsymbol{\sigma}(j)=0$.

Proof. Let $x_j \in V_1(\boldsymbol{G})$. Again, we make use of the concentration properties guaranteed by conditioning on $\mathcal{R}$. Therefore, we assume that $\mathbf{\Delta}_j = m/2 + O(\sqrt{m}\ln n)$. Then Lemma 5 ensures that

$$\begin{aligned}&\mathbb{P}\left(\boldsymbol{S}_j + \mathbf{\Delta}_j \leq \mathbb{E}\left[\boldsymbol{S}_j \mid \mathcal{E}_j,\mathcal{R}\right] + (1-\alpha)m/2 \mid \mathcal{E}_j,\mathcal{R}\right)\\ &\qquad\leq \exp\left(-\alpha^2 d/(4\gamma(1+o(1)))\ln\frac{n}{k}\right)\\ &\qquad= \exp\left(\frac{(\theta-1)\alpha^2 d}{4\gamma(1+o(1))}\ln n\right).\end{aligned}$$

Hence, the union bound shows that the first inequality holds for all k elements of $V_1(\boldsymbol{G})$ w.h.p. if

$$\frac{(\theta-1)\alpha^2 d}{4\gamma(1+o(1))} + \theta < 0. \tag{6}$$

Analogously, the second inequality holds for all $n-k$ elements of $V_0(\boldsymbol{G})$ w.h.p. if

$$\begin{aligned}\mathbb{P}\left[\boldsymbol{S}_j \geq \mathbb{E}\left[\boldsymbol{S}_j \mid \mathcal{E}_j, \mathcal{R}\right] + (1-\alpha)m/2 \mid \mathcal{E}_j, \mathcal{R}\right] \\ \leq \exp\left(\left((1-\alpha)^2 d/(4\gamma(1+o(1)))\right)\ln\frac{n}{k}\right) \\ = \exp\left(\frac{(\theta-1)(1-\alpha)^2 d}{4\gamma(1+o(1))}\ln n\right).\end{aligned}$$

Again, the union bound shows that the second inequality holds w.h.p. if

$$\frac{(\theta-1)(1-\alpha)^2 d}{4\gamma(1+o(1))} + 1 < 0. \tag{7}$$

Note that the condition in (6) is monotonically decreasing in α while the condition in (7) is monotonically increasing in α. Hence the optimal choice of α is the one that makes the two terms in (6) and (7) equal:

$$\frac{(\theta-1)\alpha^2 d}{4\gamma(1+o(1))} + \theta = \frac{(\theta-1)(1-\alpha)^2 d}{4\gamma(1+o(1))} + 1,$$

which boils down to

$$\alpha = \frac{d - 4\gamma(1+o(1))}{2d}.$$

By putting this solution for α into (6) we get

$$\frac{(\theta-1)(d-4(\gamma+o(1))^2}{16\gamma d + o(1)} + \theta < 0.$$

It now suffices to find the minimal $d = d(\theta) > 0$ such that

$$\frac{(\theta-1)(d-4\gamma+o(1))^2}{16\gamma d + o(1)} + \theta = 0.$$

Hence, we solve for (positive) d and obtain that Eqs. (6) and (7) hold w.h.p. provided

$$d \geq 4\gamma \cdot \frac{1+\sqrt{\theta}}{1-\sqrt{\theta}} + o(1),$$

which matches the assumption in the lemma statement. □

We are now ready to formally prove Theorem 1.

Proof of Theorem 1. According to Lemma 3, the event $\mathcal{R}$ is a high-probability event. Corollary 6 then immediately implies the theorem, together with the definition $m = dk\ln\frac{n}{k}$. □

IV. Information-Theoretic Achievability

In the following section we prove Theorem 2. Our approach is based on counting alternative input vectors $\sigma \neq \boldsymbol{\sigma}$ that yield the same sequence of query results as the ground truth $\boldsymbol{\sigma}$. Note that the underlying techniques are regularly employed for random constraint satisfaction problems [10].

We start with an outline of the proof. Let $S_k(\boldsymbol{G}, \boldsymbol{y})$ be the set of all vectors $\sigma \in \{0,1\}^n$ of Hamming weight k such that

$$\boldsymbol{y}_{a_i} = |\{x_j \in \partial a_i : \sigma(j) = 1\}| \quad \text{for all } i \in [m].$$

This means, we fix m queries $a_1, \dots, a_m$ and let $S_k(\boldsymbol{G}, \boldsymbol{y})$ be the set of all vectors $\sigma \in \{0,1\}^n$ with exactly k ones that are consistent with the query results. Let now $Z_k(\boldsymbol{G}, \boldsymbol{y}) = |S_k(\boldsymbol{G}, \boldsymbol{y})|$. We need to prove that $Z_k(\boldsymbol{G}, \boldsymbol{y}) = 1$ w.h.p. if the number of queries m exceeds $m^{\text{BPD}}_{\text{para}}$. Note that we can always reconstruct $\boldsymbol{\sigma}$ exactly in this case via an exhaustive search (recall that from an information-theoretic point of view the computational power is assumed to be unlimited).

In our analysis, it turns out that it is much more convenient to study $Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y})$, the number of alternative vectors that are consistent with the query results and have a so-called *overlap* of ℓ with $\boldsymbol{\sigma}$. The overlap is the number of one-entries under $\boldsymbol{\sigma}$ that are also present in an alternative vector σ. Formally, we define

$$Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) = \left|\{\sigma \in S_k(\boldsymbol{G}, \boldsymbol{y}) : \sigma \neq \boldsymbol{\sigma}, \langle \boldsymbol{\sigma}, \sigma\rangle = \ell\}\right|.$$

It now suffices to prove that $\sum_{\ell=0}^{k-1} Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) = 0$ for $m \geq (1+\varepsilon)m^{\text{BPD}}_{\text{para}}$ w.h.p. To this end, two separate arguments are needed. First, we show in Proposition 7 via a first moment argument that no second satisfying input vector σ can exist with a small overlap with $\boldsymbol{\sigma}$. Secondly, we employ in Proposition 11 the classical coupon collector argument to show that a second satisfying configuration cannot exist for large overlaps. Intuitively, this means that an entry that is flipped from zero under $\boldsymbol{\sigma}$ to one under an alternative configuration σ initiates a cascade of other changes to maintain the observed query results. The full technical proofs for the following statements can be found in the appendix.

