

A composition theorem for randomized query complexity via max conflict complexity *

Dmitry Gavinsky¹, Troy Lee², Miklos Santha^{3,4,5}, and Swagato Sanyal⁶

¹Institute of Mathematics, Czech Academy of Sciences, 115 67 Žitna 25, Praha 1, Czech Republic.

²Centre for Quantum Software and Information, Faculty of Engineering and Information Technology, University of Technology Sydney, Australia, troyjlee@gmail.com

³Centre for Quantum Technologies, National University of Singapore, Block S15, 3 Science Drive 2, Singapore 117543.

⁴IRIF, Université Paris Diderot, CNRS, 75205 Paris, France, santha@irif.fr

⁵MajuLab, UMI 3654, Singapore

⁶Indian Institute of Technology Kharagpur, swagato@cse.iitkgp.ac.in

November 27, 2018

Abstract

Let $R_\epsilon(\cdot)$ stand for the bounded-error randomized query complexity with error $\epsilon > 0$. For any relation $f \subseteq \{0, 1\}^n \times S$ and partial Boolean function $g \subseteq \{0, 1\}^m \times \{0, 1\}$, we show that $R_{1/3}(f \circ g^n) \in \Omega(R_{4/9}(f) \cdot \sqrt{R_{1/3}(g)})$, where $f \circ g^n \subseteq (\{0, 1\}^m)^n \times S$ is the composition of f and g . We give an example of a relation f and partial Boolean function g for which this lower bound is tight.

We prove our composition theorem by introducing a new complexity measure, the *max conflict complexity* $\bar{\chi}(g)$ of a partial Boolean function g . We show $\bar{\chi}(g) \in \Omega(\sqrt{R_{1/3}(g)})$ for any (partial) function g and $R_{1/3}(f \circ g^n) \in \Omega(R_{4/9}(f) \cdot \bar{\chi}(g))$; these two bounds imply our composition result. We further show that $\bar{\chi}(g)$ is always at least as large as the *sabotage complexity* of g , introduced by Ben-David and Kothari [BK16].

*This paper is a merger of [GLS18] and [San18], together with some new results.

1 Introduction

For Boolean functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $g : \{0, 1\}^m \rightarrow \{0, 1\}$, the composed Boolean function $f \circ g^n : (\{0, 1\}^m)^n \rightarrow \{0, 1\}$ is defined as $f \circ g^n(x_1, \dots, x_n) := f(g(x_1), \dots, g(x_n))$. A natural question in complexity theory is how the complexity of $f \circ g^n$ relates to the complexities of f and g . We study this question in the context of query complexity. A query algorithm for a function h queries bits of an input x , possibly in an adaptive manner, with the goal of outputting $h(x)$ after making as few queries as possible. The deterministic query complexity $D(h)$ is the minimum over all deterministic query algorithms that compute h of the worst-case number of queries made over all inputs x . For a composed function, it is easy to see that $D(f \circ g^n) \leq D(f) \cdot D(g)$, since $f \circ g^n$ can be computed by simulating an optimal query algorithm for f and serving every query of this algorithm by running an optimal query algorithm for g .

For many other natural measures of complexity as well, the complexity of $f \circ g^n$ is similarly bounded from above by the product of the complexities of f and g . As examples, this holds for randomized query complexity (up to a log factor), quantum query complexity, exact and approximate polynomial degree [She12]. Showing *lower bounds* on the complexity of $f \circ g^n$ that match these upper bounds as closely as possible is usually much more difficult; such results are often referred to as *composition theorems*. Besides answering a natural and interesting structural question, composition theorems find applications to constructing functions that are hard with respect to various measures of complexity, and separating complexity measures.

For deterministic query complexity, it can be shown by an adversary argument that $D(f \circ g^n) = D(f) \cdot D(g)$ [Mon14, Tal13], i.e., the query algorithm described above is an optimal query algorithm for $f \circ g^n$. For bounded-error quantum query complexity $Q(f)$, a perfect composition theorem is also known $Q(f \circ g^n) = \Theta(Q(f) \cdot Q(g))$ [HLS07, Rei11]. Of the three main variants of query complexity, the one that has thus far resisted a complete understanding of its nature under composition is randomized query complexity.

Recently, Ben-David and Kothari [BK16] made progress on this question by showing that for any partial function f and total function g ,

$$R_{1/3}(f \circ g^n) \in \Omega \left(R_{1/3}(f) \cdot \sqrt{\frac{R_0(g)}{\log R_0(g)}} \right). \quad (1)$$

They did this by introducing a new complexity measure, the *sabotage complexity*, $RS(g)$ of g and showing that $R_{1/3}(f \circ g^n) \in \Omega(R_{1/3}(f) \cdot RS(g))$, and then showing that $RS(g)$ is quadratically related to $R_0(g)$ for total functions g . For partial functions g , unbounded separations are known between $RS(g)$ and $R_{1/3}(g)$.

In this work, we introduce another complexity measure called the *max conflict complexity* $\bar{\chi}(g)$ of g . We show that $\bar{\chi}(g)$ is a quadratically tight lower bound on randomized query complexity, even for partial functions g .

Theorem 1. *For any partial Boolean function $g \subseteq \{0, 1\}^m \times \{0, 1\}$,*

$$\bar{\chi}(g) \in \Omega \left(\sqrt{R_{1/3}(g)} \right).$$

Our main application, and the motivation behind the definition of max conflict complexity, comes from showing a composition theorem. We use max conflict complexity to show a composition theorem in a more general setting where f can be a relation and g can be a partial function.

Theorem 2 (Main Theorem). *For any relation $f \subseteq \{0, 1\}^n \times S$ and partial Boolean function $g \subseteq \{0, 1\}^m \times \{0, 1\}$,*

$$R_{1/3}(f \circ g^n) \in \Omega(R_{4/9}(f) \cdot \bar{\chi}(g)).$$

Putting [Theorem 1](#) and [Theorem 2](#) together we have the following corollary.

Corollary 1. For any relation $f \subseteq \{0, 1\}^n \times S$ and partial Boolean function $g \subseteq \{0, 1\}^n \times \{0, 1\}$,

$$R_{1/3}(f \circ g^n) \in \Omega\left(R_{4/9}(f) \cdot \sqrt{R_{1/3}(g)}\right).$$

We further show that $\bar{\chi}(g)$ is always at least as large as the sabotage complexity of g .

Theorem 3. For any partial Boolean function $g \subseteq \{0, 1\}^m \times \{0, 1\}$,

$$\bar{\chi}(g) \geq \text{RS}(g) .$$

Thus for a partial function f and total function g , [Theorem 2](#) also implies the bound of Ben-David and Kothari [Eq. \(1\)](#).

Finally, we show that in the more general setting where f can be a relation and g a partial function, [Corollary 1](#) can be *tight*.

Theorem 4. There exists a relation $f_0 \subseteq \{0, 1\}^n \times \{0, 1\}^n$ and partial Boolean function $g_0 \subseteq \{0, 1\}^n \times \{0, 1\}$ such that

$$R_{4/9}(f_0) \in \Theta(\sqrt{n}), R_{1/3}(g_0) \in \Theta(n) \text{ and } R_{1/3}(f_0 \circ g_0^n) \in \Theta(n) .$$

1.1 Proof Technique

At a high level, the proof of [Theorem 2](#) follows the structure of the proof by Anshu et al. [\[AGJ+17\]](#) and Ben-David and Kothari [\[BK16\]](#). We show that for every probability distribution η over the input space $\{0, 1\}^n$ of f , there exists a query algorithm \mathcal{A} that makes $O(R_{1/3}(f \circ g^n) / \sqrt{R_{1/3}(g)})$ queries in the worst case, and computes f with high probability, $\Pr_{z \sim \eta}[(z, \mathcal{A}(z)) \in f] \geq 5/9$. By the minimax principle ([Fact 2](#)) this implies [Theorem 2](#).

We do this by using a query algorithm for $f \circ g^n$ to construct a query algorithm for f . We define a sampling procedure that for any $z \in \{0, 1\}^n$ samples $x = (x_1, \dots, x_n)$ such that $(z, s) \in f$ if and only if $(x, s) \in f \circ g^n$. This procedure is defined in terms of \mathcal{Q} , which is a probability distribution over pairs of distributions (μ_0, μ_1) , where μ_0 is supported on $g^{-1}(0)$ and μ_1 is supported on $g^{-1}(1)$. We define a distribution γ_η over $(\{0, 1\}^m)^n$ in terms of this sampling process as follows:

1. Sample $z = (z_1, \dots, z_n)$ from $\{0, 1\}^n$ according to η .
2. Independently sample $(\mu_0^{(i)}, \mu_1^{(i)})$ from \mathcal{Q} for $i = 1, \dots, n$.
3. Sample $x_i = (x_i^{(1)}, \dots, x_i^{(m)})$ according to $\mu_{z_i}^{(i)}$ for $i = 1, \dots, n$. Return $x = (x_1, \dots, x_n)$.

Notice that steps (1) and (2) are independent and the order in which they are performed does not matter. For future reference, for a fixed z let $\gamma_z(\mathcal{Q})$ be the probability distribution defined by the last two steps.

Now γ_η is simply a probability distribution over $(\{0, 1\}^m)^n$. Thus by the minimax principle ([Fact 2](#) below), there is a deterministic query algorithm \mathcal{A}' of worst-case complexity at most $R_{1/3}(f \circ g^n)$ such that $\Pr_{x \sim \gamma_\eta}[(x, \mathcal{A}'(x)) \in f \circ g^n] \geq 2/3$. We first use \mathcal{A}' to construct a randomized query algorithm T for f with bounded *expected* query complexity and error at most $1/3$. T is presented formally in [Algorithm 3](#). The final algorithm \mathcal{A} will be a truncation of T which has bounded worst-case complexity and error at most $4/9$.

On input z , the algorithm T seeks to sample a string x from $\gamma_z(\mathcal{Q})$, and run \mathcal{A}' on x . Put another way, $\gamma_z(\mathcal{Q})$ induces a probability distribution over the leaves of \mathcal{A}' , and the goal of T is to sample a leaf of \mathcal{A}' according to this distribution. Since for each $s \in S$, $(x, s) \in f \circ g^n$ if and

only if $(z, s) \in f$, and $\Pr_{x \sim \gamma_\eta}[(x, \mathcal{A}'(x)) \in f \circ g^n] \geq 2/3$, we have that $\Pr_{z \sim \eta}[(z, T(z)) \in f] \geq 2/3$. Thus T meets the accuracy requirement.

The catch, of course, is to specify how T samples from $\gamma_z(\mathcal{Q})$ *without making too many queries to z* . To sample x_i from $\mu_{z_i}^{(i)}$ seems to require knowledge of z_i , and thus T would have to query all of z .

To bypass this problem, we remember that \mathcal{A}' , being an efficient algorithm, will query only a few bits of x . This allows us to sample x bit by bit as and when they are queried by \mathcal{A}' . To see this more clearly, consider a run of T where the pairs of distributions $(\mu_0^{(1)}, \mu_1^{(1)}), \dots, (\mu_0^{(n)}, \mu_1^{(n)})$ were chosen in step (2) of the sampling procedure. Suppose that T is trying to simulate \mathcal{A}' at a vertex v where $x_i^{(j)}$ is queried. To respond to this query, T will sample $x_i^{(j)}$ from its marginal distribution according to $\mu_{z_i}^{(i)}$ conditioned on the event $x \in v$. Let the following be the marginal distributions of $x_i^{(j)}$ for the two possible values of z_i .

	$\Pr_{x_i \sim \mu_{z_i}^{(i)}}[x_i^{(j)} = 0 \mid x \in v]$	$\Pr_{x_i \sim \mu_{z_i}^{(i)}}[x_i^{(j)} = 1 \mid x \in v]$
$z_i = 0$	p_0	$1 - p_0$
$z_i = 1$	p_1	$1 - p_1$

Without loss of generality, assume that $p_0 \leq p_1$. T answers the query by the procedure BITSAMPLER given in Algorithm 1.

Algorithm 1: BITSAMPLER (suppose $p_0 \leq p_1$)

```

1 Sample  $r \sim [0, 1]$  uniformly at random.
2 if  $r < p_0$  then
3   | return 0.
4
5 else if  $r > p_1$  then
6   | return 1.
7
8 else
9   | query  $z_i$ .
10  | if  $r \leq p_{z_i}$  then
11  |   | return 0.
12  |   else
13  |   | return 1.

```

Note that the bit returned by BITSAMPLER has the desired distribution. The step in which BITSAMPLER returns the bit depends on the value of r sampled in step 1. In particular, z_i is queried if and only if $r \in [p_0, p_1]$, and the bit is returned in step 11 or 13. Such a query to z_i contributes to the query complexity of T . Thus the probability that T makes a query when the underlying simulation of \mathcal{A}' is at vertex v is $(p_1 - p_0)$. We refer to this quantity as $\Delta(v)$. It plays an important role in our analysis (see Section 5 and Appendix F.1).

Our sampling procedure and the tools we use to bound its cost is reminiscent of work of Barak et al. [BBCR13] in communication complexity. They look at a communication analog of our setting where two players are trying to sample a leaf in a communication protocol while communicating as little as possible.

1.1.1 Conflict complexity and max conflict complexity

Bounding the query complexity of T naturally suggests the quantities that we define in this work: the conflict complexity $\chi(g)$ and the max conflict complexity $\bar{\chi}(g)$ of a partial Boolean

function g . A formal definition can be found in Section 4; here we give the high-level idea and motivation behind these quantities.

