

# Randomised Composition and Small-Bias Minimax

Shalev Ben-David  
University of Waterloo

Eric Blais  
University of Waterloo

Mika Göös  
EPFL

Gilbert Maystre  
EPFL

**Abstract**—We prove<sup>1</sup> two results about randomised query complexity  $\mathbf{R}(f)$ . First, we introduce a *linearised complexity measure*  $\mathbf{LR}$  and show that it satisfies an *inner-optimal composition theorem*:  $\mathbf{R}(f \circ g) \geq \mathbf{LR}(f)\mathbf{LR}(g)$  for all partial  $f$  and  $g$ , and moreover,  $\mathbf{LR}$  is the largest possible measure with this property. In particular,  $\mathbf{LR}$  can be polynomially larger than previous measures that satisfy an inner composition theorem, such as the max-conflict complexity of Gavinsky, Lee, Santha, and Sanyal (ICALP 2019).

Our second result addresses a question of Yao (FOCS 1977). He asked if  $\epsilon$ -error *expected query complexity*  $\bar{\mathbf{R}}(f)$  admits a *distributional characterisation* relative to some hard input distribution. Vereshchagin (TCS 1998) answered this question affirmatively in the bounded-error case. We show that an analogous theorem *fails* in the small-bias case  $\epsilon = 1/2 - o(1)$ .

## I. INTRODUCTION

This paper is motivated by the following basic open problem in boolean function complexity theory.

**Conjecture 1.**  $\mathbf{R}(f \circ g) \geq \Omega(\mathbf{R}(f)\mathbf{R}(g))$  for all total boolean functions  $f, g$ .

Let us unpack what this conjecture is claiming. The randomised  $\epsilon$ -error query complexity  $\mathbf{R}_\epsilon(f)$  of a boolean function  $f: \{0, 1\}^n \rightarrow \{0, 1\}$  is defined (see [BdW02] for the classic reference) as the least number of queries a randomised algorithm (decision tree) needs to make, on the worst-case input, to the bits  $x_i$  of  $x \in \{0, 1\}^n$  in order to compute  $f(x)$  correctly with error at most  $\epsilon$ . We write  $\mathbf{R} := \mathbf{R}_{1/3}$  for the bounded-error case. For functions  $f$  and  $g$  over  $n$  and  $m$  bits, their composition  $f \circ g$  is defined over  $nm$  bits by

$$(f \circ g)(x) := f(g(x^1), \dots, g(x^n))$$

where  $x = (x^1, \dots, x^n) \in (\{0, 1\}^m)^n$ . In particular, we have  $\mathbf{R}(f \circ g) \leq O(\mathbf{R}(f)\mathbf{R}(g) \log \mathbf{R}(f))$  for all  $f, g$ . This holds since we can run an algorithm for  $f$  with query cost  $\mathbf{R}(f)$  and whenever it queries an input bit, we can run, as a subroutine, an  $\epsilon$ -error algorithm for  $g$  of cost  $\mathbf{R}_\epsilon(g)$ . Setting  $\epsilon \ll 1/\mathbf{R}(f)$  makes sure that the errors made by the subroutines do not add up. Moreover, we have  $\mathbf{R}_\epsilon(g) \leq O(\mathbf{R}(g) \log(1/\epsilon)) = O(\mathbf{R}(g) \log \mathbf{R}(f))$  by standard error reduction techniques. **Conjecture 1** thus postulates that a converse inequality always holds (without the log factor).

The analogue of **Conjecture 1** has been long resolved for many other well-studied complexity measures: deterministic query complexity satisfies a perfect multiplicative composition theorem,  $\mathbf{D}(f \circ g) = \mathbf{D}(f)\mathbf{D}(g)$  [Sav02], quantum query

<sup>1</sup>This is an extended abstract. For the full version of this article, please refer to [BDBGM22].

complexity satisfies  $\mathbf{Q}(f \circ g) = \Theta(\mathbf{Q}(f)\mathbf{Q}(g))$  [Rei11], [LMR<sup>+</sup>11], and yet more examples (degree, certificate complexity, sensitivity, rank) are discussed in [Tal13], [GSS16], [DM21]. In the randomised case, however, the conjecture has proved more delicate, exhibiting a far richer, and more surprising, structure.