Proposition 7. *Let $\varepsilon > 0$, $0 < \theta < 1$ and assume that $m > (1+\varepsilon)m^{\text{BPD}}_{\text{para}}$. W.h.p., we have*

$$\sum_{\ell=0}^{k(1-\exp(-1/2))} Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) = 0.$$

We now sketch the proof of Proposition 7. By Markov's inequality it suffices to show that $\mathbb{E}[Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y})] \to 0$ fast enough for all ℓ with $0 \leq \ell < k - (1-\exp(-1/2))\ln k$ if $m \geq (1+\varepsilon)m^{\text{BPD}}_{\text{para}}$ for some $\varepsilon > 0$. For $\mathbb{E}[Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y})]$ we compute

$$\begin{aligned}\mathbb{E}[Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y})] \leq \binom{k}{\ell}\binom{n-k}{k-\ell}\prod_{i=1}^{m}\sum_{j=1}^{\boldsymbol{y}_{a_i}}\left(\binom{\Gamma}{j, j, \Gamma-2j}\right. \\ \left.\cdot\left((1-\ell/k)\frac{k}{n}\right)^{2j}\cdot\left(1-2(1-\ell/k)\frac{k}{n}\right)^{\Gamma-2j}\right).\end{aligned}$$

The combinatorial meaning is the following: The binomial coefficients count the number of possible input vectors $\sigma \neq \boldsymbol{\sigma}$ of overlap ℓ with $\boldsymbol{\sigma}$. The subsequent term measures the probability that a specific such σ yields the same results on queries $a_1, \dots, a_m$ as $\boldsymbol{\sigma}$. To see this, we divide the entries $x_1, \dots, x_n$ into three categories. The first category contains those entries that exhibit the same value under σ and $\boldsymbol{\sigma}$. The second and third category feature those entries that are set to one under σ and to zero under $\boldsymbol{\sigma}$ and vice versa. Recall that ℓ determines the number of x_i that are set to one under both vectors σ and $\boldsymbol{\sigma}$. The probability for a specific entry to be in the first category is $1 - 2(1-\ell/k)k/n$, while the probability for a specific entry to be in the second or third categories

is $(1-\ell/k)k/n$ each. The key observation is that the query results are the same between σ and $\boldsymbol{\sigma}$ if and only if the number of entries in the second category is identical to the number of entries in the third category. We compute (a bound on) the sum over the number of entries which are flipped. Simplifying the term and conditioning on the high probability event $\mathcal{R}$ yields the following lemma.

Lemma 8. *For every $0 \le \ell \le k - (1-\exp(-1/2))\ln k$ and a random variable $\boldsymbol{X} \sim \mathrm{Bin}_{\ge 1}(\Gamma, 2(1-\ell/k)k/n)$, we have*

$$\begin{aligned}\mathbb{E}[Z_{k,\ell}] &\le (1+O(1))\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y}) \mid \mathcal{R}] \\ &\le (1+O(1))\binom{k}{\ell}\binom{n-k}{k-\ell}\left(\frac{1}{\sqrt{2\pi}}\mathbb{E}\left[\frac{1}{\sqrt{\boldsymbol{X}}}\right]\right)^m.\end{aligned}$$

Here, $\mathrm{Bin}_{\ge i}(n,p)$ is the binomial distribution with parameters n and p where we condition that its outcome is at least i.

Proof. The product of the two binomial coefficients simply accounts for the number of vectors σ that have overlap ℓ with $\boldsymbol{\sigma}$. Let $\mathcal{S}$ denote the event that one specific $\sigma \in \{0,1\}^n$ that has overlap ℓ with $\boldsymbol{\sigma}$ belongs to $S_{k,\ell}(\boldsymbol{G},\boldsymbol{y})$. It suffices to show for $\boldsymbol{X} \sim \mathrm{Bin}_{\ge 1}(\Gamma, 2(1-\ell/k)k/n)$ that

$$\mathbb{P}[\mathcal{S} \mid \mathcal{R}] \le (1+O(1))\left(\frac{1}{\sqrt{2\pi}}\mathbb{E}\left[\frac{1}{\sqrt{\boldsymbol{X}}}\right]\right)^m. \tag{8}$$

The remainder of the proof is dedicated to showing Eq. (8).

By the design $\boldsymbol{G}$, each query contains $\Gamma = n/2$ entries chosen uniformly at random, and we observe that all query results are statistically independent of each other. Therefore, we need only to determine the probability that for a specific σ and a specific query a_i the result is consistent with the result under $\boldsymbol{\sigma}$ such that $\boldsymbol{y}_i = y_i$. Given the overlap ℓ, we know for σ drawn uniformly at random that $\mathbb{P}[\boldsymbol{\sigma}_i = \sigma_i = 1] = \ell/n$, $\mathbb{P}[\boldsymbol{\sigma}_i = \sigma_i = 0] = (n-2k+\ell)/n$ and finally $\mathbb{P}[\boldsymbol{\sigma}_i \ne \sigma_i] = (k-\ell)/n$ holds for all x_i, $i = 1 \dots n$. We get

$$\begin{aligned}\mathbb{P}[\mathcal{S} \mid \mathcal{R}] &\le \prod_{i=1}^{m}\sum_{j=1}^{\boldsymbol{y}_i}\left(\binom{\Gamma}{j,j,\Gamma-2j}\cdot\left((1-\ell/k)\frac{k}{n}\right)^{2j}\right. \\ &\qquad\left.\cdot\left(1-2(1-\ell/k)\frac{k}{n}\right)^{\Gamma-2j}\right) \\ &\le \left(\sum_{j=1}^{\Gamma/2}\binom{\Gamma}{2j}\left(2(1-\ell/k)\frac{k}{n}\right)^{2j}\right. \\ &\qquad\left.\cdot\left(1-2(1-\ell/k)\frac{k}{n}\right)^{\Gamma-2j}\binom{2j}{j}2^{-2j}\right)^m. \end{aligned}\tag{9}$$