Forget about T for a moment and just consider a deterministic query algorithm \mathcal{B} computing the partial function $g \subseteq \{0, 1\}^m \times \{0, 1\}$. Let μ_0, μ_1 be distributions with support on $g^{-1}(0), g^{-1}(1)$, respectively. For each vertex $v \in \mathcal{B}$ let $p_0(v)$ (respectively $p_1(v)$) be the probability that the answer to the query at v is 0 on input $x \sim \mu_0$ (respectively $x \sim \mu_1$), conditioned on x reaching v . Now we can imagine a process $\mathcal{P}(\mathcal{B}, \mu_0, \mu_1)$ that runs BITSAMPLER on the tree \mathcal{B} : $\mathcal{P}(\mathcal{B}, \mu_0, \mu_1)$ begins at the root, and at a vertex v in \mathcal{B} it uniformly chooses a random real number $r \in [0, 1]$. If $r < \min\{p_0(v), p_1(v)\}$ then the query is “answered” 0 and it moves to the left child. If $r > \max\{p_0(v), p_1(v)\}$ then the query is “answered” 1 and it moves to the right child. If $r \in [\min\{p_0(v), p_1(v)\}, \max\{p_0(v), p_1(v)\}]$ then the process halts. The conflict complexity $\chi(\mathcal{B}, (\mu_0, \mu_1))$ is the expected number of vertices this process visits before halting. The conflict complexity of g is defined to be

$$\chi(g) = \max_{(\mu_0, \mu_1)} \min_T \chi(T, (\mu_0, \mu_1)) \ ,$$

where the minimum is taken over trees T that compute g . For *max conflict complexity* we enlarge the set over which we maximize. Let \mathcal{Q} be a distribution over pairs of distributions (μ_0, μ_1) , where $\text{supp}(\mu_0) \subseteq g^{-1}(0), \text{supp}(\mu_1) \subseteq g^{-1}(1)$ for each pair (μ_0, μ_1) in the support of \mathcal{Q} . Let $\chi(\mathcal{B}, \mathcal{Q}) = \mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}} [\chi(\mathcal{B}, (\mu_0, \mu_1))]$. The max conflict complexity $\bar{\chi}(g)$ is defined as

$$\bar{\chi}(g) = \max_{\mathcal{Q}} \min_T \chi(T, \mathcal{Q}) \ ,$$

where the minimum is taken over trees T that compute g . Clearly, the max conflict complexity is at least as large as the conflict complexity.

To motivate the max conflict complexity, note that the query complexity of T is the number of times step 9 in BITSAMPLER is executed, i.e. when the random number $r \in [p_0, p_1]$. In the definition of T we will choose \mathcal{Q} to achieve the optimal value in the definition of $\bar{\chi}(g)$. Then intuitively one expects that for each i , T queries z_i only after \mathcal{A} makes about $\bar{\chi}(g)$ queries into x_i . By means of a direct sum theorem for max conflict complexity we make this intuition rigorous and prove that the expected query complexity of T is at most $R_{1/3}(f \circ g^n) / \bar{\chi}(g)$. We refer the reader to Section 6 for a formal proof.

1.1.2 $\bar{\chi}(g)$ and $R(g)$

Note that applying **Theorem 2** with the outer function $f(z) = z_1$ shows that $R_{1/3}(g) \in \Omega(\bar{\chi}(g))$. We complete the proof of **Corollary 1** by showing that max conflict complexity is a quadratically tight lower bound on randomized query complexity, even for partial functions g . In fact, in **Theorem 5** we show the stronger result that this is true even for the conflict complexity $R_{1/3}(g) \in O(\chi(g)^2)$.

To prove $R(g) \in O(\chi(g)^2)$, we again resort to the minimax principle; we show that for each probability distribution μ over the valid inputs to g , there is an accurate and efficient distributional query algorithm for g . For $b \in \{0, 1\}$, let μ_b be the distribution obtained by conditioning μ on the event $g(x) = b$. By the definition of $\chi(g)$, there is a query algorithm \mathcal{B} such that the following is true: if its queries are served by BITSAMPLER, step 9 is executed within expected $\chi(\mathcal{B}, \mu_0, \mu_1) \leq \chi(g)$ queries. Note that at a vertex v which queries i , the probability that step 9 is executed is $\Delta(v) = |\Pr_{\mu_0}[x_i = 0 \mid x \text{ at } v] - \Pr_{\mu_1}[x_i = 0 \mid x \text{ at } v]|$. This roughly implies that for a typical vertex v of \mathcal{B} , $\Delta(v)$ is at least about $\frac{1}{\chi(g)}$. By a technical claim that we prove (**Claim 15**) this implies that the query outcome at v carries about $\frac{1}{\chi(g)^2}$ bits of information about $g(x)$. Using the *chain rule of mutual information*, we can show that the mutual information between $g(x)$ and the outcomes of first $O(\chi(g)^2)$ queries by \mathcal{B} is $\Omega(1)$. This enables us to conclude that we can infer the value of $g(x)$ with success probability $1/2 + \Omega(1)$

from the transcript of \mathcal{B} restricted to the first $O(\chi(g)^2)$ queries. The distributional algorithm of g for μ is simply the algorithm \mathcal{B} terminated after $O(\chi(g)^2)$ queries.

1.1.3 $\bar{\chi}(g)$ and $\text{RS}(g)$

To see why $\bar{\chi}(g) \geq \text{RS}(g)$, we first give an alternative characterization of $\text{RS}(g)$. For a deterministic tree T computing g and strings x, y such that $g(x) \neq g(y)$, let $\text{sep}_T(x, y)$ be the depth of the node v in T such that x and y both reach v yet $x_{q(v)} \neq y_{q(v)}$, where $q(v)$ is the index queried at v . Let \mathcal{T} be a zero-error randomized protocol for g , i.e. \mathcal{T} is a probability distribution supported on deterministic trees that compute g . Then we have (for a proof see [Appendix B](#))

$$\text{RS}(g) = \min_{\mathcal{T}} \max_{\substack{x, y \\ g(x) \neq g(y)}} \mathbb{E}_{T \sim \mathcal{T}}[\text{sep}_T(x, y)] .$$

By von Neumann's minimax theorem [[von28](#)], this is equal to

$$\text{RS}(g) = \max_p \min_T \mathbb{E}_{(x, y) \sim p}[\text{sep}_T(x, y)] .$$

Here, the max is taken over distributions p on pairs (x, y) where $g(x) \neq g(y)$, and the min is taken over deterministic trees T computing g .

We have seen that the definition of $\bar{\chi}(g)$ is

$$\bar{\chi}(g) = \max_{\mathcal{Q}} \min_T \mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}} [\chi(T, (\mu_0, \mu_1))] ,$$

where \mathcal{Q} is a distribution over pairs (μ_0, μ_1) and T is a deterministic tree computing g . When (μ_0, μ_1) are taken to be singleton distributions, i.e. μ_0 puts all its weight on a single x with $g(x) = 0$, and μ_1 puts all its weight on a single y with $g(y) = 1$, it is easy to see that $\chi(T, (\mu_0, \mu_1)) = \text{sep}_T(x, y)$ (see [Claim 3](#)). Thus $\bar{\chi}(g)$ is at least as large as the sabotage complexity of g as \mathcal{Q} is allowed to be a distribution over general (μ_0, μ_1) , not just singleton distributions.

2 Preliminaries

Let $g \subseteq \{0, 1\}^m \times \{0, 1\}$ be a partial Boolean function. For $b \in \{0, 1\}$, $g^{-1}(b)$ is defined to be the set of strings x in $\{0, 1\}^m$ for which $(x, b) \in g$ and $(x, \bar{b}) \notin g$. We refer to $g^{-1}(0) \cup g^{-1}(1)$ as the set of valid inputs to g . We assume that for all strings $y \notin g^{-1}(0) \cup g^{-1}(1)$, both $(y, 0)$ and $(y, 1)$ are in g . For a string $x \in g^{-1}(0) \cup g^{-1}(1)$, $g(x)$ refers to the unique bit b such that $(x, b) \in g$. All the probability distributions μ over the domain of a partial Boolean function g in this paper are assumed to have support on $g^{-1}(0) \cup g^{-1}(1)$. Thus $g(x)$ is well-defined for any x in the support of μ .

Let S be any set. Let $h \subseteq \{0, 1\}^k \times S$ be any relation. Consider query algorithms \mathcal{A} that accept a string $x \in \{0, 1\}^k$ as input, query various bits of x , and produce an element of S as output. We denote the output by $\mathcal{A}(x)$.

Definition 1 (Deterministic query complexity). A deterministic query algorithm \mathcal{A} is said to compute h if $(x, \mathcal{A}(x)) \in h$ for all $x \in \{0, 1\}^k$. The deterministic query complexity $\text{D}(h)$ of h is the minimum over all deterministic query algorithms \mathcal{A} computing h of the maximum number of queries made by \mathcal{A} over $x \in \{0, 1\}^k$.

Definition 2 (Bounded-error randomized query complexity). Let $\epsilon \in [0, 1/2)$. We say that a randomized query algorithm \mathcal{A} computes h with error ϵ if $\Pr[(x, \mathcal{A}(x)) \in h] \geq 1 - \epsilon$ for all $x \in \{0, 1\}^k$. The bounded-error randomized query complexity $\text{R}_\epsilon(h)$ of h is the minimum over all randomized query algorithms \mathcal{A} computing h with error ϵ of the maximum number of queries made by \mathcal{A} over all $x \in \{0, 1\}^k$ and the internal randomness of \mathcal{A} .

Definition 3 (Distributional query complexity). Let μ a distribution on the input space $\{0, 1\}^k$ of h , and $\epsilon \in [0, 1/2)$. We say that a deterministic query algorithm \mathcal{A} computes h with distributional error ϵ on μ if $\Pr_{x \sim \mu}[(x, \mathcal{A}(x)) \in h] \geq 1 - \epsilon$. The distributional query complexity $D_\epsilon^\mu(h)$ of h is the minimum over deterministic algorithms \mathcal{A} computing h with distributional error ϵ on μ of the maximum over $x \in \{0, 1\}^k$ of the number of queries made by \mathcal{A} on x .

We will use the minimax principle in our proofs to go between distributional and randomized query complexity.

Fact 2 (Minimax principle). For any integer $k > 0$, set S , and relation $h \subseteq \{0, 1\}^k \times S$,

$$R_\epsilon(h) = \max_{\mu} D_\epsilon^\mu(h).$$

We present a proof of Fact 2 in Appendix A.

Let μ be a probability distribution over $\{0, 1\}^k$. We use $\text{supp}(\mu)$ to denote the support of μ . By $x \sim \mu$ we mean that x is a random string drawn from μ . Let $C \subseteq \{0, 1\}^k$ be an arbitrary set such that $\Pr_{x \sim \mu}[x \in C] = \sum_{y \in C} \mu(y) > 0$. Then $\mu \mid C$ is defined to be the probability distribution obtained by conditioning μ on the event that the sampled string belongs to C , i.e.,

$$(\mu \mid C)(x) = \begin{cases} 0 & \text{if } x \notin C \\ \frac{\mu(x)}{\sum_{y \in C} \mu(y)} & \text{if } x \in C \end{cases}$$

For a distribution \mathcal{Q} over pairs of distributions (μ_0, μ_1) , let $\text{supp}_0(\mathcal{Q}) = \cup\{\text{supp}(\mu_0) : \exists \mu_1, (\mu_0, \mu_1) \in \text{supp}(\mathcal{Q})\}$. Similarly let $\text{supp}_1(\mathcal{Q}) = \cup\{\text{supp}(\mu_1) : \exists \mu_0, (\mu_0, \mu_1) \in \text{supp}(\mathcal{Q})\}$. We say that \mathcal{Q} is *consistent* if $\text{supp}_0(\mathcal{Q})$ and $\text{supp}_1(\mathcal{Q})$ are disjoint sets. We say that \mathcal{Q} is consistent with a (partial) function g if $\text{supp}_0(\mathcal{Q}) \subseteq g^{-1}(0)$ and $\text{supp}_1(\mathcal{Q}) \subseteq g^{-1}(1)$.

Definition 4 (Subcube, co-dimension). A subset $C \subseteq \{0, 1\}^m$ is called a subcube if there exists a set $S \subseteq \{1, \dots, m\}$ of indices and an *assignment function* $\sigma : S \rightarrow \{0, 1\}$ such that $C = \{x \in \{0, 1\}^m : \forall i \in S, x_i = \sigma(i)\}$. The co-dimension $\text{codim}(C)$ of C is defined to be $|S|$.

Now we define the composition of a relation and a partial Boolean function.

Definition 5 (Composition of a relation and a partial Boolean function). Let $f \subseteq \{0, 1\}^n \times S$ and $g \subseteq \{0, 1\}^m \times \{0, 1\}$ be a relation and a partial Boolean function respectively. The composed relation $f \circ g^n \subseteq (\{0, 1\}^m)^n \times S$ is defined as follows: For $x = (x^{(1)}, \dots, x^{(n)}) \in (\{0, 1\}^m)^n$ and $s \in S$, $(x, s) \in f \circ g^n$ if and only if one of the following holds:

- $x_i \notin g^{-1}(0) \cup g^{-1}(1)$ for some $i \in \{1, \dots, n\}$.
- $x_i \in g^{-1}(0) \cup g^{-1}(1)$ for each $i \in \{1, \dots, n\}$ and $((g(x_1), \dots, g(x_n)), s) \in f$.

We will often view a deterministic query algorithm as a binary decision tree. In each vertex v of the tree, an input variable is queried. Depending on the outcome of the query, the computation goes to a child of v . The child of v corresponding to outcome b of the query is denoted by v_b .

The set of inputs that lead the computation of a decision tree to a certain vertex is a subcube. We will use the same symbol (e.g. v) to refer to a vertex as well as the subcube associated with it.