**Partial counterexamples.** **Conjecture 1** is known to be false if we relax the requirement that  $f, g$  are total and instead consider *partial* functions (promise problems), which are undefined on some inputs  $x$ ,  $f(x) = *$ . Indeed, works by Gavinsky, Lee, Santha, and Sanyal [GLSS19] and Ben-David and Blais [BB20b] have culminated in examples of partial functions  $f, g$  such that  $\mathbf{R}(f \circ g) \leq o(\mathbf{R}(f)\mathbf{R}(g))$ . Motivated by these counterexamples, we ask: *What is the best possible composition theorem one can prove for partial functions?*

### A. A new composition theorem

Our first result is an *inner-optimal* composition theorem for partial functions. To state this result, we start by introducing a new *linearised complexity measure* defined for a partial function  $f: \{0, 1\}^n \rightarrow \{0, 1, *\}$  by

$$\mathbf{LR}(f) := \min_R \max_x \frac{\text{cost}(R, x)}{\text{bias}_f(R, x)},$$

- where  $R$  ranges over randomised decision trees;
- $x$  ranges over the domain of  $f$ , namely,  $\text{Dom}(f) := f^{-1}(\{0, 1\})$ ;
- $\text{cost}(R, x)$  denotes the *expected* number of queries  $R$  makes on input  $x$ ; and
- $\text{bias}_f(R, x)$  denotes the bias  $R$  has of guessing the value  $f(x)$  correctly; formally,  $\text{bias}_f(R, x) := \max\{1 - 2 \text{err}_f(R, x), 0\}$  where  $\text{err}_f(R, x) := \Pr_R[R(x) \neq f(x)]$ . We often omit the subscript  $f$  for brevity.

This definition might seem mysterious at first sight. To get better acquainted with it, let us first note that

$$\forall f: \quad \Omega(\sqrt{\mathbf{R}(f)}) \leq \mathbf{LR}(f) \leq O(\mathbf{R}(f)). \quad (1)$$

Indeed, the second inequality follows by considering a bounded-error decision tree  $R$ , with  $\text{cost}(R, x) \leq \mathbf{R}(f)$  and  $\text{bias}(R, x) \geq 1/3$ . For the first inequality, if we let  $R$  be a randomised tree that achieves the minimum in the definition of  $\mathbf{LR}(f)$ , we can amplify the bias of  $R$ , which is possibly tiny, as follows. On input  $x$  we run  $R(x)$  repeatedly until we have made a total of  $\mathbf{LR}(f)^2$  queries, and then output the majority answer over all runs. We expect this simulation to run  $R(x)$  for  $\mathbf{LR}(f)^2 / \text{cost}(R, x) \geq 1/\text{bias}(R, x)^2$  many times, which,

by standard Chernoff bounds, is enough to amplify the bias to a constant. This shows  $R(f) \leq O(\text{LR}(f)^2)$ .

Both extremes in (1) can be realised. First, consider the  $n$ -bit parity function  $\text{XOR}_n$ . It is not hard to see that any randomised tree that achieves bias  $\delta$  for  $\text{XOR}_n$  needs to query all the  $n$  bits with probability at least  $\delta$ , resulting in expected query cost at least  $\delta n$ . This shows  $\text{LR}(\text{XOR}_n) = R(\text{XOR}_n) = n$ . Second, consider the partial  $n$ -bit *gap-majority* function (here  $|x|$  denotes the Hamming weight)

$$\text{GAPMAJ}_n(x) := \begin{cases} 1 & \text{if } |x| \geq n/2 + \sqrt{n}, \\ 0 & \text{if } |x| \leq n/2 - \sqrt{n}, \\ * & \text{otherwise.} \end{cases}$$

It is well known that  $R(\text{GAPMAJ}_n) = \Theta(n)$ . By contrast, the algorithm  $R$  that queries and outputs a uniform random bit of  $x$  has  $\text{cost}(R, x) = 1$  and  $\text{bias}(R, x) \geq \Omega(1/\sqrt{n})$ , which shows  $\text{LR}(\text{GAPMAJ}_n) \leq O(\sqrt{n})$ .