The last two components of (9) describe the probability that a one-dimensional simple random walk returns to its original position after $2j$ steps, which is by Lemma 14 equal to $(1+O(j^{-1}))/\sqrt{\pi j}$. The former term describes the probability that a $\mathrm{Bin}_{\ge 1}(\Gamma, 2(1-\ell/k)k/n))$ random variable $\boldsymbol{X}$ takes the value $2j$. For $\ell \le k-(1-\exp(-1/2))\ln k$ the expectation of $\boldsymbol{X}$ given $\boldsymbol{G}$ is at least of order $\ln k$ such that the asymptotic description of the random walk return probability is feasible. Note that if ℓ gets closer to k, the expectation of $\boldsymbol{X}$ gets finite, s.t. the random walk approximation is not feasible anymore. Therefore, using Lemma 15, we can, as long as $\Gamma(2(1-\ell/k)k/n) = \Omega(\ln n)$, simplify (9) to

$$\begin{aligned}\mathbb{P}[\mathcal{S} \mid \mathcal{R}] &\le (1+O(1))\left(\sum_{j=1}^{\Gamma/2}\binom{\Gamma}{2j}\left(2(1-\ell/k)\frac{k}{n}\right)^{2j}\right. \\ &\qquad\left.\cdot\left(1-2(1-\ell/k)\frac{k}{n}\right)^{\Gamma-2j}\frac{1}{\sqrt{\pi j}}\right)^m \\ &= (1+O(1))\left(\frac{1}{2}\sum_{j=1}^{\Gamma}\binom{\Gamma}{j}\left(2(1-\ell/k)\frac{k}{n}\right)^{j}\right. \\ &\qquad\left.\cdot\left(1-2(1-\ell/k)\frac{k}{n}\right)^{\Gamma-j}\frac{1}{\sqrt{\pi j/2}}\right)^m \\ &= (1+O(1))\left(\frac{1}{\sqrt{2\pi}}\mathbb{E}\left[\frac{1}{\sqrt{\boldsymbol{X}}}\right]\right)^m\end{aligned}$$

for large $n \gg 1$ which implies Lemma 8. □

While the expression given through Lemma 8 might look hard to work with, it can be simplified using standard asymptotic arguments as follows.

Lemma 9. *For every $0 \le \ell \le k-(1-\exp(-1/2))\ln k$, $m = ck\frac{\ln(n/k)}{\ln(k)}$ and $n \gg 1$, we have*

$$\begin{aligned}&\frac{1}{n}\ln\left(\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y}) \mid \mathcal{R}]\right) \\ &\qquad\le (1+o(1))\left(\frac{k}{n}H\left(\frac{\ell}{k}\right)+\left(1-\frac{k}{n}\right)H\left(\frac{k-\ell}{n-k}\right)\right. \\ &\qquad\left.-\frac{ck/n\ln(n/k)}{2\ln k}\ln\left(2\pi\left(1-\frac{\ell}{k}\right)k\right)\right).\end{aligned}$$

The key is to choose $c = c(n)$ such that $Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y}) \to 0$ for every $\ell \le k-(1-\exp(-1/2))\ln k$ when $n \to \infty$. Asymptotically, $\ln\left(\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y})]/n\right)$ takes its maximum at $\ell = \Theta\left(k^2/n\right)$. Therefore, the r.h.s. of (9) becomes negative if and only if the number of queries m parametrized by c exceeds $m_{\text{para}}^{\text{BPD}}$. This is formalized in the following lemma and concludes the proof of Proposition 7.

Lemma 10. *For every $0 \le \ell \le k-(1-\exp(-1/2))\ln k$, $0 < \theta < 1$ and $\varepsilon > 0$ it holds if $m \ge (1+\varepsilon)m_{\text{para}}^{\text{BPD}}$ that*

$$\frac{1}{n}\ln\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y}) \mid \mathcal{R}] < 0.$$

Proof of Proposition 7. The proposition is a direct consequence of Lemmas 8 to 10 and Markov's inequality. □

While we could already establish that there are w.h.p. no feasible vectors $\sigma \in \{0,1\}^n$ that have a small overlap with the ground truth $\boldsymbol{\sigma}$, we still need to ensure that there are w.h.p. no feasible vectors that have a large overlap with $\boldsymbol{\sigma}$. Indeed, we exclude such vectors with the next proposition.

Proposition 11. *Let $\varepsilon > 0$ and $0 < \theta \le 1$ and assume that $m > (1+\varepsilon)m_{\text{para}}^{\text{BPD}}$. Given $\mathcal{R}$ we have $Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y}) = 0$ for all $k-(1-\exp(-1/2))\ln k < \ell < k$ w.h.p.*

The proof is fundamentally easy as it follows the classical coupon collector argument. However, it needs some technical attention. If we consider a vector σ of length n different from $\boldsymbol{\sigma}$ with the same Hamming weight k, at least one entry that is set to one under $\boldsymbol{\sigma}$ is labeled zero under σ. Given the event $\mathcal{R}$, this entry is part of at least $\boldsymbol{\Delta}_i^\star > m/4$ different queries whose results all change by at least -1, depending on how often the entry participates. To compensate for these changes, we need to find $x_1 \dots x_\ell$ that are zero under $\boldsymbol{\sigma}$ and one under σ such that their joint neighborhood is a super-set of the changed queries. We show that this only happens with probability $o(1)$ following a classical balls-into-bins argument. We now give the full technical proof.

Proof of Proposition 11. Assume that $\sigma \in \{0,1\}^n$ is a second vector that is consistent with the query results $\boldsymbol{y}$. By definition, there is an index $j \in \{1,\dots,n\}$ for which $\boldsymbol{\sigma}(j) = 1$ but $\sigma(j) = 0$. By Lemma 3 the size of $\partial^\star x_j$ is at least

$$\boldsymbol{\Delta}_i^\star \geq (1-\exp(-1/2))\, m - O\left(\sqrt{m}\ln n\right)$$

and for any query $a_l \in \partial x_j$ we have $|y_l(\sigma) - y_l(\boldsymbol{\sigma})| \geq 1$. To guarantee that $y(\boldsymbol{\sigma}) = y(\sigma)$ it is necessary to identify a set of h entries $\mathcal{X}$ for which $\sigma(i) = 1 - \boldsymbol{\sigma}(i)$ for all $i \in \mathcal{X}$ with the property that $\mathcal{X} \supseteq \partial x_j$.