The execution of a decision tree terminates at some leaf. If the tree computes some relation $h \subseteq \{0, 1\}^k \times S$, the leaves are labelled by elements of S , and the tree outputs the label of the leaf at which it terminates. We will also consider decision tree with unlabelled leaves (see Section 4).

3 Conflict Complexity

In this section, we define the conflict complexity and max conflict complexity of a partial Boolean function g on m bits. For this, we will need to introduce some notation related to a deterministic decision tree T . For a node $v \in T$, let $\pi(v) = \perp$ if v is the root and $\pi(v)$ be the parent of v otherwise. Let $q(v)$ be the index that is queried at v in T , and let $d_T(v)$ be the number of vertices on the unique path in T from the root to v . The depth of the root is 1.

Now fix a partial function $g \subseteq \{0, 1\}^m \times \{0, 1\}$ and probability distributions μ_0, μ_1 over $g^{-1}(0), g^{-1}(1)$, respectively. Let T be a tree that computes g . For a node $v \in T$ let $p_0(v) = \Pr_{\mu_0}[x_{q(v)} = 0 | x \text{ at } v]$ and $p_1(v) = \Pr_{\mu_1}[x_{q(v)} = 0 | x \text{ at } v]$, and

$$R(v) = \begin{cases} 1 & \text{if } v \text{ is the root} \\ R(\pi(v)) \cdot \min\{\Pr_{\mu_0}[x \rightarrow v | x \text{ at } \pi(v)], \Pr_{\mu_1}[x \rightarrow v | x \text{ at } \pi(v)]\} & \text{otherwise .} \end{cases}$$

Also define

$$\Delta(v) = |p_0(v) - p_1(v)| .$$

To gather intuition about these quantities, imagine a random walk on T that begins at the root. At a node v , this walk moves to the left child with probability $\min\{p_0(v), p_1(v)\}$, and it moves to the right child with probability $1 - \max\{p_0(v), p_1(v)\}$. With the remaining probability, $\Delta(v)$, it terminates at v . Note that for any tree T computing g we have $\sum_{v \in T} \Delta(v)R(v) = 1$. This is because the walk always terminates before it reaches a leaf of T . In particular, this means that $\sum_{v \in T} d_T(v)\Delta(v)R(v)$ —the expected number of steps the walk takes before it terminates—is always at most the depth of the tree T .

Definition 6 (Conflict complexity and max conflict complexity). Let g be a partial function. For distributions μ_0, μ_1 with $\text{supp}(\mu_b) \subseteq g^{-1}(b)$ for $b \in \{0, 1\}$, and a deterministic decision tree T computing g , define

$$\chi(T, (\mu_0, \mu_1)) = \sum_{v \in T} d_T(v)\Delta(v)R(v) .$$

The conflict complexity of g is

$$\chi(g) = \max_{\mu_0, \mu_1} \min_T \chi(T, (\mu_0, \mu_1)) ,$$

where the maximum is over all pairs of distributions (μ_0, μ_1) supported on $g^{-1}(0)$ and $g^{-1}(1)$ respectively, and the minimum is taken over all deterministic trees T computing g . For \mathcal{Q} a distribution over pairs satisfying $\text{supp}_b(\mathcal{Q}) \subseteq g^{-1}(b)$ for $b \in \{0, 1\}$, and T a deterministic tree computing g , let $\chi(T, \mathcal{Q}) = \mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}}[\chi(T, (\mu_0, \mu_1))]$. Finally, the max conflict complexity of g is

$$\bar{\chi}(g) = \max_{\mathcal{Q}} \min_T \chi(T, \mathcal{Q}) ,$$

where the maximum is taken over \mathcal{Q} with $\text{supp}_b(\mathcal{Q}) \subseteq g^{-1}(b)$ for $b \in \{0, 1\}$, and the minimum is taken over deterministic trees T computing g .

We can extend the definition of conflict complexity and max conflict complexity to more general query processes that do not necessarily compute a function. We first need the notion of FULL.

Definition 7. For a deterministic tree T and pair of distributions (μ_0, μ_1) with disjoint support, we say that $(T, (\mu_0, \mu_1))$ is FULL if $\sum_{v \in T} \Delta(v)R(v) = 1$, i.e. if the random walk described above terminates with probability 1. We say that (T, \mathcal{Q}) is FULL if $(T, (\mu_0, \mu_1))$ is FULL for each $(\mu_0, \mu_1) \in \text{supp}(\mathcal{Q})$.

Definition 8. For a deterministic tree T and pair of distributions (μ_0, μ_1) such that $(T, (\mu_0, \mu_1))$ is FULL, define $\chi(T, (\mu_0, \mu_1)) = \sum_{v \in T} d_T(v)\Delta(v)R(v)$. For a distribution \mathcal{Q} such that (T, \mathcal{Q}) is FULL, define $\chi(T, \mathcal{Q}) = \mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}}[\chi(T, (\mu_0, \mu_1))]$.

3.1 Comparison with other query measures

Li [Li18] shows that the conflict complexity of a total Boolean function g is at least the block sensitivity of g . Here we show that the max conflict complexity of a (partial) function g is at least as large as the sabotage complexity of g . For a total Boolean function g , Ben-David and Kothari [BK16] show that the sabotage complexity of g is at least as large as the fractional block sensitivity of g [Aar08, Tal13, GSS16], which in turn is at least as large as the block sensitivity. They also show examples where the sabotage complexity is much larger than the partition bound, quantum query complexity and approximate polynomial degree, thus the same holds for max conflict complexity as well.

We first need the following simple claim. Let δ_x be the probability distribution that puts weight 1 on the string x .

Claim 3. Let T be a deterministic tree computing the partial function g and let x, y be such that $g(x) = 0, g(y) = 1$. Then

$$\chi(T, (\delta_x, \delta_y)) = \text{sep}_T(x, y) .$$

Proof. Let v_1 be the root of T , and v_1, v_2, \dots, v_t be the longest sequence of vertices in T that are visited both by x and y , i.e. $x_{q(v_i)} \neq y_{q(v_i)}$. For each $i = 1, \dots, t-1$ we see that $\Delta(v_i) = 0$, while $\Delta(v_t) = 1$. Also $R(v_i) = 1$ for each $i = 1, \dots, t$, while $R(v) = 0$ for any other vertex. Thus $\sum_{v \in T} d(v)\Delta(v)R(v) = t = \text{sep}_T(x, y)$. ■

Theorem 4. Let $g \subseteq \{0, 1\}^m \times \{0, 1\}$ be a partial function. Then $\bar{\chi}(g) \geq \text{RS}(g)$.

Proof. By Theorem 10,

$$\text{RS}(g) = \max_p \min_T \mathbf{E}_{(x,y) \sim p} [\text{sep}_T(x, y)] .$$

By definition of max conflict complexity we have

$$\bar{\chi}(g) = \max_{\mathcal{Q}} \min_T \mathbf{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}} [\chi(T, (\mu_0, \mu_1))] .$$

The distribution p in sabotage complexity is a special case of \mathcal{Q} where all the pairs of distributions in the support are singleton distributions. The theorem now follows from Claim 3. ■

4 Query Process

We now come to the most important definition of the paper, that of the query process $\mathcal{P}(\mathcal{B}, \mathcal{Q})$. Let $t > 0$ be any integer and \mathcal{B} be any deterministic query algorithm that runs on inputs in $(\{0, 1\}^m)^t$. Let $x = (x_i^{(j)})_{\substack{i=1, \dots, t \\ j=1, \dots, m}}$ be a generic input to \mathcal{B} , and let x_i stand for $(x_i^{(j)})_{j=1, \dots, m}$. For a vertex v of \mathcal{B} , let $v^{(i)}$ denote the subcube in v corresponding to x_i , i.e., $v = v^{(1)} \times \dots \times v^{(t)}$. Recall from Section 2 that v_b stands for the child of v corresponding to the query outcome being b , for $b \in \{0, 1\}$.

The query process $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ runs on an input $z \in \{0, 1\}^t$ and uses the BITSAMPLER (Algorithm 1) routine to simulate the queries of \mathcal{B} to x when it can. This process is the heart of how we will transform an algorithm for $f \circ g^n$ into a query efficient algorithm for f .

Definition 9 (Query process $\mathcal{P}(\mathcal{B}, \mathcal{Q})$). Let \mathcal{B} be a decision tree that runs on inputs $(\{0, 1\}^m)^t$. Let \mathcal{Q} be a consistent probability distribution over pairs of distributions (μ_0, μ_1) . The query process $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ is run on an input $z \in \{0, 1\}^t$ and is defined by Algorithm 2.

A few comments about Definition 9. First, we think of \mathcal{B} and \mathcal{P} as query procedures that query input variables and terminate. In particular, they do not have to produce outputs, i.e. their leaves do not have to be labeled. Also note that in Algorithm 2 the segment from line 9

to line 19 corresponds to the BITSAMPLER procedure in Algorithm 1. Queries to the input bits z_i are made in line 15, which corresponds to step 9 of BITSAMPLER.

We now present an important structural result about $\mathcal{P}(\mathcal{B}, \mathcal{Q})$. In particular, this formally proves that the procedure BITSAMPLER given in Algorithm 1 samples the bits from the right distribution.

Theorem 5 (Simulation Theorem). *Let \mathcal{B} be a deterministic decision tree running on inputs from $(\{0, 1\}^m)^t$, and let v be a vertex in \mathcal{B} . Let $A_z(v, \mathcal{Q})$ be the event that $\mathcal{P}(\mathcal{B}, \mathcal{Q})$, when run on z , reaches node v . Let $B_z(v, \mathcal{Q})$ be the event that for a random input x sampled from $\gamma_z(\mathcal{Q})$, the computation of \mathcal{B} reaches v . Then for every $z \in \{0, 1\}^t$ and each vertex v of \mathcal{B} ,*

$$\Pr[A_z(v, \mathcal{Q})] = \Pr[B_z(v, \mathcal{Q})] .$$

The proof of the claim is given in [Appendix D](#).

Algorithm 2: $\mathcal{P}(\mathcal{B}, \mathcal{Q})$

```

Input:  $z = (z_1, \dots, z_t) \in \{0, 1\}^t$ .
1 for  $1 \leq k \leq t$  do
2    $\text{QUERY}_k \leftarrow 0$ .
3    $\text{N}_k \leftarrow 0$ .
4   Sample  $(\mu_0^{(k)}, \mu_1^{(k)})$  from  $\mathcal{Q}$ .
5  $v \leftarrow \text{Root of } \mathcal{B}$  // Corresponds to  $(\{0, 1\}^m)^t$ 
6 while  $v$  is not a leaf of  $\mathcal{B}$  do
7   Let  $q(v) = (i, j)$ , the  $j^{\text{th}}$  coordinate of  $x_i$ 
8   if  $\text{QUERY}_i = 0$  then
9     Sample a fresh real number  $r \sim [0, 1]$  uniformly at random.
10    if  $r < \min_b \Pr_{x_i \sim \mu_b^{(i)}}[x_i^{(j)} = 0 \mid x_i \in v^{(i)}]$  then
11       $v \leftarrow v_0$ .
12    else if  $r > \max_b \Pr_{x_i \sim \mu_b^{(i)}}[x_i^{(j)} = 0 \mid x_i \in v^{(i)}]$  then
13       $v \leftarrow v_1$ .
14    else
15      Query  $z_i$ .  $\text{QUERY}_i \leftarrow 1$ .
16      if  $r \leq \Pr_{x_i \sim \mu_{z_i}^{(i)}}[x_i^{(j)} = 0 \mid x_i \in v^{(i)}]$  then
17         $v \leftarrow v_0$ .
18      else
19         $v \leftarrow v_1$ .
20     $\text{N}_i \leftarrow \text{N}_i + 1$ .
21  else
22     $b \leftarrow \begin{cases} 1 & \text{with probability } \Pr_{x_i \sim \mu_{z_i}^{(i)}}[x_i^{(j)} = 1 \mid x_i \in v^{(i)}] \\ 0 & \text{with probability } \Pr_{x_i \sim \mu_{z_i}^{(i)}}[x_i^{(j)} = 0 \mid x_i \in v^{(i)}] \end{cases}$ 
23     $v \leftarrow v_b$ 

```

We will be interested in the number of queries $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ is able to simulate before making a query to z_i . To this end, let the random variable $\mathcal{N}_i(\mathcal{B}, z, \mathcal{Q})$ stand for the value of the variable N_i in Algorithm 2 after the termination of $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ on input z . Note that \mathcal{N}_i depends on the randomness in the choices of r (Line 9) and also on the randomness in \mathcal{Q} in the choice of distributions $(\mu_0^{(k)}, \mu_1^{(k)})$ (Line 4).

4.1 Relating $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ to max conflict complexity

A key to our composition theorem will be relating the number of simulated queries made by $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ to max conflict complexity, which we do in this section. Let \mathcal{B} be a query algorithm taking inputs from $\{0, 1\}^m$. In this case, $\mathcal{N}_1(\mathcal{B}, 1, \mathcal{Q}) = \mathcal{N}_1(\mathcal{B}, 0, \mathcal{Q})$. This is because the behavior of $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ on input 0 is exactly the same as the behavior on input 1 before a query to z is made, and after z is queried the value of N_i does not change.

Claim 6. Let \mathcal{B} be an algorithm taking inputs from $\{0, 1\}^m$. Then $(\mathcal{B}, \mathcal{Q})$ is FULL if and only if $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ queries z with probability 1. If $(\mathcal{B}, \mathcal{Q})$ is FULL then

$$\chi(\mathcal{B}, \mathcal{Q}) = \mathbb{E}[\mathcal{N}_1(T, 1, \mathcal{Q})]$$

Proof. Note that until z is queried, $\mathcal{P}(\mathcal{B}, (\mu_0, \mu_1))$ exactly executes the random walk described in [Section 3](#), and querying z in $\mathcal{P}(\mathcal{B}, (\mu_0, \mu_1))$ corresponds to this random walk terminating. The first part of the claim then follows as $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ queries z with probability 1 if and only if $\mathcal{P}(\mathcal{B}, (\mu_0, \mu_1))$ queries z with probability 1 for every $(\mu_0, \mu_1) \in \text{supp}(\mathcal{Q})$.