Our first main result shows that a multiplicative composition theorem holds when the inner function is measured according to LR, and moreover, our choice of LR is optimal among all inner complexity measures. Ultimately, these theorems are what lends naturalness to our definition of LR.

**Theorem 1.**  $R(f \circ g) \geq \Omega(R(f)\text{LR}(g))$  for all partial boolean functions  $f, g$ .

**Theorem 2.** *Theorem 1 is optimal: If  $M$  is any complexity measure such that  $R(f \circ g) \geq \Omega(R(f)M(g))$  for all partial  $f, g$ , then  $\text{LR}(g) \geq \Omega(M(g))$  for all partial  $g$ .*

Additionally, LR itself satisfies a composition theorem as well.

**Theorem 3.**  $\text{LR}(f \circ g) \geq \Omega(\text{LR}(f)\text{LR}(g))$  for all partial boolean functions  $f, g$ .

### B. Comparison with previous work

The randomised composition conjecture for general boolean functions was first explicitly raised in [BK16]. Several complexity measures have since been shown to satisfy an inner composition theorem, including:

- 1) (block-)sensitivity  $\mathfrak{s}$ ,  $\text{bs}$  [ABK16],
- 2) randomised sabotage complexity  $\text{RS}$  [BK16],
- 3) randomised complexity  $R_\delta$  with small-bias error  $\delta := 1/2 - 1/n^4$  [AGJ<sup>+</sup>18],
- 4) max-conflict complexity  $\bar{\chi}$  [GLSS19] (also studied in [Li21]).

By our optimality theorem, we have  $\text{LR}(f) \geq \Omega(M(f))$  for all  $M \in \{\mathfrak{s}, \text{bs}, \text{RS}, R_\delta, \bar{\chi}\}$  and all  $f$ . In fact, we can show that the largest of the above measures, namely  $\bar{\chi}$ , can sometimes be polynomially smaller than LR.<sup>2</sup>

<sup>2</sup>Technically, it does not seem to be known in the literature whether  $R$  is always at most  $\bar{\chi}$ ; this doesn't matter much for our purposes, as LR is larger than both and it is easy to separate LR from  $R$  (for example with the OR function).

**Lemma 4.** *There exists a partial  $f$  such that  $\text{LR}(f) \geq \Omega(\bar{\chi}(f)^{1.5})$ .*

Previous work has also investigated complexity measures  $M$  that admit an *outer* composition theorem, that is,  $R(f \circ g) \geq \Omega(M(f)R(g))$  for all partial  $f, g$ . These measures include:

- 1) sensitivity  $\mathfrak{s}$  [GJPW18] (which was applied in [AKK16]),
- 2) fractional block sensitivity  $\text{fbs}$  [BDG<sup>+</sup>20],
- 3) noisy randomised complexity  $\text{noisyR}$  [BB20b] (also studied in [GTW21]).

In particular,  $\text{noisyR}$  is known to be *outer-optimal*: if we have  $R(f \circ g) \geq \Omega(M(f)R(g))$  for all partial  $f, g$ , then  $\text{noisyR}(f) \geq \Omega(M(f))$  for all partial  $f$ . Our result can be viewed as an inner analogue of this.

Finally, we mention that randomised composition has also been studied in the *super-multiplicative* regime, where we have examples of functions  $f, g$  with  $R(f \circ g) \geq \omega(R(f)R(g))$ . Tight bounds exist when the outer function is identity [BB19] (building on [JKS10], [BK16]), parity [BKLS20], or majority [BGKW20], [GM21].