By construction of $\boldsymbol{G}$, the number of queries in $\partial^\star x_j$ that do not contain any of the entries in $\mathcal{X}$, i.e., $\boldsymbol{H} = |\{a \in \partial^\star x_j : \mathcal{X} \cap \partial a = \emptyset\}|$, can be coupled with the number of empty bins in a balls-into-bins experiment as follows. Given $\boldsymbol{G}$, throw $b = \sum_{i=1}^{h} \deg(x_i)$ balls into $\deg(x_i)$ bins. Observe that

$$\deg(x_i) \geq (1-\exp(-1/2))\, m - O\left(\sqrt{m}\ln n\right)$$

and denote by $\boldsymbol{H'}$ the number of empty bins in this experiment. Since for any x_i the $\deg(x_i)$ edges are not only distributed over the $(1-o(1))\,(1-\exp(-1/2))\,m$ query-nodes in ∂x_j but over all m query-nodes in $\boldsymbol{G}$, we get

$$\mathbb{P}\left[\boldsymbol{H} = 0 \mid \mathcal{R}\right] \leq \mathbb{P}\left[\boldsymbol{H'} = 0 \mid \mathcal{R}\right]. \tag{10}$$

We condition on $\mathcal{R}$ and therefore $b = (1+o(1))hm/2$. Furthermore, set $L = \ln(m)h^{-1}$ and let $\gamma = (1-\exp(-1/2))$. Then the r.h.s. of (10) becomes

$$\begin{aligned}\mathbb{P}\left[\boldsymbol{H'} = 0 \mid \mathcal{R}\right] &\leq \left(1 - \left(1 - \frac{1}{\gamma m}\right)^{hm/2}\right)^{\gamma m} \\ &= (1+o(1))\exp\left(-\gamma m^{1-L/(2\gamma)}\right).\end{aligned}$$

Therefore, if $L < 2\gamma$, or equivalently,

$$h < 2\gamma \ln(m) \sim 2\gamma\left(\ln k + \ln\ln k\right),$$

we have $\mathbb{P}\left[\boldsymbol{H'} = 0 \mid \mathcal{R}\right] \leq n^{-\omega(1)}$.

Thus, a Hamming distance of at least one between $\boldsymbol{\sigma}$ and σ immediately implies that the Hamming distance is at least $2\gamma\left(\ln k + \ln\ln k\right)$ with probability $1 - n^{-\omega(1)}$. A union bound over all k one-entries implies the proposition. □

Proof of Theorem 2. The theorem follows directly from Propositions 7 and 11. □

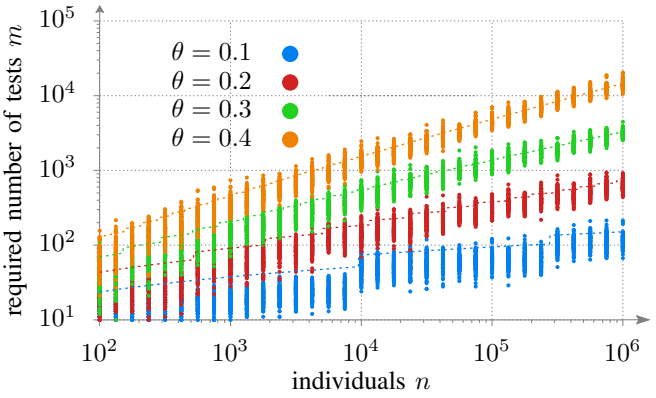

Fig. 2. The required number of queries until $\boldsymbol{\sigma}$ can be exactly reconstructed for different vector lengths n and θ regimes. For each value of n, 100 simulations were carried out independently.

V. Empirical Analysis and Simulation Results

In this section we present simulation results for the MN-Algorithm (Algorithm 1). Our simulation software is implemented in the C++ programming language. It performs a faithful simulation of the parallel system. To generate the random structures, we resort to the Mersenne Twister `mt19937_64` as provided by the C++11 `<random>` library. All of our simulations have been carried out on machines equipped with 20 Intel(R) Xeon(R) E5-2630 v4 CPU cores, backed by 128GiB memory, and running the linux 5.11 kernel. All required code to reproduce our figures, including the gnuplot scripts and various helper tools, can be obtained from our public github repository.

In our first empirical result in Fig. 2 we analyze the number of queries required to reconstruct $\boldsymbol{\sigma}$ for $n \in [10^2, 10^6]$ and different values of θ. The dotted lines show our theoretical asymptotic bounds. Note that the discontinuities in the theoretical bound stem from rounding the number of one-entries k to the closest integer. We remark that our simulation results align well with the theoretical predictions for larger values of n. For smaller values of n, our theoretical results are too optimistic: the lower-order term hidden in the $o(1)$ in Eq. (4) scales as $\Theta\left(\frac{\sqrt{\ln n}}{k}\right)$, and while this expression decreases polynomially fast in n, it is far from vanishing for small values of n and θ.

In Figs. 3 and 4 we analyze the success probability for exact reconstruction of $\boldsymbol{\sigma}$ and the number of correctly identified one-entries. For different numbers of queries we conducted 100 independent simulation runs for $n = 10^3$ and $n = 10^4$ and different values of θ. The dashed lines show the phase-transitions predicted by Theorem 1. The data in Fig. 4 indicate that all but a small fraction of one-entries are correctly detected, even if the exact reconstruction of $\boldsymbol{\sigma}$ is still quite unlikely according to Fig. 3. Overall, the implementation hints at the practical usability of the MN-Algorithm, even for small values of n.

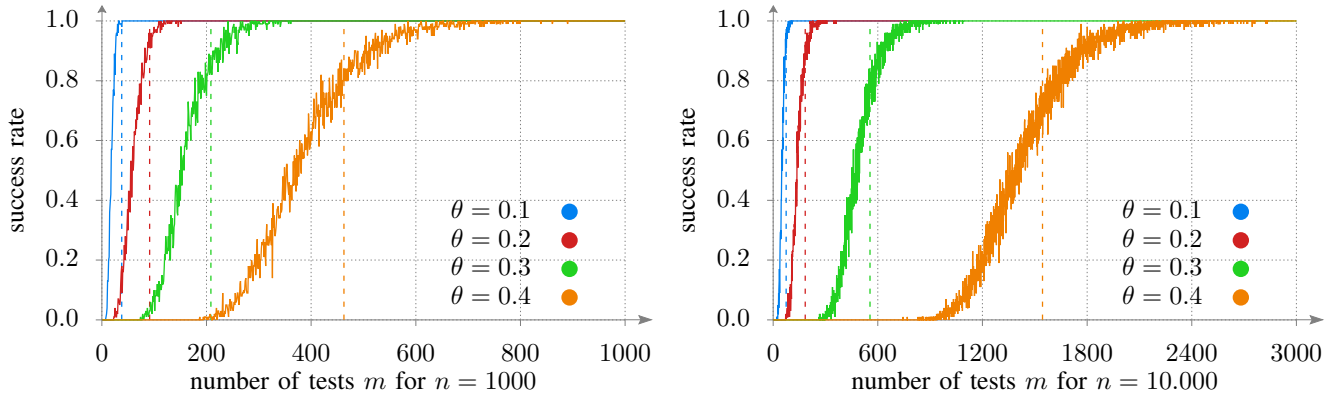

Fig. 3. The plot shows the rate of successful recovery of $\boldsymbol{\sigma}$ among 100 independent simulation runs over the number of queries m for different values of θ and $n = 10^3$ (left) and $n = 10^4$ (right).