Also because $\mathcal{P}(\mathcal{B}, (\mu_0, \mu_1))$ exactly executes the random walk described in [Section 3](#) we see that $\chi(\mathcal{B}, (\mu_0, \mu_1)) = \mathbb{E}[\mathcal{N}_1(T, 1, (\mu_0, \mu_1))]$. The second part of the claim follows by taking the expectation of this equality over $(\mu_0, \mu_1) \sim \mathcal{Q}$. \blacksquare

The correspondence of [Claim 6](#) prompts us to define FULL in a more general setting.

Definition 10 (FULL). Let \mathcal{B} be a query algorithm taking inputs from $(\{0, 1\}^m)^t$. The pair $(\mathcal{B}, \mathcal{Q})$ is said to be FULL if for every $z \in \{0, 1\}^t$ it holds that $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ queries z_i with probability 1, for every $i = 1, \dots, t$.

5 Conflict Complexity and Randomized Query Complexity

In this section, we will prove [Theorem 5](#) (restated below). Our proof relates the conflict complexity to the expected amount of information that is learned about the function value through each query via Pinsker's Inequality. At a high level, our proof is reminiscent of the result of [\[BBCR13\]](#) on compressing communication protocols in that both look at a random sampling process to navigate a tree, and relate the probability of this process needing to query or communicate at a node to the amount of information that is learned at the node.

Theorem 5. For any partial Boolean function $g \subseteq \{0, 1\}^m \times \{0, 1\}$,

$$\chi(g) \in \Omega\left(\sqrt{R_{1/3}(g)}\right).$$

Proof. We will show that there exists a constant $\varepsilon < 1/2$ such that for each input distribution μ , $D_\varepsilon^\mu(g) \leq 10\chi(g)^2$. [Theorem 5](#) will follow from the *minimax principle* ([Fact 2](#)), and the observation that the error can be brought down to $1/3$ by constantly many independent repetitions followed by a selection of the majority of the answers. It is enough to consider distributions μ supported on valid inputs of g . To this end, fix a distribution μ supported on $g^{-1}(0) \cup g^{-1}(1)$.

Let $\chi(g) = d$. Let \mathcal{B} be a deterministic query algorithm for inputs in $\{0, 1\}^m$ such that $(\mathcal{B}, \mu_0, \mu_1)$ is FULL and $\chi(\mu_0, \mu_1) = \chi(\mathcal{B}, \mu_0, \mu_1)$. We call such a decision tree an *optimal* decision tree for μ_0, μ_1 . Thus in $\mathcal{P}(\mathcal{B}, \mu_0, \mu_1)$, $\mathbb{E}[\mathcal{N}_1] = \chi(\mu_0, \mu_1) \leq d$. Recall from [Section 4](#) that the leaves of \mathcal{B} can be labelled by bits such that \mathcal{B} computes g on the supports of μ_0 and μ_1 . We assume \mathcal{B} 's leaves to be labelled as such.

Consider the following query algorithm \mathcal{B}' : Start simulating \mathcal{B} . Terminate the simulation if one of the following events occurs. The output in each case is specified below.

1. If $10d^2$ queries have been made and $v_{10d^2+1} \neq \perp$, terminate and output $\arg \max_b \Pr_{x \sim \mu}[g(x) = b \mid x \in v_{10d^2+1}]$.

2. If \mathcal{B} terminates, terminate and output what \mathcal{B} outputs.

By construction, \mathcal{B}' makes at most $10d^2$ queries in the worst case. The following claim bounds the error of \mathcal{B}' .

Claim 7. There exists constant $\varepsilon < 1/2$ such that $\Pr_{x \sim \mu}[\mathcal{B}'(x) \neq g(x)] \leq \varepsilon$. Furthermore, the constant ε is independent of μ .

Claim 7 is proven in Appendix F. This completes the proof of Theorem 5. \blacksquare

6 The Composition Theorem

In this section we prove Theorem 6 (restated below).

Theorem 6. Let \mathcal{S} be an arbitrary set, $f \subseteq \{0, 1\}^n \times \mathcal{S}$ be a relation and $g \subseteq \{0, 1\}^m \times \{0, 1\}$ a partial Boolean function. Then,

$$R_{1/3}(f \circ g^n) \in \Omega(R_{4/9}(f) \cdot \bar{\chi}(g)).$$

Our proof will make use of the following *direct sum theorem*, which we prove in Appendix E.

Theorem 7. Let \mathcal{B} be a query algorithm acting on inputs from $(\{0, 1\}^m)^t$. Let \mathcal{Q} be a consistent distribution over pairs of distributions (μ_0, μ_1) on m -bit strings. If $(\mathcal{B}, \mathcal{Q})$ is FULL then for any $z \in \{0, 1\}^t$

$$\sum_{i=1}^t \mathbb{E}[\mathcal{N}_i(\mathcal{B}, z, \mathcal{Q})] \geq t \cdot \min_{\mathcal{C}} \chi(\mathcal{C}, \mathcal{Q}) ,$$

where the minimum is taken over deterministic trees \mathcal{C} acting on inputs from $\{0, 1\}^m$ such that $(\mathcal{C}, \mathcal{Q})$ is FULL.

Proof of Theorem 6. We shall prove that for each distribution η on the inputs to f , there is a randomized query algorithm \mathcal{A} making at most $9R_{1/3}(f \circ g^n)/\bar{\chi}(g)$ queries in the worst case, for which $\Pr_{z \in \eta}[(z, \mathcal{A}(z)) \in f] \geq \frac{5}{9}$ holds. \mathcal{A} can be made deterministic with the same complexity and accuracy guarantees by appropriately fixing its randomness. This will imply the theorem by the *minimax principle* (Fact 2). To this end let us fix a distribution η over $\{0, 1\}^n$.

Let \mathcal{Q} be consistent with g such that for any deterministic decision tree \mathcal{C} computing g we have $\chi(\mathcal{C}, \mathcal{Q}) \geq \bar{\chi}(g)$. We will use distributions η and \mathcal{Q} to set up a distribution γ_η over the input space of $f \circ g^n$. This distribution is defined as follows:

1. Sample $z = (z_1, \dots, z_n)$ from η .
2. Sample $(\mu_0^{(i)}, \mu_1^{(i)})$ independently from \mathcal{Q} for $i = 1, \dots, t$.
3. Sample x_i from $\mu_{z_i}^{(i)}$ for $i = 1, \dots, t$. Return $x = (x_1, \dots, x_t)$.

Recall from Section 1.1 the observation that for each z, x sampled as above, for each $s \in \mathcal{S}$, $(z, s) \in f$ if and only if $(x, s) \in f \circ g^n$.

Assume that $R_{1/3}(f \circ g^n) = c$. The minimax principle (Fact 2) implies that there is a deterministic query algorithm \mathcal{A}' for inputs from $(\{0, 1\}^m)^n$, that makes at most c queries in the worst case, such that $\Pr_{x \in \gamma_\eta}[(x, \mathcal{A}'(x)) \in f \circ g^n] \geq \frac{2}{3}$. We will first use \mathcal{A}' to construct a randomized algorithm T for f whose accuracy under the distribution η is as desired and which, for every input z , makes few queries in expectation. T is described in Algorithm 3.

First we bound the probability of error by T . By Theorem 5, we have that $\Pr[(z, T(z)) \in f] = \Pr_{x \sim \gamma_z(\mathcal{Q})}[(x, \mathcal{A}'(x)) \in f \circ g^n]$ for each $z \in \{0, 1\}^n$. Thus, $\Pr_{z \sim \eta}[(z, T(z)) \in f] = \Pr_{x \sim \gamma_\eta}[(x, \mathcal{A}'(x)) \in f \circ g^n] \geq \frac{2}{3}$.

Next, we bound the expected number of queries made by T in the worst-case.

Algorithm 3: T

Input: $z \in \{0, 1\}^n$

- 1 Run $\mathcal{P}(\mathcal{A}', \mathcal{Q})$ on z .
 - 2 Return the output of \mathcal{A}' .
-

Claim 8. The expected number of queries made by T on each input z is at most $\frac{c}{\bar{\chi}(g)}$.

Proof. Fix an input $z \in \{0, 1\}^n$. For each leaf $\ell = \ell^{(1)} \times \dots \times \ell^{(n)}$ of \mathcal{A}' and for each $i = 1, \dots, n$ define $\mathcal{E}'_{i,\ell}$ to be the event that the computation of $\mathcal{P}(\mathcal{A}', \mathcal{Q})$ finishes at ℓ with $\text{QUERY}_i = 0$. For $i = 1, \dots, n$ define \mathcal{F}'_i to be the event that QUERY_i is set to 1 in $\mathcal{P}(\mathcal{A}', \mathcal{Q})$. Let \mathcal{Q}^n stand for the (product) distribution of n pairs of probability distributions each independently sampled from \mathcal{Q} . For i, ℓ such that g is not constant on $\ell^{(i)}$, let $\mathcal{D}_{i,\ell}$ be the distribution given by the following sampling procedure:

1. Sample $(\mu_0^{(1)}, \mu_1^{(1)}), \dots, (\mu_0^{(n)}, \mu_1^{(n)})$ from \mathcal{Q}^n conditioned on $\mathcal{E}'_{i,\ell}$,
2. return $(\mu_0^{(i)} | \ell^{(i)}, \mu_1^{(i)} | \ell^{(i)})$.

Let $\mathcal{B}_{i,\ell}$ be an optimal tree for $\mathcal{D}_{i,\ell}$, i.e., $\chi(\mathcal{B}_{i,\ell}, \mathcal{D}_{i,\ell}) = \min_{\mathcal{C}} \chi(\mathcal{C}, \mathcal{D}_{i,\ell})$, where the minimization is over all algorithms \mathcal{C} that output 0 on $\text{supp}_0(\mathcal{D}_{i,\ell})$ and output 1 on $\text{supp}_1(\mathcal{D}_{i,\ell})$. Now, consider the query algorithm H defined in Algorithm 4. Note that (H, \mathcal{Q}) is FULL. Now consider a

Algorithm 4: H

Input: $x \in (\{0, 1\}^m)^n$

- 1 Run \mathcal{A}' on x .
 - 2 Let \mathcal{A}' terminate at leaf $\ell = \ell^{(1)} \times \dots \times \ell^{(n)}$.
 - 3 **for** $1 \leq i \leq n$ **do**
 - 4 **if** g is not constant on $\ell^{(i)}$ **then**
 - 5 Run $\mathcal{B}_{i,\ell}$ on x_i .
-

run of the query process $\mathcal{P}(H, \mathcal{Q})$ on input z . **Theorem 7** implies that $\sum_{i=1}^n \mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q})] \geq n \cdot \min_{\mathcal{C}} \chi(\mathcal{C}, \mathcal{Q}) = n \cdot \bar{\chi}(g)$, by the choice of \mathcal{Q} .

Let \mathcal{F}_i to be the event that QUERY_i is set to 1 in $\mathcal{P}(H, \mathcal{Q})$ when it reaches a leaf of \mathcal{A}' , and for each leaf ℓ of \mathcal{A}' let $\mathcal{E}_{i,\ell}$ be the event that $\mathcal{P}(H, \mathcal{Q})$ reaches ℓ and $\text{QUERY}_i = 0$ when it does. Observe that for each $i = 1, \dots, n$, the events $\{\mathcal{F}_i, (\mathcal{E}_{i,\ell})_{\ell}\}$ are mutually exclusive and exhaustive.