### C. On small-bias minimax

Our second result addresses a question of Yao [Yao77]. Yao-style minimax theorems are routinely used to construct and analyse hard input distributions (including in our proof of the new composition theorem). For example,  $R_\epsilon$  admits a distributional characterisation as

$$R_\epsilon(f) = \max_{\mu} \min_{R \in \mathcal{R}(f, \epsilon, \mu)} \text{depth}(R), \quad (2)$$

where  $\mu$  ranges over distributions on  $\text{Dom}(f)$ ; the set  $\mathcal{R}(f, \epsilon, \mu)$  consists of trees  $R$  with  $\mathbb{E}_{x \sim \mu}[\text{err}(R, x)] \leq \epsilon$ ; and  $\text{depth}(R)$  is the worst-case cost of  $R$ , that is, maximum number of queries over all inputs (and internal randomness if  $R$  is randomised). While the worst-case cost setting is perhaps what is most widely studied up to this day, Yao's original paper discussed, in fact, exclusively the expected cost setting. It is the expected cost setting that is currently undergoing a renaissance as it has proven important in the randomised composition literature surveyed above (Section I-B).

**Minimax for expected cost.** We define the  $\epsilon$ -error expected query complexity and the  $\epsilon$ -error distributional expected query complexity by

$$\begin{aligned} \bar{R}_\epsilon(f) &:= \min_{R \in \mathcal{R}(f, \epsilon)} \max_x \text{cost}(R, x), \\ \bar{D}_\epsilon(f) &:= \max_{\mu} \min_{R \in \mathcal{R}(f, \epsilon, \mu)} \text{cost}(R, \mu), \end{aligned}$$

where  $\mathcal{R}(f, \epsilon)$  is the set of randomised trees  $R$  such that  $\text{err}(R, x) \leq \epsilon$  for all inputs  $x$ ; and  $\text{cost}(R, \mu) := \mathbb{E}_{x \sim \mu}[\text{cost}(R, x)]$  is the expected cost over  $\mu$  (and internal randomness of  $R$ ). We note that the set  $\mathcal{R}(f, \epsilon, \mu)$  is sometimes restricted to contain only deterministic algorithms wlog (as can be done in (2)), but in the expected cost setting this may not necessarily be the case (see Open Problem 4); hence we allow  $\mathcal{R}(f, \epsilon, \mu)$  to contain randomised trees.

Yao showed an exact distributional characterisation for zero-error algorithms, namely,  $\bar{R}_0(f) = \bar{D}_0(f)$ , and moreover, the

optimal distributional algorithm is deterministic. He asked if a similar characterisation holds in the case  $\epsilon > 0$ . He observed that the “easy” direction of  $\overline{D}_\epsilon(f) \leq \overline{R}_\epsilon(f)$ , certainly holds (although Yao’s version of this inequality had some loss in parameters as he was restricted to deterministic algorithms). Vereshchagin [Ver98] proved the “hard” direction with a modest loss in parameters; in summary,

$$\overline{D}_\epsilon(f) \leq \overline{R}_\epsilon(f) \leq 2\overline{D}_{\epsilon/2}(f).$$

These bounds give a satisfying distributional characterisation in the bounded-error case. What happens in the small-bias case  $\epsilon = 1/2 - o(1)$ ? Our second result shows that, surprisingly, the distributional characterisation fails in a particularly strong sense. We write  $\delta = (1 - \epsilon)/2$  for short.

**Theorem 5.** *There is an  $n$ -bit partial function  $f$  and a bias  $\delta(n) = o(1)$  such that  $\overline{R}_\delta(f) \geq \overline{D}_\delta(f)^{1+\Omega(1)}$ .*

This theorem says that there is no way to capture  $\overline{R}_\epsilon(f)$  relative to a *single* hard distribution. However, there does exist a distributional characterisation using a pair of distributions, as we explore next.