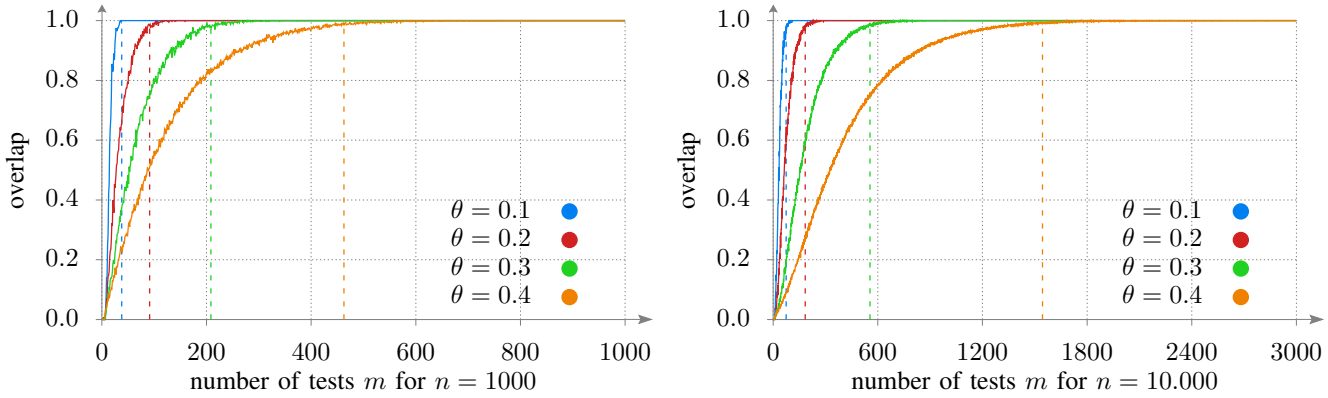

Fig. 4. The plots show the *overlap* – the fraction of correctly classified one-entries – among 100 independent simulation runs over the numbers of queries m for different values of θ and $n = 10^3$ (left) and $n = 10^4$ (right).

Remark. The formal proof of the algorithmic bound directly gives an insight about the convergence speed and thus about the expected performance of the MN-Algorithm for finite n: we can compute that the MN-Algorithm requires an additional multiplicative factor of at least

$$\left(1 + \sqrt{2}\ln n \left(4\left(1 - \exp\left(-1/2\right)\right) mk\right)^{-1/2}\right)$$

queries in addition to the asymptotic analysis for $n \to \infty$. This explains the (slight) deviation of the theoretical and the empirical results for small values of n. See the proof of Corollary 6 in Section III for the rigorous analysis.

VI. Conclusions and Open Problems

In this paper we analyze the binary pooled data problem with additive queries both from an information-theoretic and an algorithmic point of view. Our first result is a simple greedy reconstruction scheme that performs well even close to the information-theoretic boundaries. Our main concern is the design of a reconstruction scheme that works well when all queries are conducted in parallel. In a series of simulations we show that this scheme is applicable to a large range of parameters that can be expected from real-world instances. For example, our data indicate that on average we correctly identify 99% of the one-entries when conducting only 220 queries for $n = 1000$ and $\theta = 0.3$. Our second result sheds light on the information-theoretic achievability threshold, where our theorem closes the open gap between the results of [11] and [17] by establishing a sharp phase transition.

An immediate open problem is to close the gap between the algorithmic and the information theoretic threshold. Furthermore, there are similar reconstruction problems in which parallel conductance of all queries is crucial. As discussed in the introduction, group testing is such a prime example which was recently fully understood using similar techniques as in the present work. A less well understood reconstruction problem is *threshold group testing* [8], [22], in which a query outputs 1 if and only if the number of positive entries exceeds a threshold $T > 0$. It is very likely that the techniques of the present contribution can be applied to threshold group testing as well, as they were previously applied to various reconstruction problems, but the tailor-made application remains a highly non-trivial challenge. Another exciting avenue for future research are partially parallelizable designs. Suppose that, for instance, L processing units can be used to evaluate queries in parallel. Then it is a natural requirement for a design

to always conduct up to L queries in parallel. An interesting open question then is to analyze the trade-offs that arise in such partially parallelized schemes. In particular, there might be designs providing efficient reconstruction algorithms that outperform the completely parallel design studied in this paper.

ACKNOWLEDGEMENTS

The authors thank Uriel Feige for various detailed comments which improved the quality of the paper significantly. Furthermore, the authors thank Petra Berenbrink and Amin Coja-Oghlan for helpful discussions and important hints.

REFERENCES

[1] A. Alaoui, A. Ramdas, F. Krzakala, L. Zdeborová, and M. I. Jordan, "Decoding from pooled data: Phase transitions of message passing," *IEEE Trans. Information Theory*, vol. 65, no. 1, pp. 572–585, 2019.

[2] M. Aldridge, O. Johnson, and J. Scarlett, "Group testing: An information theory perspective," *Foundations and Trends in Communications and Information Theory*, vol. 15, no. 3–4, pp. 196–392, 2019.

[3] Y. Arjoune, N. Kaabouch, H. E. Ghazi, and A. Tamtaoui, "Compressive sensing: Performance comparison of sparse recovery algorithms," *Proc. 7th IEEE CCWC*, 2017.

[4] R. Ben-Ami, A. Klochendler *et al.*, "Large-scale implementation of pooled rna extraction and rt-pcr for sars-cov-2 detection," *Clinical Microbiology and Infection*, vol. 26, no. 9, pp. 1248–1253, 2020.

[5] R. W. Benz, S. J. Swamidass, and P. Baldi, "Discovery of power-laws in chemical space," *Journal of Chemical Information and Modeling*, vol. 48, no. 6, pp. 1138–1151, 2008.

[6] N. H. Bshouty, "Optimal algorithms for the coin weighing problem with a spring scale," *Proc. 22nd COLT*, 2009.

[7] C. C. Cao, C. Li, and X. Sun, "Quantitative group testing-based overlapping pool sequencing to identify rare variant carriers," *BMC Bioinformatics*, vol. 15, p. 195, 2014.