We have that

$$\begin{aligned} n \cdot \bar{\chi}(g) &\leq \sum_{i=1}^n \mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q})] \\ &= \sum_{i=1}^n \sum_{\ell} \Pr[\mathcal{E}_{i,\ell}] \cdot \mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q}) | \mathcal{E}_{i,\ell}] + \sum_{i=1}^n \Pr[\mathcal{F}_i] \cdot \mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q}) | \mathcal{F}_i] \end{aligned} \quad (2)$$

Let $d_i(\ell)$ be the number of queries into x_i made in the unique path from the root of \mathcal{A}' to ℓ . Now, condition on the n pairs of distributions $(\mu_0^{(j)}, \mu_1^{(j)})_{j=1, \dots, n}$ that are used in $\mathcal{P}(H, \mathcal{Q})$. We have that,

$$\mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q}) | \mathcal{E}_{i,\ell}, (\mu_0^{(j)}, \mu_1^{(j)})_{j=1, \dots, n}] = d_i(\ell^{(i)}) + \mathbb{E}[\mathcal{N}_1(\mathcal{B}_{i,\ell}, z_i, (\mu_0^{(i)} | \ell^{(i)}, \mu_1^{(i)} | \ell^{(i)}))]. \quad (3)$$

Averaging over $(\mu_0^{(j)}, \mu_1^{(j)})_{j=1, \dots, n}$ we have from (3) that

$$\begin{aligned} \mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q}) \mid \mathcal{E}_{i,\ell}] &= d_i(\ell^{(i)}) + \mathbb{E}[\mathcal{N}_1(\mathcal{B}_{i,\ell}, z_i, \mathcal{D}_{i,\ell})] \\ &= d_i(\ell^{(i)}) + \min_{\mathcal{C}} \chi(\mathcal{C}, \mathcal{D}_{i,\ell}) \quad (\text{By the choice of } \mathcal{B}_{i,\ell}). \\ &\leq d_i(\ell^{(i)}) + \bar{\chi}(g). \end{aligned} \tag{4}$$

Observing that $\sum_{\ell} \Pr[\mathcal{E}_{i,\ell}] = 1 - \Pr[\mathcal{F}_i]$, we have from (2) and (4) that

$$\begin{aligned} n \cdot \bar{\chi}(g) &\leq \sum_{i=1}^n \sum_{\ell} \Pr[\mathcal{E}_{i,\ell}] \cdot (d_i(\ell^{(i)}) + \bar{\chi}(g)) + \sum_{i=1}^n \Pr[\mathcal{F}_i] \cdot \mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q}) \mid \mathcal{F}_i] \\ &= \sum_{i=1}^n (1 - \Pr[\mathcal{F}_i]) \cdot \bar{\chi}(g) + \sum_{i=1}^n \left(\sum_{\ell} (\Pr[\mathcal{E}_{i,\ell}] \cdot d_i(\ell^{(i)}) + \Pr[\mathcal{F}_i] \cdot \mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q}) \mid \mathcal{F}_i]) \right) \\ \Rightarrow \sum_{i=1}^n \Pr[\mathcal{F}_i] &\leq \frac{1}{\bar{\chi}(g)} \cdot \sum_{i=1}^n \left(\sum_{\ell} (\Pr[\mathcal{E}_{i,\ell}] \cdot d_i(\ell^{(i)}) + \Pr[\mathcal{F}_i] \cdot \mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q}) \mid \mathcal{F}_i]) \right). \end{aligned} \tag{5}$$

We will show that $\sum_{i=1}^n (\sum_{\ell} (\Pr[\mathcal{E}_{i,\ell}] \cdot d_i(\ell^{(i)}) + \Pr[\mathcal{F}_i] \cdot \mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q}) \mid \mathcal{F}_i])) \leq c$. Since $\sum_{i=1}^n \Pr[\mathcal{F}_i]$ is exactly the expected number of queries made by T , the claim will follow from (5).

Consider a run of $\mathcal{P}(H, \mathcal{Q})$ on input z , and let c_i be a random variable denoting the number of times Line 7 of Algorithm 2 (with $\mathcal{B} = H$) is a query into x_i before a leaf of \mathcal{A}' is reached, for $i = 1, \dots, n$. Thus $\sum_{i=1}^n \mathbb{E}[c_i] \leq c$. Further, for each i, ℓ we have $d_i(\ell^{(i)}) = \mathbb{E}[c_i \mid \mathcal{E}_{i,\ell}]$ and $\mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q}) \mid \mathcal{F}_i] = \mathbb{E}[\mathcal{N}_i(\mathcal{A}', z, \mathcal{Q}) \mid \mathcal{F}_i] \leq \mathbb{E}[c_i \mid \mathcal{F}_i]$. Thus,

$$\begin{aligned} \sum_{i=1}^n \left(\sum_{\ell} (\Pr[\mathcal{E}_{i,\ell}] \cdot d_i(\ell^{(i)}) + \Pr[\mathcal{F}_i] \cdot \mathbb{E}[\mathcal{N}_i(H, z, \mathcal{Q}) \mid \mathcal{F}_i]) \right) \\ \leq \sum_{i=1}^n \left(\sum_{\ell} (\Pr[\mathcal{E}_{i,\ell}] \cdot \mathbb{E}[c_i \mid \mathcal{E}_{i,\ell}] + \Pr[\mathcal{F}_i] \cdot \mathbb{E}[c_i \mid \mathcal{F}_i]) \right) \\ = \sum_{i=1}^n \mathbb{E}[c_i] \leq c. \end{aligned}$$

■

Now we finish the proof of Theorem 6 by constructing the query algorithm \mathcal{A} . Let \mathcal{A} be the algorithm obtained by terminating T after $9t/d$ queries. By Markov's inequality, for each z , the probability that T makes more than $9t/d$ queries is at most $1/9$. Thus \mathcal{A} computes f with probability at least $2/3 - 1/9 = 5/9$ on a random input from η . ■

7 Tightness: $\mathbf{R}_{1/3}(f \circ g^n) \in O(\mathbf{R}_{4/9}(f) \cdot \sqrt{\mathbf{R}_{1/3}(g)})$ is possible

In this section we prove Theorem 4. We construct a relation $f_0 \subseteq \{0, 1\}^n \times \{0, 1\}^n$ (i.e., $\mathcal{S} = \{0, 1\}^n$) and a promise function $g_0 \subseteq \{0, 1\}^n \times \{0, 1\}$ (i.e., $m = n$), such that $\mathbf{R}_{4/9}(f_0) \in \Theta(\sqrt{n})$, $\mathbf{R}_{1/3}(g_0) \in \Theta(n)$ and $\mathbf{R}_{1/3}(f_0 \circ g_0^n) \in \Theta(n)$.

For strings $x = (x_1, \dots, x_n), z = (z_1, \dots, z_n)$ in $\{0, 1\}^n$, let $x \oplus z$ be the string $(x_1 \oplus z_1, \dots, x_n \oplus z_n)$ obtained by taking their bitwise XOR. Let $|x|$ stand for the *Hamming weight* $|\{i \in [n] : x_i = 1\}|$ of x . We define f_0 as follows:

$$f_0(z) \stackrel{\text{def}}{=} \left\{ (a, z) \in \{0, 1\}^n \times \{0, 1\}^n \mid |a \oplus z| \leq \frac{n}{2} - \sqrt{n} \right\}$$

Now we define g_0 by specifying $g_0^{-1}(0)$ and $g_0^{-1}(1)$.

$$\begin{aligned} g_0^{-1}(0) &\stackrel{\text{def}}{=} \{(x, 0) \mid x \in \{0, 1\}^n, |x| \leq n/2 - \sqrt{n}\}, \\ g_0^{-1}(1) &\stackrel{\text{def}}{=} \{(x, 1) \mid x \in \{0, 1\}^n, |x| \geq n/2 + \sqrt{n}\}. \end{aligned}$$

We now determine the randomized query complexities of f_0, g_0 and $f_0 \circ g_0^n$.

Claim 9. (i) $R_{4/9}(f_0) \in \Omega(\sqrt{n})$.

(ii) $R_{1/3}(g_0) \in \Omega(n)$.

(iii) $R_\varepsilon(f_0 \circ g_0^n) \in O\left(n \cdot \sqrt{\log(1/\varepsilon)}\right)$.

We prove Claim 9 in Appendix G. Theorem 4 follows from Theorem 2 and Claim 9 with ε set to $1/3$.

Acknowledgements. We thank Rahul Jain for useful discussions. We thank Srijita Kundu and Jevgēnijs Vihrovs for their helpful comments on the manuscript. T.L. would like to thank the Simons Institute and the organizers of the “Workshop on Interactive Complexity” where part of this work took place. In particular, we thank Yuval Filmus for the suggestion during this workshop to look at the min-max version of conflict complexity, which led to the development of max conflict complexity.

Part of this work was conducted while T.L. and S.S. were at the Nanyang Technological University and the Centre for Quantum Technologies, supported by the Singapore National Research Foundation under NRF RF Award No. NRF-NRFF2013-13. This work was additionally supported by the Singapore National Research Foundation, the Prime Minister’s Office, Singapore and the Ministry of Education, Singapore under the Research Centres of Excellence programme under research grant R 710-000-012-135. This research was supported in part by the QuantERA ERA-NET Cofund project QuantAlgo. D.G. is partially funded by the grant P202/12/G061 of GA ĀR and by RVO: 67985840.

A Minimax principle: proof of Fact 2

Fix an integer ℓ . Let \mathcal{D}_ℓ be the finite set of all deterministic query algorithms on k bits with worst-case complexity at most ℓ . Let $\mathcal{H}_k := \{0, 1\}^k$. For algorithm $A \in \mathcal{D}_\ell$ and input $x \in \mathcal{H}_k$, let $E(A, x) = 1$ if $(x, A(x)) \notin h$, and 0 otherwise. By von Neumann’s minimax principle,

$$\min_{\sigma} \max_{\mu} \sum_{A \in \mathcal{D}_\ell, x \in \mathcal{H}_k} \sigma(A) E(A, x) \mu(x) = \max_{\mu} \min_{\sigma} \sum_{A \in \mathcal{D}_\ell, x \in \mathcal{H}_k} \sigma(A) E(A, x) \mu(x), \quad (6)$$

where σ and μ range over probability distributions over \mathcal{D}_ℓ and \mathcal{H}_k respectively. Note that in equation (6), we can assume that the maximum in the left hand side is over point distributions on \mathcal{H}_k , i.e., distributions that assign weight 1 to some input $x \in \mathcal{H}_k$. Similarly we can assume that the minimum in the right hand side is over point distributions on \mathcal{D}_ℓ . Thus we have that,

$$\min_{\sigma} \max_{x \in \mathcal{H}_k} \sum_{A \in \mathcal{D}_\ell} \sigma(A) E(A, x) = \max_{\mu} \min_{A \in \mathcal{D}_\ell} \sum_{x \in \mathcal{H}_k} E(A, x) \mu(x). \quad (7)$$

From equation (7) it follows that

$$R_\varepsilon(h) = \min \left\{ \ell \left| \min_{\sigma} \max_{x \in \mathcal{H}_k} \sum_{A \in \mathcal{D}_\ell} \sigma(A) E(A, x) \leq \varepsilon \right. \right\}$$

$$\begin{aligned}
& \text{(where } \sigma \text{ ranges over all probability distributions on } \mathcal{D}_\ell) \\
& = \min \left\{ \ell \left| \max_{\mu} \min_{A \in \mathcal{D}_\ell} \sum_{x \in \mathcal{H}_k} \mathbb{E}(A, x) \mu(x) \leq \epsilon \right. \right\} \\
& \quad \text{(where } \mu \text{ ranges over all probability distributions on } \mathcal{H}_k) \\
& = \max_{\mu} \min \left\{ \ell \left| \min_{A \in \mathcal{D}_\ell} \sum_{x \in \mathcal{H}_k} \mathbb{E}(A, x) \mu(x) \leq \epsilon \right. \right\} \\
& = D_\epsilon^\mu(h).
\end{aligned}$$

B Alternative characterization of sabotage complexity

We first go over the standard definition of sabotage complexity from [BK16]. Let $g \subseteq \{0, 1\}^m \times \{0, 1\}$ be a partial function. From g , define a partial function $g_{\text{sab}} : P \rightarrow \{*, \dagger\}$, where now $P \subseteq \{0, 1, *, \dagger\}^n$ is defined in the following way. Let $P^* \subseteq \{0, 1, *\}$ be the largest set such that for all $z \in P^*$ there exist x, y with $g(x) \neq g(y)$ and both x and y are consistent with the non-star coordinates of z . Define $P^\dagger \subseteq \{0, 1, \dagger\}$ analogously with \dagger instead of $*$. Then $P = P^* \cup P^\dagger$. Finally, define $g_{\text{sab}}(z) = *$ if $z \in P^*$ and $g_{\text{sab}}(z) = \dagger$ if $z \in P^\dagger$. The sabotage complexity of g is defined as $\text{RS}(g) = R_0(g_{\text{sab}})$.

For a tree T computing g , and strings x, y such that $g(x) \neq g(y)$, let $\text{sep}_T(x, y)$ denote the depth of the node v in T such that x and y both reach v yet $x_{q(v)} \neq y_{q(v)}$ where $q(v)$ is the index queried at node v . We have the following alternative characterization of sabotage complexity.

Theorem 10. *Let $g \subseteq \{0, 1\}^m \times \{0, 1\}$ be a partial function. Then*

$$\begin{aligned}
\text{RS}(g) &= \min_{\mathcal{T}} \max_{\substack{x, y \\ g(x) \neq g(y)}} \mathbb{E}_{T \sim \mathcal{T}}[\text{sep}_T(x, y)] \\
&= \max_p \min_T \mathbb{E}_{(x, y) \sim p}[\text{sep}_T(x, y)] .
\end{aligned}$$

In the first equation the minimum is taken over zero-error randomized algorithms \mathcal{T} for g . In the second equation, the maximum is taken over distributions over pairs (x, y) where $g(x) = 0, g(y) = 1$, and the minimum is taken over deterministic trees T computing g .

Proof. That the right hand side of the first line is equal to the second line follows by von Neumann's minimax theorem [von28].

Now we focus on establishing the first line. We first show that $\text{RS}(g)$ is at most the right hand side of the first line. Let \mathcal{T}^* achieve the minimum of the expression on the right hand side. Let $z \in P$ be any sabotaged input. Then there are x^*, y^* with $g(x^*) \neq g(y^*)$ such that x^* and y^* only differ where z has special symbols. Thus any query that separates x^* and y^* will also find a special symbol. The expected number of queries to separate x^* and y^* is at most $\max_{x, y} \mathbb{E}_{T \sim \mathcal{T}^*}[\text{sep}_T(x, y)]$, thus the left hand side is at most the right hand side.

For the other direction, let \mathcal{T}^* be an optimal zero-error randomized algorithm computing g_{sab} . For any x, y with $g(x) \neq g(y)$ we can create $z^* \in P^*$ such that z^* has $*$ in those positions where x, y disagree, and z^* agrees with x, y in those positions where they agree with each other.. Let z^\dagger equal z^* with $*$ replaced by \dagger . Now \mathcal{T}^* is able to distinguish between z^* and z^\dagger using an expected number of queries that is at most $\text{RS}(g)$. Any query that distinguishes z^* and z^\dagger is also a query that separates x and y , as z_* and z_\dagger only differ where x and y do. This means

$$\mathbb{E}_{T \sim \mathcal{T}^*}[\text{sep}_T(x, y)] \leq \text{RS}(g) ,$$

showing that the right hand side is at most the left hand side. ■

C Information Theory

Let X be a random variable supported on a finite set $\{x_1, \dots, x_s\}$. Let \mathcal{E} be any event in the same probability space. Let $\mathbb{P}[\cdot]$ denote the probability of any event. The *conditional entropy* $H(X | \mathcal{E})$ of X conditioned on \mathcal{E} is defined as follows.