#### D. Discussion: How are our two results related?

Suppose we want to prove an inner composition theorem. All the previous proofs [BK16], [AGJ<sup>+</sup>18], [GLSS19] revolve around the following high-level idea. Let  $R$  be a randomised tree that on input  $x$  seeks to compute  $f(g(x^1), \dots, g(x^n))$ . The tree can invest different numbers of queries  $q_i$  to different components  $x^i$ , making  $q = \sum_i q_i$  queries in total. If we had a complexity measure  $M(g)$  that allowed us to bound the bias the tree has for the  $i$ -th component  $g(x^i)$  as a *linear function* of  $q_i$ —say, the bias for  $g(x^i)$  is at most  $q_i/M(g)$ —then, by *linearity of expectation*, the expected total sum of the biases for all components  $g(x^1), \dots, g(x^n)$  is at most  $q/M(g)$ . This would allow us to track the total progress  $R$  is making across all the inner functions.

What is the largest such “linearised” measure  $M$ ? The most natural attempt at a definition (which the authors of this paper studied for a long time before finding the correct definition of LR) runs as follows. The measure should be such that with  $q := \overline{R}_\delta(f)$  queries one gets bias at most  $\delta \leq q/M(f)$ . Optimising for  $M(f)$  this suggests the following definition (a competitor for LR)

$$\text{ULR}(f) := \min_{\delta > 0} \frac{\overline{R}_\delta(f)}{\delta} = \min_R \max_{x,y} \frac{\text{cost}(R,x)}{\text{bias}(R,y)}.$$

We call it *uniform-LR*, since the tree  $R$  that achieves the minimum has an upper bound on  $\text{cost}(R,x)$  that is uniformly the same for all  $x$ , and similarly there is a uniform lower bound on  $\text{bias}(R,x)$  for all  $x$ . By contrast, the definition of  $\text{LR}(f)$  is *non-uniform*: a tree  $R$  that achieves the minimum for  $\text{LR}(f)$  has only a bound on the cost/bias *ratio*, but the individual cost and bias functions can vary wildly as a function of  $x$ .

We clearly have  $\text{LR}(f) \leq \text{ULR}(f)$  by definition. How about the converse? It is enlightening to compare the distributional characterisations of these two measures, which can be derived

using the recent minimax theorem for ratios of bilinear functions [BB20a]:

$$\text{LR}(f) := \min_R \max_x \frac{\text{cost}(R,x)}{\text{bias}(R,x)} = \max_\mu \min_R \frac{\text{cost}(R,\mu)}{\text{bias}(R,\mu)},$$

$$\text{ULR}(f) := \min_R \max_{x,y} \frac{\text{cost}(R,x)}{\text{bias}(R,y)} = \max_{\mu,\nu} \min_R \frac{\text{cost}(R,\mu)}{\text{bias}(R,\nu)}.$$

Here, LR is captured using a single hard distribution  $\mu$  such that both cost and bias are measured against it. By contrast, ULR needs a pair of distributions  $\mu, \nu$ , one to measure the cost, one to measure the bias. The upshot is that we are able to show that the two measures are polynomially separated.

**Theorem 6.** *There is an  $n$ -bit partial function  $f$  such that  $\text{ULR}(f) \geq \Omega(\text{LR}(f)^{5/4}) \geq n^{\Omega(1)}$ .*

Our optimality theorem thus implies that ULR *cannot* satisfy an inner composition theorem. This means that our attempt at finding a “linearised” measure at the start of this section missed a subtlety, namely, Yao’s question: can we capture our measure relative to a single hard distribution? Our proof of the composition theorem will rely heavily on the fact that LR admits a single hard distribution. Our separation of LR and ULR is what allows us to prove the impossibility of capturing  $\overline{R}_\epsilon(f)$  relative to a single distribution. Indeed, [Theorem 5](#) can be derived from [Theorem 6](#) simply as follows.