[8] C. L. Chan, S. Cai, M. Bakshi, S. Jaggi, and V. Saligrama, "Stochastic threshold group testing," *2013 IEEE Information Theory Workshop (ITW)*, 2013.

[9] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick, "Optimal group testing," *Combinatorics, Probability and Computing*, p. 1–38, 2021.

[10] A. Coja-Oghlan, F. Krzakala, W. Perkins, and L. Zdeborová, "Information-theoretic thresholds from the cavity method," *Advances in Mathematics*, vol. 333, pp. 694–795, 2018.

[11] A. G. Djackov, "On a search model of false coins," in *Topics in Information Theory. Hungarian Acad. Sci*, vol. 16, 1975, pp. 163–170.

[12] D. Donoho and J. Tanner, "Thresholds for the recovery of sparse solutions via l1 minimization," in *Proc. 40th CISS*, 2006, pp. 202–206.

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

[14] U. Feige and A. Lellouche, "Quantitative group testing and the rank of random matrices," 2020.

[15] S. Foucart and H. Rauhut, *An Invitation to Compressive Sensing*. New York, NY: Springer New York, 2013, pp. 1–39.

[16] X. Gao, M. Sitharam, and A. E. Roitberg, "Bounds on the jensen gap, and implications for mean-concentrated distributions," *The Australian Journal of Mathematical Analysis and Applications*, vol. 16, no. 16, pp. 1–16, 2019.

[17] V. Grebinski and G. Kucherov, "Optimal reconstruction of graphs under the additive model," *Algorithmica*, vol. 28, no. 1, pp. 104–124, 2000.

[18] E. Karimi, F. Kazemi, A. Heidarzadeh, K. R. Narayanan, and A. Sprintson, "Non-adaptive quantitative group testing using irregular sparse graph codes," *Proc. 2019 IEEE Allerton*, pp. 608–614, 2019.

[19] E. Karimi, F. Kazemi, A. Heidarzadeh, K. R. Narayanan, and A. Sprintson, "Sparse graph codes for non-adaptive quantitative group testing," in *IEEE Information Theory Workshop (ITW)*, 2019, pp. 1–5.

[20] W. Liang and J. Zou, "Neural group testing to accelerate deep learning," in *Proc. ISIT*, 2021, pp. 958–963.

[21] G. D. Marco and D. Kowalski, "Searching for a subset of counterfeit coins: Randomization vs determinism and adaptiveness vs non-adaptiveness," *Random Struct. Algorithms*, vol. 42, pp. 97–109, 2013.

[22] G. D. Marco, T. Jurdziński, D. R. Kowalski, M. Różański, and G. Stachowiak, "Subquadratic non-adaptive threshold group testing," *Journal of Computer and System Sciences*, vol. 111, pp. 42–56, 2020.

[23] J. P. Martins, R. Santos, and R. Sousa, "Testing the maximum by the mean in quantitative group tests," in *New Advances in Statistical Modeling and Applications*. Springer, 2014, pp. 55–63.

[24] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition," in *Proc. 27th ACSSC*, 1993, pp. 40–44 vol.1.

[25] J. Scarlett and V. Cevher, "Phase transitions in the pooled data problem," in *Proc. 30th NEURIPS*, 2017, pp. 376–384.

[26] P. Sham, J. S. Bader, I. Craig, M. O'Donovan, and M. Owen, "Dna pooling: a tool for large-scale association studies," *Nature Reviews Genetics*, vol. 3, pp. 862–871, 2002.

[27] H. S. Shapiro, "Problem e 1399," *Amer. Math. Monthly*, vol. 67, p. 82, 1960.

[28] D. P. Singh, I. Joshi, and J. Choudhary, "Survey of GPU based sorting algorithms," *Int. J. Parallel Program.*, vol. 46, no. 6, pp. 1017–1034, 2018.

[29] S. Sparrer and R. F. H. Fischer, "Soft-feedback omp for the recovery of discrete-valued sparse signals," in *Proc. 23rd EUSIPCO*, 2015, pp. 1461–1465.

[30] J. Spencer, *Asymptopia*. American Mathematical Society, 2014.

[31] L. Wang, X. Li, Y. Zhang, and K. Zhang, "Evolution of scaling emergence in large-scale spatial epidemic spreading," *Public Library of Science ONE*, vol. 6, 2011.

[32] Y. Wang and W. Yin, "Sparse signal reconstruction via iterative support detection," *SIAM Journal on Imaging Sciences*, vol. 3, no. 3, pp. 462–491, Jan. 2010.

[33] Y. Zhou, U. Porwal, C. Zhang, H. Q. Ngo, X. Nguyen, C. Ré, and V. Govindaraju, "Parallel feature selection inspired by group testing," in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, Eds., vol. 27. Curran Associates, Inc., 2014.

APPENDIX

We start this appendix with standard concentration bounds that we use throughout this paper.

Lemma 12 ([30])**.** *Let $X \sim \mathrm{Bin}(n,p)$ and $\delta \in (0,1)$.*

$$\text{Then} \qquad \mathbb{P}[X > (1+\delta)np] \le \exp\left(-np\delta^2/(2+\delta)\right)$$
$$\text{and} \qquad \mathbb{P}[X < (1-\delta)np] \le \exp\left(-np\delta^2/2\right).$$

For binomial random variables, the Jensen gap provides good approximations.

Lemma 13 (follows from Eq. (1.1) of [16])**.** *Let $\mathrm{Bin}_{\ge i}(n,p)$ be the binomial distribution with parameters n and p where we condition that its outcome is at least i. Let $\boldsymbol{X} \sim \mathrm{Bin}_{x\ge 1}(n,p)$ with $np \to \infty$. Then, for $\ell \in \{1/2, 1\}$, we have*

$$\mathbb{E}\left[\boldsymbol{X}^{-\ell}\right] = \left(1 + o(n^{-1})\right) \mathbb{E}[\boldsymbol{X}]^{-\ell}.$$

The following lemmas are results on random walks.