Definition 11 (Conditional entropy).

$$H(X | \mathcal{E}) := \sum_{i=1}^s \mathbb{P}[X = x_i | \mathcal{E}] \log_2 \frac{1}{\mathbb{P}[X = x_i | \mathcal{E}]}.$$

An important special case is when \mathcal{E} is the entire sample space. In that case the above conditional entropy is referred to as the entropy $H(X)$ of X .

Definition 12 (Entropy).

$$H(X) := \sum_{i=1}^s \mathbb{P}[X = x_i] \log_2 \frac{1}{\mathbb{P}[X = x_i]}.$$

Let Y be another random variable in the same probability space as X , taking values from a finite set $\{y_1, \dots, y_t\}$. Then the conditional entropy of X conditioned on Y , $H(X | Y)$, is defined as follows.

Definition 13.

$$H(X | Y) = \sum_{i=1}^t \mathbb{P}[Y = y_i] \cdot H(X | Y = y_i).$$

Definition 14 (Mutual information). Let X , Y and Z be two random variables in the same probability space, taking values from finite sets. The mutual information between X and Y conditioned on Z , $I(X; Y | Z)$, is defined as follows.

$$I(X; Y | Z) := H(X | Z) - H(X | Y, Z).$$

It can be shown that $I(X; Y | Z)$ is symmetric in X and Y : $I(X; Y | Z) = I(Y; X | Z) = H(Y | Z) - H(Y | X, Z)$.

Theorem 11 (Chain rule of mutual information). *Let X_1, \dots, X_k, Y, Z be random variables in the same probability space, taking values from finite sets. Then,*

$$I(X_1, \dots, X_k : Y | Z) = \sum_{i=1}^k I(X_i : Y | Z, X_1, \dots, X_{i-1}).$$

Definition 15 (Kullback-Leibler Divergence). Given two probability distributions P and Q on a finite set \mathcal{U} , the *Kullback-Leibler divergence* from Q to P , denoted by $D(P||Q)$, is defined as:

$$D(P||Q) := - \sum_{u \in \mathcal{U}} P(u) \log \frac{P(u)}{Q(u)}.$$

Given two random variables X and Y taking values in a finite set \mathcal{U} , Let $X \otimes Y$ be distribution over ordered pairs of elements of \mathcal{U} (i.e., over elements of $\mathcal{U} \times \mathcal{U}$), where the elements are sampled independently according to distributions of X and Y respectively. Let (X, Y) denote the joint distribution of X and Y . The following fact can be easily verified.

Fact 12. $I(X : Y) = D((X, Y) || (X \otimes Y))$.

Definition 16. Given two probability distributions P and Q on a finite set \mathcal{U} , the L_1 -distance between P and Q , denoted by $\|P - Q\|_1$, is defined as:

$$\|P - Q\|_1 := \sum_{u \in \mathcal{U}} |P(u) - Q(u)|.$$

Pinsker's inequality, stated below, bounds $D(P||Q)$ in terms of $\|P - Q\|_1$ from below.

Theorem 13 (Pinsker's inequality). *Given two probability distributions P and Q on a finite set \mathcal{U} ,*

$$D(P||Q) \geq \frac{1}{2} \|P - Q\|_1^2.$$

D Proof of Theorem 5

In this section we show that $\Pr[A_z(v, \mathcal{Q})] = \Pr[B_z(v, \mathcal{Q})]$. To save writing, we fix $z \in \{0, 1\}^t$ and \mathcal{Q} and let $A(v) := A_z(v, \mathcal{Q})$ be the event that \mathcal{P} reaches node v on input z , and $B(v) := B_z(v, \mathcal{Q})$ be the event that \mathcal{B} reaches node v under the distribution $\gamma_z(\mathcal{Q})$. Additionally, we write $\overline{(\mu_0, \mu_1)} = ((\mu_0^{(1)}, \mu_1^{(1)}), \dots, (\mu_0^{(t)}, \mu_1^{(t)}))$ for a t -tuple of pairs of distributions. In the following when we write $\mathbb{E}_{\overline{(\mu_0, \mu_1)} \sim \mathcal{Q}^t}$ this expectation is taken with respect to drawing each $(\mu_0^{(i)}, \mu_1^{(i)})$ independently from \mathcal{Q} .

Now notice that $\Pr[A(v)] = \mathbb{E}_{\overline{(\mu_0, \mu_1)} \sim \mathcal{Q}^t} \Pr[A(v) | \overline{(\mu_0, \mu_1)}]$ and $\Pr[B(v)] = \mathbb{E}_{\overline{(\mu_0, \mu_1)} \sim \mathcal{Q}^t} \Pr[B(v) | \overline{(\mu_0, \mu_1)}]$. We prove by induction on $d(v)$, the depth of a node v , that

$$\Pr[A(v) | \overline{(\mu_0, \mu_1)}] = \Pr[B(v) | \overline{(\mu_0, \mu_1)}] \quad (8)$$

for any $\overline{(\mu_0, \mu_1)}$. This will give the claim.

Towards the aim of showing Eq. (8), fix an arbitrary $\overline{(\mu_0, \mu_1)}$.

Base case: $d(v) = 1$, i.e. v is the root of \mathcal{B} . Thus $\Pr[A(v) | \overline{(\mu_0, \mu_1)}] = \Pr[B(v) | \overline{(\mu_0, \mu_1)}] = 1$.

Inductive step: Assume that $d(v) \geq 2$, and that the statement is true for all vertices of depth at most $d(v) - 1$. Since $d(v) \geq 2$, v is not the root of \mathcal{B} . Let $u = u^{(1)} \times \dots \times u^{(t)}$ be the parent of v , and say variable $x_i^{(j)}$ is queried at u . Without loss of generality we assume that $v = u_0$. We split the proof into the following two cases.

- **Case 1:** $\Pr_{x_i \sim \mu_{z_i}^{(i)}}[x_i^{(j)} = 0 | x_i \in u^{(i)}] \leq \Pr_{x_i \sim \mu_{\bar{z}_i}^{(i)}}[x_i^{(j)} = 0 | x_i \in u^{(i)}]$.

Conditioned on $A(u)$, $\overline{(\mu_0, \mu_1)}$ and $\text{QUERY}_i = 1$, the probability that \mathcal{P} reaches v is $\Pr_{x_i \sim \mu_{z_i}^{(i)}}[x_i^{(j)} = 0 | x_i \in u^{(i)}]$. Also, conditioned on $A(u)$, $\overline{(\mu_0, \mu_1)}$ and $\text{QUERY}_i = 0$ the probability that \mathcal{P} reaches v is exactly equal to the probability that the real number r sampled at u lies in $[0, \Pr_{x_i \sim \mu_{z_i}^{(i)}}[x_i^{(j)} = 0 | x_i \in u^{(i)}]]$, which is equal to $\Pr_{x_i \sim \mu_{z_i}^{(i)}}[x_i^{(j)} = 0 | x_i \in u^{(i)}]$. Thus,

$$\begin{aligned} \Pr[A(v) | \overline{(\mu_0, \mu_1)}] &= \Pr[A(u) | \overline{(\mu_0, \mu_1)}] \cdot \Pr[A(v) | A(u), \overline{(\mu_0, \mu_1)}] \\ &= \Pr[A(u) | \overline{(\mu_0, \mu_1)}] \cdot \Pr_{x_i \sim \mu_{z_i}^{(i)}}[x_i^{(j)} = 0 | x_i \in u^{(i)}]. \end{aligned} \quad (9)$$

Now condition on $B(u)$ and $\overline{(\mu_0, \mu_1)}$. The probability that \mathcal{B} reaches v is exactly equal to the probability that $x_i^{(j)} = 0$ when x is sampled according to the distribution $\gamma_z(\overline{(\mu_0, \mu_1)})$ conditioned on the event that $x \in u$. Note that in the distribution $\gamma_z(\overline{(\mu_0, \mu_1)})$, the x_k 's are independently distributed. Thus,

$$\Pr[B(v) | \overline{(\mu_0, \mu_1)}] = \Pr[B(u) | \overline{(\mu_0, \mu_1)}] \cdot \Pr[B(v) | B(u), \overline{(\mu_0, \mu_1)}]$$

$$= \Pr[B(u) \mid \overline{(\mu_0, \mu_1)}] \cdot \Pr_{x_i \sim \mu_{z_i}^i} [x_i^{(j)} = 0 \mid x_i \in u^{(i)}]. \quad (10)$$

By the inductive hypothesis, $\Pr[A(u) \mid \overline{(\mu_0, \mu_1)}] = \Pr[B(u) \mid \overline{(\mu_0, \mu_1)}]$. It follows from (9) and (10) that $\Pr[A(v) \mid \overline{(\mu_0, \mu_1)}] = \Pr[B(v) \mid \overline{(\mu_0, \mu_1)}]$.

- **Case 2:** $\Pr_{x_i \sim \mu_{z_i}^{(i)}} [x_i^{(j)} = 0 \mid x_i \in u^{(i)}] > \Pr_{x_i \sim \mu_{z_i}^{(i)}} [x_i^{(j)} = 0 \mid x_i \in u^{(i)}]$. Let $v' = u_1$. By an argument similar to Case 1, we have that

$$\Pr[A(v') \mid \overline{(\mu_0, \mu_1)}] = \Pr[B(v') \mid \overline{(\mu_0, \mu_1)}]. \quad (11)$$

Now,

$$\begin{aligned} \Pr[A(v) \mid \overline{(\mu_0, \mu_1)}] &= \Pr[A(u) \mid \overline{(\mu_0, \mu_1)}] - \Pr[A(v') \mid \overline{(\mu_0, \mu_1)}] \\ &= \Pr[B(u) \mid \overline{(\mu_0, \mu_1)}] - \Pr[A(v') \mid \overline{(\mu_0, \mu_1)}] && \text{(By inductive hypothesis)} \\ &= \Pr[B(u) \mid \overline{(\mu_0, \mu_1)}] - \Pr[B(v') \mid \overline{(\mu_0, \mu_1)}] && \text{(By (11))} \\ &= \Pr[B(v) \mid \overline{(\mu_0, \mu_1)}]. \end{aligned}$$

E Proof of Theorem 7

Towards a contradiction, assume that

$$\sum_{i=1}^t \mathbb{E}[\mathcal{N}_i(\mathcal{B}, z, \mathcal{Q})] < t \cdot \min_{\mathcal{C}} \mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}}[\chi(\mathcal{C}, (\mu_0, \mu_1))]. \quad (12)$$

By averaging, there exists a k such that $\mathbb{E}[\mathcal{N}_k(\mathcal{B}, z, \mathcal{Q})] < \min_{\mathcal{C}} \mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}}[\chi(\mathcal{C}, (\mu_0, \mu_1))]$. Let us focus on the expression on the left hand side. Recall that there are two kinds of randomness in this expectation, the choice of the random numbers r in $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ and the choice of $(\mu_0, \mu_1) \sim \mathcal{Q}^t$. We separate out these two as follows:

$$\begin{aligned} \mathbb{E}[\mathcal{N}_k(\mathcal{B}, z, \mathcal{Q})] &= \mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}^t} \mathbb{E}_r[\mathcal{N}_k(\mathcal{B}, z, \overline{(\mu_0, \mu_1)})] \\ &= \mathbb{E}_r \mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}^t}[\mathcal{N}_k(\mathcal{B}, z, \overline{(\mu_0, \mu_1)})] \\ &= \mathbb{E}_r \mathbb{E}_{(\mu_0, \mu_1)^{- (k)} \sim \mathcal{Q}^{t-1}} \mathbb{E}_{(\mu_0^{(k)}, \mu_1^{(k)}) \sim \mathcal{Q}}[\mathcal{N}_k(\mathcal{B}, z, \overline{(\mu_0, \mu_1)})], \end{aligned}$$

where $(\mu_0, \mu_1)^{- (k)}$ is a $t - 1$ -tuple of pairs of distributions without the k^{th} coordinate. This further means that there is a fixing of the randomness r and the $(t - 1)$ -tuple of pairs distributions $(\mu_0, \mu_1)^{- (k)}$ such that $\mathbb{E}_{(\mu_0^{(k)}, \mu_1^{(k)}) \sim \mathcal{Q}}[\mathcal{N}_k(\mathcal{B}, z, \overline{(\mu_0, \mu_1)})] < \min_{\mathcal{C}} \mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}}[\chi(\mathcal{C}, (\mu_0, \mu_1))]$. With such a fixed setting, however, $\mathcal{P}(\mathcal{B}, \mathcal{Q})$ creates a query process equivalent to $\mathcal{P}(\mathcal{B}', \mathcal{Q})$ run on $z_i \in \{0, 1\}$ for a deterministic query algorithm \mathcal{B}' running on inputs from $\{0, 1\}^m$ and such that $(\mathcal{B}', \mathcal{Q})$ is FULL. The distribution $\mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}}[\mathcal{N}_1(\mathcal{B}', 1, \mu_0, \mu_1)]$ is the same as that as $\mathbb{E}_{(\mu_0^{(k)}, \mu_1^{(k)}) \sim \mathcal{Q}}[\mathcal{N}_k(\mathcal{B}, z, ((\mu_0, \mu_1)^{- (k)}, (\mu_0^{(k)}, \mu_1^{(k)})))]$ conditioned on the earlier fixing of $(\mu_0, \mu_1)^{- (k)}$ and the randomness r . Thus $\mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}}[\chi(\mathcal{B}', (\mu_0, \mu_1))] < \min_{\mathcal{C}} \mathbb{E}_{(\mu_0, \mu_1) \sim \mathcal{Q}}[\chi(\mathcal{C}, (\mu_0, \mu_1))]$, a contradiction.