*Proof of Theorem 5.* Let  $f$  be as in [Theorem 6](#) and let  $R$  be a randomised tree witnessing  $\text{LR}(f)$ . We may assume wlog that  $\text{cost}(R,x) \geq 1$  for all  $x$ . (If  $R$  places a lot of weight on a 0-cost tree, we may re-weight  $R$  without affecting the cost/bias ratio; see [Lemma 9](#) for details.) Thus  $\text{bias}(R,x) \geq 1/n =: \delta$  for all  $x$ . We show the following inequalities, which would prove [Theorem 5](#).

$$\overline{R}_\delta(f) \geq \delta \cdot \text{ULR}(f), \tag{3}$$

$$\overline{D}_\delta(f) \leq \delta \cdot \text{LR}(f), \tag{4}$$

Indeed, (3) holds since  $\text{ULR}(f) \leq \overline{R}_\delta(f)/\delta$  by the definition of ULR. For (4) consider any input distribution  $\mu$ . Define  $R'$  as the randomised tree that with probability  $\lambda := \delta/\text{bias}(R,\mu)$  runs  $R$ , and with probability  $1 - \lambda$  makes no queries and outputs a random 0/1 answer. Then  $\text{bias}(R',\mu) = \lambda \text{bias}(R,\mu) = \delta$  and  $\text{cost}(R',\mu) = \lambda \text{cost}(R,\mu) = \delta \text{cost}(R,\mu)/\text{bias}(R,\mu) \leq \delta \text{LR}(f)$ , as desired.  $\square$

#### E. Techniques

**Composition theorem.** Our first result, the inner-optimal composition theorem, is proved in [Section III](#). As in other composition theorems for randomised algorithms, we start with a randomised algorithm  $R$  for the composition  $f \circ g$  as well as hard distributions  $\mu_0$  and  $\mu_1$  for  $g$  (corresponding to distributions on  $g^{-1}(0)$  and  $g^{-1}(1)$ ), and we construct a randomised algorithm  $R'$  for  $f$  whose cost is significantly lower than that of  $R$  (we need the cost to decrease by a factor of  $\text{LR}(g)$ ). The algorithm  $R'$  will simulate  $R$ , but not every query that  $R$  makes to the large,  $mn$ -sized input to  $f \circ g$  will turn into a query to the smaller,  $n$ -sized input to  $f$  that  $R'$

has access to. Instead,  $R'$  will attempt to delay making a true query as long as possible, and instead when  $R$  makes a query  $(i, j)$  (querying position  $j$  inside copy  $i$  of an input to  $g$ ),  $R'$  will return an answer that is generated according to  $\mu_0$  and  $\mu_1$ , so long as these two distributions approximately agree on the answer to that query.

So far, this is the same strategy employed by several other composition theorems, including in particular that of [GLSS19]. Our innovation comes from the precise way we choose when to query the bit  $i$  versus when to return an artificially-generated query answer to the query  $(i, j)$ . Specifically, in Section IV, we prove the following simulation theorem for decision trees. Suppose we are given two distributions  $\mu_0$  and  $\mu_1$ , we are asked to answer online queries to the bits of a string sampled from  $\mu_b$  without knowing the value of  $b$ ; moreover, suppose we have access to a big red button that, when pressed, provides the value of  $b \in \{0, 1\}$ . Then there is a strategy to answer these online queries with perfect soundness (i.e. with distribution identical to sampling a string from  $\mu_b$ ) with the following guarantee: if the decision tree that is making the online queries is  $D$ , then the probability we press the button is at most  $\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))$  (the total variation distance between the query outputs  $D$  receives when run on  $\mu_0$  and the query outputs  $D$  receives when run on  $\mu_1$ ).