Lemma 14 ([30], Section 1.5)**.** *The probability that a simple random walk on $\mathbb{Z}$ with $2j$ steps will end at its original position is given by $(\pi j)^{-1/2} + O(j^{-3/2})$.*

Lemma 15. *The following asymptotic equivalence holds for every $0 < p = p(n) < 1$ when $np \to \infty$.*

$$\sum_{j=1}^{n/2} \binom{n}{2j} p^{2j} (1-p)^{n-2j} j^{-1/2}$$
$$= 2^{-1/2} \sum_{j=1}^{n} \binom{n}{j} p^j (1-p)^{n-j} j^{-1/2} + O((np)^{-1})$$

Proof. Let $\boldsymbol{X}\sim\mathrm{Bin}_{\geq 1}(n,p)$ and define $a_j=\mathbb{P}(\boldsymbol{X}=j)/\sqrt{j/2}$ for $j=1\dots n$. Then

$$a_{j+1}/a_j=(p/(1-p))\left(j/(j+1)^3\right)^{1/2}(n-j)$$

is larger than 1 up to $j^\star\in\{\lfloor(n+1)p\rfloor,\lfloor(n+1)p-1\rfloor\}$, depending on n being even or odd, and strictly less than 1 for $j=j^\star+1,...,n$. Furthermore, $a_j=o(1)$ for every j. Define j' as the largest even integer s.t. $j'\leq j^\star$. Then

$$\begin{aligned}\sum_{j=1}^{n/2}a_{2j}&\geq\frac{1}{2}\left(\sum_{j=1}^{j'/2}a_{2j}+a_{2j-1}+\sum_{j=j'/2+1}^{n/2-1}a_{2j}+a_{2j+1}\right)\\&=\left(\frac{1}{2}\sum_{j=1}^{n}a_j\right)+O((np)^{-2}),\end{aligned}$$

The upper bound follows similarly, and together they imply the lemma. □

We now prove the concentration results for the random regular pooling design.

Proof of Lemma 3. Fix an index $i\in[n]$. From the construction of $\boldsymbol{G}$ it follows that $\boldsymbol{\Delta}_i$ is distributed as $\mathrm{Bin}(mn/2,1/n)$. Then Lemma 12 implies

$$\mathbb{P}\left(\boldsymbol{\Delta}_i>m/2+O\left(\sqrt{m}\ln^2 n\right)\right)=n^{-\omega(1)}.$$

Furthermore, the probability that an entry x_i is contained in a specific query a_j is given by

$$p=1-\left(1-n^{-1}\right)^{\Gamma}=\left(1+n^{-\Omega(1)}\right)\left(1-1/\sqrt{e}\right).$$

Since queries select their participating entries independently of each other, we observe that $\boldsymbol{\Delta}_i^\star\sim\mathrm{Bin}\,(m,p)$. Thus, Lemma 12 implies

$$\mathbb{P}\left(\boldsymbol{\Delta}_i^\star>\left(1-1/\sqrt{e}\right)m+O\left(\sqrt{m}\ln n\right)\right)=n^{-\omega(1)}.$$

The union bound over all n entries concludes the proof. □

Proof of Lemma 9. Let $\boldsymbol{X}\sim\mathrm{Bin}_{\geq 1}(\Gamma,2(1-\ell/k)k/n)$. Then

$$\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y})\,|\,\mathcal{R}]\leq O(1)\cdot\binom{k}{\ell}\binom{n-k}{k-\ell}\left(\frac{1}{2\pi\mathbb{E}[\boldsymbol{X}]}\right)^{\frac{m}{2}}\quad(11)$$

by Lemmas 8 and 13. We use the well known fact [30] that as $n\to\infty$ we have for $p\in(0,1)$ that

$$n^{-1}\ln\binom{n}{np}\to H(p):=-p\ln(p)-(1-p)\ln(1-p).$$

We apply the $\ln(\cdot)$ to (11) and divide it by n. Then we calculate using $m=ck\ln(n/k)\ln^{-1}(k)$

$$\begin{aligned}&n^{-1}\ln\left(\mathbb{E}\left[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y})\,|\,\mathcal{R}\right]\right)\\&\quad\leq(1+o(1))\left(\frac{k}{n}H\left(\frac{\ell}{k}\right)+\left(1-\frac{k}{n}\right)H\left(\frac{k-\ell}{n-k}\right)\right.\\&\qquad\left.-\frac{ck/n\ln(n/k)}{2\ln k}\ln\left(2\pi k\left(1-\ell/k\right)\right)\right).\end{aligned}$$

□

Proof of Lemma 10. Let $\gamma=1-\exp(-1/2)$ and recall

$$m=ck\ln\left(\frac{n}{k}\right)/\ln k=c\frac{1-\theta}{\theta}k\text{ and }0\leq\ell\leq k-\gamma\ln k\quad(12)$$

for a constant $c>0$. Then define $f_{n,k}:[0,k-\gamma\ln k]\to\mathbb{R}$ as

$$\begin{aligned}\ell\mapsto\left(\frac{k}{n}H\left(\frac{\ell}{k}\right)+\left(1-\frac{k}{n}\right)H\left(\frac{k-\ell}{n-k}\right)\right.\\\left.-\frac{ck/n\ln(n/k)}{2\ln k}\ln(2\pi(1-\ell/k)k)\right)\quad(13)\end{aligned}$$

and assume, as usual, $0\ln 0=0$. By Lemma 9 we get

$$n^{-1}\ln\left(\mathbb{E}\left[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y})\,|\,\boldsymbol{G}\right]\right)\leq(1+o(1))f_{n,k}(\ell).$$

Expanding the entropy yields

$$\begin{aligned}f_{n,k}(\ell)=\frac{1}{n}\left(-\ell\ln\left(\frac{\ell}{k}\right)-(k-\ell)\ln\left(1-\frac{\ell}{k}\right)\right.\\-(k-\ell)\ln\left(\frac{k-\ell}{n-k}\right)-(n-2k+\ell)\ln\left(1-\frac{k-\ell}{n-k}\right)\\\left.+\frac{ck\ln(k/n)}{2\ln k}\ln\left(2\pi k\left(1-\frac{\ell}{k}\right)\right)\right),\\f'_{n,k}(\ell)=\frac{1}{n}\left(-\ln\left(\frac{\ell}{k}\right)+\ln\left(1-\frac{\ell}{k}\right)+\ln\left(\frac{k-\ell}{n-k}\right)\right.\\\left.-\ln\left(1-\frac{k-\ell}{n-k}\right)-\frac{ck\ln(k/n)}{2(k-\ell)\ln k}\right),\quad\text{and}\\f''_{n,k}(\ell)=\frac{1}{n}\left(-\frac{1}{\ell}-\frac{2}{k-\ell}-\frac{1}{n-2k+\ell}-\frac{ck\ln(k/n)}{2\ln k(k-\ell)^2}\right).\end{aligned}$$