F Proof of Claim 7

Let v_k be the random vertex at which the \mathcal{B} makes its k -th query when it is run on x ; If \mathcal{B} terminates before making k queries, define $v_k := \perp$. Let \mathcal{E} denote the event that in at most $10d^2$ queries, the computation of \mathcal{B} does not reach a vertex v such that $\Pr_{x \sim \mu}[g(x) = 0 \mid x \in$

$v] \cdot \Pr_{x \sim \mu}[g(x) = 1 \mid x \in v] \leq \frac{1}{9}$. Since \mathcal{B} computes g on the supports of μ_0 and μ_1 , therefore if \mathcal{E} happens then the computation of \mathcal{B} does not reach a leaf within $10d^2$ queries. We split the proof into the following two cases.

Case 1: $\Pr[\mathcal{E}] < \frac{3}{4}$.

Condition on the event that the computation reaches a vertex v of \mathcal{B} for which $\Pr_{x \sim \mu}[g(x) = 0 \mid x \in v] \cdot \Pr_{x \sim \mu}[g(x) = 1 \mid x \in v] \leq \frac{1}{9}$ holds. In this case, one of $\Pr_{x \sim \mu}[g(x) = 0 \mid x \in v]$ and $\Pr_{x \sim \mu}[g(x) = 1 \mid x \in v]$ is at most $1/3$. Hence, $|\Pr_{x \sim \mu}[g(x) = 0 \mid x \in v] - \Pr_{x \sim \mu}[g(x) = 1 \mid x \in v]| \geq 1/3$. Let w be the random leaf of the subtree of \mathcal{B}' rooted at v at which the computation ends. The probability that \mathcal{B}' errs is at most

$$\begin{aligned} & \mathbb{E}_w \left[\left| \frac{1}{2} - \frac{1}{2} \left| \Pr_{x \sim \mu}[g(x) = 0 \mid x \in w] - \Pr_{x \sim \mu}[g(x) = 1 \mid x \in w] \right| \right| \right] \\ & \leq \frac{1}{2} - \frac{1}{2} \left| \mathbb{E}_w \left[\Pr_{x \sim \mu}[g(x) = 0 \mid x \in w] \right] - \mathbb{E}_w \left[\Pr_{x \sim \mu}[g(x) = 1 \mid x \in w] \right] \right| \\ & \quad \text{(By Jensen's inequality and linearity of expectation)} \\ & = \frac{1}{2} - \frac{1}{2} \left| \Pr_{x \sim \mu}[g(x) = 0 \mid x \in v] - \Pr_{x \sim \mu}[g(x) = 1 \mid x \in v] \right| \leq \frac{1}{3}. \end{aligned}$$

Thus we have shown that conditioned on $\bar{\mathcal{E}}$ the probability that \mathcal{B}' errs is at most $\frac{1}{3}$. Hence, the probability that \mathcal{B}' errs is at most $\frac{1}{4} \cdot \frac{1}{3} + \frac{3}{4} \cdot \frac{1}{2} = \frac{11}{24} < \frac{1}{2}$.

Case 2: $\Pr[\mathcal{E}] \geq \frac{3}{4}$.

Let $a_j := (i_j, x_{i_j})$ be the tuple formed by the index and value of the random input variable queried at the j -th step by \mathcal{B}' ; if \mathcal{B}' terminates before making j queries (i.e. $v_j = \perp$) or v_j is a leaf of \mathcal{B} , then define $i_j, x_{i_j} := \perp$. Note that the sequence (a_1, \dots, a_{10d^2}) uniquely specifies a leaf of \mathcal{B}' , and vice versa. Let $I(\cdot, \cdot)$ denote the *mutual information*. (See Appendix C for the definitions and results from information theory used in this work). We prove the following claim in Appendix F.1.

Claim 14. If $\Pr[\mathcal{E}] \geq \frac{3}{4}$, then $I(a_1, \dots, a_{10d^2} : g(x)) \geq \frac{1}{40}$.

Thus if $\Pr[\mathcal{E}] \geq \frac{3}{4}$, Claim 14 implies that

$$\mathbf{H}(g(x) \mid a_1, \dots, a_{10d^2}) \leq 1 - \frac{1}{40} = \frac{39}{40}. \quad (13)$$

Let \mathcal{L} be the set of leaves ℓ of \mathcal{B}' such that $\mathbf{H}(g(x) \mid x \in \ell) \leq \frac{79}{80}$. For each $\ell \in \mathcal{L}$, $\min_b \Pr_{x \sim \mu}[g(x) = b \mid x \in \ell] \leq \frac{9}{20}$. Conditioned on $(a_1, \dots, a_{10d^2}) \in \mathcal{L}$, the probability that \mathcal{B}' errs is at most $\frac{9}{20}$. By *Markov's inequality* and (13), it follows that $\Pr[(a_1, \dots, a_{10d^2}) \in \mathcal{L}] \geq \frac{1}{79}$. Thus \mathcal{B}' errs with probability at most $\frac{1}{79} \cdot \frac{9}{20} + \frac{18}{19} \cdot \frac{1}{2} < \frac{1}{2}$.

F.1 Proof of Claim 14

Let v be a vertex in \mathcal{B} . Define $\Delta(v)$ as follows.¹

$$\Delta(v) := \begin{cases} \left| \Pr_{x \sim \mu_0}[x_i = 0 \mid x \in v] - \Pr_{x \sim \mu_1}[x_i = 0 \mid x \in v] \right| & \text{if } v \neq \perp \text{ and } \Pr_{x \sim \mu_b}[x \in v] > 0 \\ & \text{for } b \in \{0, 1\}, \\ 1 & \text{otherwise.} \end{cases}$$

The following claim shows that if $\Delta(v)$ is large, the query outcome of v contains significant information about $g(x)$.

¹Recall that we mentioned $\Delta(v)$ in Section 1.1.

Claim 15. Let v be a vertex in \mathcal{B} . Let variable x_i be queried at v . Then,

$$I(g(x) : x_i \mid x \in v) \geq 8 \left(\Pr_{x \sim \mu} [g(x) = 0 \mid x \in v] \cdot \Pr_{x \sim \mu} [g(x) = 1 \mid x \in v] \cdot \Delta(v) \right)^2.$$

Proof. Define $b := g(x)$. Condition on the event $x \in v$. Recall from Appendix C that $(b \otimes x_i)$ be the distribution over pairs of bits, where the bits are distributed independently according to the distributions of b and x_i respectively. Fact 12 implies that $I(b : x_i) = D((b, x_i) \parallel (b \otimes x_i))^2$. Now, *Pinsker's inequality* (Theorem 13) implies that

$$D((b, x_i) \parallel (b \otimes x_i)) \geq \frac{1}{2} \|(b, x_i) - (b \otimes x_i)\|_1^2. \quad (14)$$

Next, we bound $\|(b, x_i) - (b \otimes x_i)\|_1$. To this end, we fix bits $z_1, z_2 \in \{0, 1\}$, and bound $|\Pr[(b, x_i) = (z_1, z_2)] - \Pr[(b \otimes x_i) = (z_1, z_2)]|$. We have that,

$$\Pr[(b, x_i) = (z_1, z_2)] = \Pr[b = z_1] \Pr[x_i = z_2 \mid b = z_1]. \quad (15)$$

Now,

$$\begin{aligned} \Pr[(b \otimes x_i) = (z_1, z_2)] &= \Pr[b = z_1] \Pr[x_i = z_2] \\ &= \Pr[b = z_1] (\Pr[b = z_1] \Pr[x_i = z_2 \mid b = z_1] + \\ &\quad \Pr[b = \bar{z}_1] \Pr[x_i = z_2 \mid b = \bar{z}_1]). \end{aligned} \quad (16)$$

Taking the absolute difference of (16) and (15) we have that,

$$\begin{aligned} &|\Pr[(b, x_i) = (z_1, z_2)] - \Pr[(b \otimes x_i) = (z_1, z_2)]| \\ &= \Pr[b = z_1] \cdot \Pr[b = \bar{z}_1] \cdot \Delta(v) = \Pr[b = 0] \cdot \Pr[b = 1] \cdot \Delta(v) \end{aligned} \quad (17)$$

The Claim follows by adding (17) over z_1, z_2 and using (14). \blacksquare

Let \mathcal{B} be run on a random input x sampled from μ . The next claim proves a lower bound on the expected sum of $\Delta(v)$ for the random vertices v in the transcript of \mathcal{B} . Recall from Appendix F that v_k is the random vertex at which the k -th query is made; If \mathcal{B} terminates before making k queries, define $v_k := \perp$.

Claim 16. Let c be any positive integer. Then,

$$\sum_{k=1}^{10dc} \mathbb{E}[\Delta(v_k) \mid \mathcal{E}] \geq \frac{13c}{20}.$$

To prove Claim 16 we need the following claim.

Claim 17.

$$\sum_{k=1}^{10d} \mathbb{E}[\Delta(v_k) \mid \mathcal{E}] \geq \frac{13}{20}.$$

Note that if \mathcal{B} terminates before making t queries, $v_t = \perp$ and $\Delta(v_t) = 1$.

Proof of Claim 17. Let us sample vertices u_k of \mathcal{B} as follows:

1. Set $z = \begin{cases} 1 & \text{with probability } \Pr_{x \sim \mu}[g(x) = 1], \\ 0 & \text{with probability } \Pr_{x \sim \mu}[g(x) = 0] \end{cases}$
2. Run $\mathcal{P}(\mathcal{B}, \mu_0, \mu_1)$ on the 1-bit input z .

²See Appendix C for definition of Kullback-Leibler divergence and L₁-distance.

3. Let u_k be the vertex v of \mathcal{B} in the beginning of the k -th iteration of the *while* loop of Algorithm 2. If the simulation stops before k iterations, set $u_k := \perp$. Return $(u_k)_{k=1, \dots}$.

By Theorem 5, the transcripts $(u_k)_{k=1, 2, \dots}$ and $(v_k)_{k=1, 2, \dots}$ have the same distribution.

Now, since $\mathbb{E}[\mathcal{N}_1] \leq \chi(g) = d$, we have by *Markov's inequality* that the probability that $\mathcal{P}(\mathcal{B}, \mu_0, \mu_1)$ sets QUERY₁ to 1 within first $10d$ iterations of the *while* loop, is at least $9/10$. Note that conditioned on the event that the computation of $\mathcal{P}(\mathcal{B}, \mu_0, \mu_1)$ is at vertex v of \mathcal{B} that queries the input bit x_i , the probability that the random real number r generated in the same iteration lies in the interval $[\min_b \Pr_{x_i \sim \mu_b}[x_i = 0 \mid x \in v], \max_b \Pr_{x_i \sim \mu_b}[x_i = 0 \mid x \in v]]$ is exactly $\Delta(v)$. We have,

$$\begin{aligned} \sum_{k=1}^{10d} \mathbb{E}[\Delta(v_k) \mid \mathcal{E}] &= \sum_{k=1}^{10d} \mathbb{E}[\Delta(u_k) \mid \mathcal{E}] \\ &\geq \Pr[\text{QUERY}_1 \text{ is set to to 1 within first } 10d \text{ iterations} \mid \mathcal{E}] \quad (\text{by union bound}) \\ &\geq \Pr[\text{QUERY}_1 \text{ is set to to 1 within first } 10d \text{ iterations}] - \Pr[\bar{\mathcal{E}}] \\ &\geq \frac{9}{10} - \frac{1}{4} = \frac{13}{20}. \end{aligned}$$

■

The following observation will be useful in the proof of Claim 16.

Observation 18. Let v be any node of \mathcal{B} , such that the associated subcube has non-empty intersections with the supports of both μ_0 and μ_1 . Let $\mu'_0 := \mu_0 \mid v$ and $\mu'_1 := \mu_1 \mid v$. Let \mathcal{B}_v denote the subtree of \mathcal{B} rooted at v . Then \mathcal{B}_v is an optimal decision tree for μ'_0 and μ'_1 .

Proof. If \mathcal{B}_v is not an optimal decision tree for μ'_0 and μ'_1 then we could replace it by an optimal decision tree for μ'_0 and μ'_1 , and for the resultant decision tree \mathcal{B}' , the expected value of \mathcal{N}_1 in $\mathcal{P}(\mathcal{B}', \mu_0, \mu_1)$ will be smaller than that in $\mathcal{P}(\mathcal{B}, \mu_0, \mu_1)$. This will contradict the optimality of \mathcal{B} . ■

Proof of Claim 16. For $i = 0, \dots, c-1$, let w be any vertex at depth $10id + 1$ consistent with \mathcal{E} , such that $\Pr_{x \sim \mu_0}[x \in w], \Pr_{x \sim \mu_1}[x \in w] \neq 0$. Consider the subtree \mathbb{T} of \mathcal{B} rooted at w . Let $w_1 := w$ and w_ℓ be the random vertex at depth ℓ of \mathbb{T} , when \mathbb{T} is run on a random input from $\mu \mid w$, or \perp if \mathbb{T} terminates before ℓ queries. By Observation 18, \mathbb{T} is an optimal decision tree for distributions $\mu'_0 := \mu_0 \mid w, \mu'_1 := \mu_1 \mid w$. From Claim 17 we have that,

$$\sum_{\ell=1}^{10d} \mathbb{E}[\Delta(w_\ell) \mid \mathcal{E}] \geq \frac{13}{20}. \quad (18)$$

Where $\Delta(w_\ell)$ is with respect to distributions μ'_0 and μ'_1 . Since $\mu'_0 \mid w_\ell = \mu_0 \mid w_\ell$ and $\mu'_1 \mid w_\ell = \mu_1 \mid w_\ell$, $\Delta(w_\ell)$ in (18) is also with respect to distributions μ_0 and μ_1 . Now, when w is the random vertex v_{10id+1} , w_ℓ is the random vertex $v_{10id+\ell}$. Thus from (18) we have that,

$$\sum_{k=10id+1}^{10(i+1)d} \mathbb{E}[\Delta(v_k) \mid \mathcal{E}] \geq \frac{13}{20}. \quad (19)$$

The claim follows by adding (19) over $i = 0, \dots, c-1$. ■

Now we are ready to prove Claim 14. By setting $c = d$ and invoking Claim 16 we have,

$$\sum_{t=1}^{10d^2} \mathbb{E}[\Delta(v_t) \mid \mathcal{E}] \geq \frac{13d}{20}. \quad (20)$$

Let E_j be the event $\Pr_{x \sim \mu_0}[x \in v_j] \neq 0 \wedge \Pr_{x \sim \mu_1}[x \in v_j] \neq 0 \wedge v_j \neq \perp$ (i.e. v_j is a vertex of \mathcal{B} and is not a leaf). Note that if $v_j \neq \perp$ and v_j is not a leaf of \mathcal{B} , v_j is determined by (a_1, \dots, a_{j-1}) and vice versa, and hence $I(a_j : g(x) \mid a_1, \dots, a_{j-1}) = I(x_{i_j} : g(x) \mid v_j)$. If $v_j = \perp$ or v_j is a leaf of \mathcal{B} , then $g(x)$ is determined by (a_1, \dots, a_{i-1}) , and $x_{i_j} = \perp$; thus, $I(a_j : g(x) \mid a_1, \dots, a_{j-1}) = I(x_{i_j} : g(x) \mid v_j) = 0$. Thus we have,

$$\begin{aligned}
& I(a_1, \dots, a_{10d^2} : g(x)) \\
&= \sum_{j=1}^{10d^2} I(a_j : g(x) \mid a_1, \dots, a_{j-1}) \quad (\text{By the chain rule of mutual information (Theorem 11)}) \\
&= \sum_{j=1}^{10d^2} I(x_{i_j} : g(x) \mid v_j) \quad (\text{From the discussion above}) \\
&\geq 8 \sum_{j=1}^{10d^2} \mathbb{E}[\mathbf{1}_{E_j} \cdot [\Pr[g(x) = 0 \mid x \in v_j] \cdot \Pr[g(x) = 1 \mid x \in v_j] \cdot \Delta(v_j)]^2] \\
&\quad (\text{From Claim 15}) \\
&\geq 8 \sum_{j=1}^{10d^2} \Pr[\mathcal{E}] \cdot \mathbb{E}[[\Pr[g(x) = 0 \mid x \in v_j] \cdot \Pr[g(x) = 1 \mid x \in v_j] \cdot \Delta(v_j)]^2 \mid \mathcal{E}] \\
&\quad (\text{Conditioned on } \mathcal{E}, E_j \text{ happens with probability 1 for each } j \leq 10d^2) \\
&\geq 8 \sum_{j=1}^{10d^2} \frac{3}{4} \cdot \frac{1}{9} \cdot \mathbb{E}[\Delta(v_j)^2 \mid \mathcal{E}] \quad (\text{By the assumption } \Pr[\mathcal{E}] \geq \frac{3}{4}) \\
&= \frac{2}{3} \sum_{j=1}^{10d^2} \mathbb{E}[\Delta(v_j)^2 \mid \mathcal{E}] \\
&\geq \frac{2}{3} \sum_{j=1}^{10d^2} (\mathbb{E}[\Delta(v_j) \mid \mathcal{E}])^2 \quad (\text{By Jensen's inequality}) \\
&\geq \frac{2}{3} \cdot \frac{1}{10d^2} \left(\sum_{j=1}^{10d^2} \mathbb{E}[\Delta(v_j) \mid \mathcal{E}] \right)^2 \quad (\text{By Cauchy-Schwarz inequality}) \\
&\geq \frac{1}{40}. \quad (\text{From (20)}) \tag{21}
\end{aligned}$$

G Proof of Claim 9

Proof. (i) Assume that a deterministic protocol of cost k solves f_0 with respect to the uniform input distribution with error at most $4/9$. Such a protocol partitions $\{0, 1\}^n$ into (at most) 2^k sub-cubes, each marked by some “answer” (an element from $\{0, 1\}^n$). In particular, more than $2^n - 2^{n-4}$ points belong to sub-cubes of size at least 2^{n-k-4} – in other words, to sub-cubes of co-dimension at most $k+4$. As more than $15/16$ fraction of all points belong to such sub-cubes and the total protocol error is at most $4/9$, there exists at least one sub-cube of co-dimension $k+4$, on which the protocol errs with probability less than $\frac{4}{9} \cdot \frac{16}{15} < \frac{1}{2}$.

The symmetry in the definition of f_0 allows us to assume without loss of generality that the sub-cube is the set $\tau \stackrel{\text{def}}{=} 0^{k+4} \circ \{0, 1\}^{n-k-4}$, where “ \circ ” denotes string concatenation. It is easy to see that the “answer” that would minimize the error probability with respect to this sub-cube can be any binary string starting with “ 0^{k+4} ”, so let us assume that the

actual label is 0^n . Let \mathcal{U}_ℓ denote uniform distribution on $\{0, 1\}^n$. Then

$$\Pr[\text{error} | Z \in \tau] = \Pr_{Z' \sim \mathcal{U}_{n-k-4}} \left[|Z'| \leq \frac{n}{2} - \sqrt{n} \right] < \frac{1}{2},$$

which implies that $k+4 \geq 2\sqrt{n}$, as a uniformly-random binary string of length more than $n - 2\sqrt{n}$ would have more than $n/2 - \sqrt{n}$ “ones” with probability at least $1/2$.

- (ii) A randomized query protocol of cost k and error $1/3$ for g_0 would trivially imply existence of a randomized communication protocol of cost at most $2k$ and error $1/3$ for the bipartite problem *Gap-Hamming-Distance*:

$$\text{GHD}(X, Y) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } |X \oplus Y| \leq n/2 - \sqrt{n}; \\ 1 & \text{if } |X \oplus Y| \geq n/2 + \sqrt{n}; \\ * & \text{otherwise,} \end{cases}$$

and it has been demonstrated by Chakrabarti and Regev [CR11] that the complexity of this problem for any constant error is $\Omega(n)$.

- (iii) Consider the following protocol for computing $f_0(g_0(x_1), \dots, g_0(x_n))$, where $x_i \in \{0, 1\}^n$: For every $i \in [n]$, let $a_i = x_i(j_i)$, where j_i is chosen uniformly at random from $\{1, \dots, n\}$ – that is, a_i is a uniformly-random bit of x_i . Then $|\{i | a_i = g_0(x_i)\}|$ – the expected number of “correctly guessed” a_i -s is at least $n/2 + \sqrt{n}$; intuitively, this means that the probability that a_1, \dots, a_n is a right answer to $f_0(g_0(x_1), \dots, g_0(x_n))$ is “non-trivially high” – to “boost” this probability, we will use several “probes” from every x_i and take their majority vote.

Protocol: For an odd integer t_ε as defined next, independently choose $j_{i,k} \in [n]$ for $i \in [n]$ and $k \in [t_\varepsilon]$. Let $a_i \stackrel{\text{def}}{=} \text{maj}(x_i(j_{i,1}), \dots, x_i(j_{i,t_\varepsilon}))$ and output “ a_1, \dots, a_n ”.

To analyse it, we consider for every $i \in [n]$:

$$\begin{aligned} & \Pr[a_i = g_0(x_i)] - \Pr[a_i \neq g_0(x_i)] \\ & \geq \sum_{i=0}^{\frac{t_\varepsilon-1}{2}} \binom{t_\varepsilon}{i} \cdot \left(\left(\frac{1}{2} - \frac{1}{\sqrt{n}} \right)^i \left(\frac{1}{2} + \frac{1}{\sqrt{n}} \right)^{t_\varepsilon-i} - \left(\frac{1}{2} - \frac{1}{\sqrt{n}} \right)^{\frac{t_\varepsilon+1}{2}+i} \left(\frac{1}{2} + \frac{1}{\sqrt{n}} \right)^{\frac{t_\varepsilon-1}{2}-i} \right) \\ & = \left(1 - \left(\frac{1-2/\sqrt{n}}{1+2/\sqrt{n}} \right)^{\frac{t_\varepsilon+1}{2}} \right) \cdot \sum_{i=0}^{\frac{t_\varepsilon-1}{2}} \binom{t_\varepsilon}{i} \cdot \left(\frac{1}{2} - \frac{1}{\sqrt{n}} \right)^i \left(\frac{1}{2} + \frac{1}{\sqrt{n}} \right)^{t_\varepsilon-i}, \end{aligned}$$

where the equality occurs when $|x_i| - n/2 = \pm\sqrt{n}$. As

$$\sum_{i=0}^{\frac{t_\varepsilon-1}{2}} \binom{t_\varepsilon}{i} \cdot \left(\frac{1}{2} - \frac{1}{\sqrt{n}} \right)^i \left(\frac{1}{2} + \frac{1}{\sqrt{n}} \right)^{t_\varepsilon-i} = \Pr \left[a_i = g_0(x_i) \mid |x_i| - \frac{n}{2} = \pm\sqrt{n} \right] > \frac{1}{2},$$

we get

$$\begin{aligned} & \Pr[a_i = g_0(x_i)] - \Pr[a_i \neq g_0(x_i)] \\ & > \frac{1}{2} \cdot \left(1 - \left(\frac{1-2/\sqrt{n}}{1+2/\sqrt{n}} \right)^{\frac{t_\varepsilon+1}{2}} \right) > \frac{1}{2} \cdot \left(1 - \left(1 - \frac{2}{\sqrt{n}} \right)^{t_\varepsilon/2} \right) \geq \min \left\{ \frac{t_\varepsilon}{4\sqrt{n}}, \frac{1}{4} \right\}. \end{aligned}$$

Our t_ε will be small enough to guarantee that $\frac{t_\varepsilon}{4\sqrt{n}} \leq \frac{1}{4}$, so we can write

$$\Pr[a_i = g_0(x_i)] > \frac{1}{2} + \frac{t_\varepsilon}{8\sqrt{n}}. \quad (22)$$

Now let us estimate the probability that a_1, \dots, a_n is a wrong answer to $f_0(g_0(x_1), \dots, g_0(x_n))$: This occurs only if $|\{i|a_i = g_0(x_i)\}| < n/2 + \sqrt{n}$, so by the Chernoff bound (in a form given in [DM05]),

$$\Pr[\text{the protocol errs}] < \exp\left(-\frac{1}{2} \cdot \left(\frac{t_\varepsilon}{8} - 1\right)^2\right),$$

so that choosing $t_\varepsilon \in \Theta\left(\sqrt{\log(1/\varepsilon)}\right)$ would suffice for our needs and the result follows. ■

References

- [Aar08] Scott Aaronson. Quantum certificate complexity. *Journal of Computer and System Sciences*, 74(3):313–322, 2008.
- [AGJ⁺17] Anurag Anshu, Dmitry Gavinsky, Rahul Jain, Srijita Kundu, Troy Lee, Priyanka Mukhopadhyay, Miklos Santha, and Swagato Sanyal. A composition theorem for randomized query complexity. In *37th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2017, December 11-15, 2017, Kanpur, India*, pages 10:1–10:13, 2017.
- [BBCR13] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. *SIAM J. Comput.*, 42(3):1327–1363, 2013.
- [BK16] Shalev Ben-David and Robin Kothari. Randomized query complexity of sabotaged and composed functions. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, pages 60:1–60:14, 2016.
- [CR11] A. Chakrabarti and O. Regev. An Optimal Lower Bound on the Communication Complexity of Gap-Hamming-Distance. *Proceedings of the 43rd Symposium on Theory of Computing*, pages 51–60, 2011.
- [DM05] E. Drukh and Y. Mansour. Concentration Bounds for Unigram Language Models. *Journal of Machine Learning Research* 6, pages 1231–1264, 2005.
- [GLS18] Dmitry Gavinsky, Troy Lee, and Miklos Santha. On the randomised query complexity of composition. Technical Report arXiv:1801.02226, arXiv, 2018.
- [GSS16] Justin Gilmer, Michael Saks, and Srikanth Srinivasan. Composition limits and separating examples for some Boolean function complexity measures. *Combinatorica*, 36(3):265–311, 2016.
- [HLS07] Peter Høyer, Troy Lee, and Robert Spalek. Negative weights make adversaries stronger. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 526–535, 2007.
- [Li18] Yaqiao Li. Conflict complexity is lower bounded by block sensitivity. Technical Report arXiv:1810.08873, arXiv, 2018.
- [Mon14] Ashley Montanaro. A composition theorem for decision tree complexity. *Chicago J. Theor. Comput. Sci.*, 2014, 2014.
- [Rei11] Ben Reichardt. Reflections for quantum query algorithms. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 560–569, 2011.

- [San18] Swagato Sanyal. A composition theorem via conflict complexity. Technical Report arXiv:1801.03285, arXiv, 2018.
- [She12] Alexander A. Sherstov. Making polynomials robust to noise. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 747–758, 2012.
- [Tal13] Avishay Tal. Properties and applications of boolean function composition. In *Innovations in Theoretical Computer Science, ITCS '13*, pages 441–454, 2013.
- [von28] John von Neumann. Zur theorie der gessellschaftsspiele. *Math. Ann.*, 100:295–320, 1928.