If $\ell=o(k)$ we get $\left|\frac{1}{k-\ell}\left(2-\frac{ck(1-\theta)}{2\theta(k-\ell)}\right)\right|\ll\frac{1}{\ell}$ and therefore

$$nf''_{n,k}(\ell)=-\frac{1}{\ell}-\frac{1}{n-2k+\ell}-\frac{1}{k-\ell}\left(2-\frac{ck(1-\theta)}{2\theta(k-\ell)}\right)<0.$$

This shows that $f'_{n,k}$ is monotonically decreasing in ℓ for large enough n. Furthermore, $f'_{n,k}$ is continuous on $(0,k-\gamma\ln k]$. Let $\tilde{c}>0$ be an arbitrary constant. Then

$$nf'_{n,k}\left(\tilde{c}\frac{k^2}{n}\right)=-\ln(\tilde{c})+\frac{c(1-\theta)}{\theta}+o(1).$$

This implies that there are $0<\tilde{c}_1<\tilde{c}_2<\infty$ s.t.

$$nf'_{n,k}\left(\tilde{c}_1\frac{k^2}{n}\right)>0\quad\text{and}\quad nf'\left(\tilde{c}_2\frac{k^2}{n}\right)<0.$$

By the intermediate value theorem it follows that there is $\hat{c}\in[\tilde{c}_1,\tilde{c}_2]$ s.t. $\hat{c}\frac{k^2}{n}$ is the unique maximizer of $f_{n,k}$ for $\ell=o(k)$. Finally, by putting this value into Eq. (13) we obtain that the highest order terms satisfy

$$nf_{n,k}\left(\hat{c}\frac{k^2}{n}\right)<0\Longleftrightarrow c>-2\frac{H(k/n)}{k/n\ln(k/n)}=2+o(1).\quad(14)$$

Furthermore, if $k-\gamma\ln k\geq\ell=\Theta(k)$, we get

$$nf_{n,k}(\ell)=-\frac{c(1-\theta)}{2\theta}k\ln(k)+O(k)\quad(15)$$

by definition, which is negative. Therefore, the lemma follows from Eqs. (12), (14) and (15). □

Oliver Gebhard
Zimmerweg 3 60325 Frankfurt

Lebenslauf

Persönliche Daten

Geburtstag	01.08.1993
Geburtsort	Frankfurt am Main
Nationalität	Deutsch

Studium

04/2019 – 08/2022	Promotion „Mathematik“ an der Goethe Universität Frankfurt Gutachter: Prof. Dr. A.Coja-Oghlan, Prof. Dr. J. Scarlett
09/2017 – 03/2019	Master of Science „Mathematik“ Goethe Universität Frankfurt Nebenfach: Volkswirtschaft Abschlussarbeit: „The Group Testing Problem“ Gutachter: Prof. Dr. A.Coja-Oghlan, Dr. M. Hahn-Klimroth
09/2017 – 12/2017	Auslandssemester in „Economics“ an der University of Toronto
10/2013 – 03/2017	Bachelor of Science „Mathematik“ Goethe Universität Frankfurt Nebenfach: Finanzwirtschaft Abschlussarbeit: „Volumenberechnung konvexer Mengen“ Gutachter: Prof. Dr. A.Coja-Oghlan, Dr. N. Jafaari
07/2013	Abitur an Heinrich-von-Kleist Schule Eschborn

Berufliche Erfahrung

09/2021 – 09/2022	Wissenschaftlicher Mitarbeiter an der TU Dortmund in der Arbeitsgruppe von Prof. Dr. A.Coja-Oghlan im Fachgebiet „Effizient Algorithmen und Komplexitätstheorie“
04/2019 – 09/2021	Wissenschaftlicher Mitarbeiter an der Goethe Universität Frankfurt am Main in der Arbeitsgruppe von Prof. Dr. A.Coja-Oghlan im Fachgebiet „Probabilistische Kombinatorik“
01/2018 – 04/2018	Deutsche Bundesbank; Bereich bankenaufsichtliche Stresstests, Risikomodellierung und Forschung
08/2014 – 09/2014	KPMG International Cooperative; Tax Consulting

Ehrenamtliches Engagement

2012 – heute	Rettungsschwimmer bei der DLRG
2014 – 2019	Pfadfinderleiter bei der DPSG Eschborn
2015 - 2019	Mitglied des Buddy-Program der Goethe Universität

Vordrucke und Publikationen

2022 Statistical and Computational Phase Transitions in Group Testing
(Eingereicht und zur Zeit in Revision),
Autoren: A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth,
A. Wein, I. Zadik

2022 On the Parallel Reconstruction from Pooled Data
(zur Publikation in den Proceedings of IPDPS 2022 akzeptiert)
Autoren: O. Gebhard, M. Hahn-Klimroth, P. Loick, D. Kaaser

2022 On the Hierarchy of Distributed Majority Protocols
(Eingereicht und zur Zeit in Revision),
Autoren: P. Berenbrink, A. Coja-Oghlan, O. Gebhard,
M. Hahn-Klimroth, D. Kaaser, M. Rau

2021 Improved bounds for noisy group testing with constant tests per item
(in IEEE Transactions on Information Theory akzeptiert)
Autoren: O. Gebhard, O. Johnson, P. Loick, M. Rolvien

2021 Near optimal sparsity-constrained group testing: improved bounds and algorithms
(in IEEE Transactions on Information Theory akzeptiert)
Autoren: O. Gebhard, M. Hahn-Klimroth, O. Parczyk, M. Penschuck,
M. Rolvien, J. Scarlett, N. Tan

2020 Optimal group testing
(Proceedings of Machine Learning Research (COLT'20), Vol. 125,
Seiten: 1374-1388),
Autoren: A. Coja-Oghlan, O.Gebhard, M.Hahn-Klimroth, P. Loick

2019 Information-theoretic and algorithmic thresholds
for group testing (IEEE Transaction on Information Theory, Vol. 66,
Ausgabe: 12, Seiten: 7911-7928)

Wissenschaftliche Vorträge

2021	IEEE European School on Information Theory (Universität Stuttgart)
2020	Combinatorics Seminar (University of Illinois)
2019	Workshop on local algorithms (ETH Zürich)

Organisationskomitee

2022	Summer School on Algorithms, Dynamics and Information Flow in Networks (TU Dortmund)
2020	Workshop: Inference Problems: Algorithms and Lower Bounds (Goethe Universität)

Akademische Lehrer (Auswahl)

Prof. Dr. A. Bernig, Prof. Dr. A. Coja-Oghlan, Prof. Dr. C. Schnorr, Prof. Dr. T. Theobald, Prof. Dr. A. Wakolbinger