This simulation theorem, though somewhat technical, ends up being stronger than the simulation guarantee used by Gavinsky, Lee, Santha, and Sanyal [GLSS19] to provide their composition result for max-conflict complexity. To get a composition theorem, we need to convert this total variation distance between transcripts into a more natural measure; this can be done via some minimax arguments, and the resulting measure is LR. We note that a similarly structured argument occurred in [BB20b], but the squared-Hellinger distance between the transcripts appeared instead of the total variation distance; in that result, the authors showed that this squared-Hellinger distance between transcripts characterized  $\mathbf{R}(g)$ , but they failed to construct a randomised algorithm  $R'$  for  $f$ , instead constructing only a “noisy” randomised algorithm. This gave them the result  $\mathbf{R}(f \circ g) = \Omega(\text{noisyR}(f)\mathbf{R}(g))$ . In contrast, the total variation distance allows us to get  $\mathbf{R}(f)$  on the outside, at the cost of getting only LR( $g$ ) on the inside.

The measure LR is arguably more natural than max-conflict complexity, but the real advantage is that our composition theorem turns out to be the best possible of its type: if  $\mathbf{R}(f \circ g) = \Omega(\mathbf{R}(f)\mathbf{M}(g))$  for all partial functions  $f$  and  $g$ , then  $\text{LR}(g) = \Omega(\mathbf{M}(g))$ . To show this, we give a characterization of LR( $g$ ) in terms of randomised query complexity: there is a family of partial functions  $f_m$  such that for all partial functions  $g$ , we have

$$\text{LR}(g) = \Theta\left(\frac{\mathbf{R}(f_m \circ g)}{\mathbf{R}(f_m)}\right),$$

where  $m$  is the input size of  $g$ . Once we have this, it clearly follows that  $\mathbf{R}(f \circ g) = \Omega(\mathbf{R}(f)\mathbf{M}(g))$  implies  $\text{LR}(g) = \Omega(\mathbf{M}(g))$ . The function family  $f_m$  turns out to be the same as the one introduced in [BB20b] (based on a

family of relations introduced in [GLSS19]); the randomised query complexity  $\mathbf{R}(f_m)$  was already established in that paper, so all we need is an upper bound on  $\mathbf{R}(f_m \circ g)$  which uses the existence of an LR-style algorithm for  $g$ . The linear dependence on the bias which is built into the definition of LR( $g$ ) turns out to be precisely what is needed to upper bound  $\mathbf{R}(f_m \circ g)$  (see Section VI for details).

**Failure of small-bias minimax.** Our second result, separation of LR and ULR, is proved in the full version of this article [BDBGM22]. The function  $f$  that witnesses the separation  $\text{ULR}(f) \geq \Omega(\text{LR}(f)^{5/4})$  is not hard to define. For simplicity, we denote its input length by  $N := Bn$  and think of the input as being composed of  $B = n^c$  blocks (for some large constant  $c$ ) of  $n$  bits each. We define  $f$  as a composition of  $\text{MAJ}_B$  as an outer function, and  $\text{XOR}_n$  as an inner function, where we are able to switch individual XOR-blocks to be easy (requiring  $O(1)$  queries) or hard (requiring  $n$  queries). Moreover, we make the following promises about the input. Either

- (1) all blocks are easy, and a random block has a value with bias  $1/n$  towards the majority value; or
- (2)  $b := n^{-3/4}$  fraction of the blocks are hard, and a random block has bias  $\Omega(b)$  towards the majority.

We claim that this function is easy for LR, namely,  $\text{LR}(f) = O(n)$ . To see this, consider the algorithm  $R$  that chooses a block at random, computes it, and outputs its value. For inputs  $x$  of type (1) we have  $\text{cost}(R, x) = O(1)$  and  $\text{bias}(R, x) \geq 1/n$  so that  $\text{cost}/\text{bias}$  ratio is  $O(n)$ . For inputs  $x$  of type (2) we have  $\text{cost}(R, x) = bn + (1-b)O(1) \leq O(bn)$  and  $\text{bias}(R, x) = b$  so that  $\text{cost}/\text{bias}$  ratio is  $O(n)$  again.

The difficult part is to show that  $\text{ULR}(f) \geq \Omega(n^{5/4})$ . For example, the above algorithm  $R$  has ULR-style measure  $\max_{x,y} \text{cost}(R, x)/\text{bias}(R, y) = O(bn)/(1/n) = O(n^{5/4})$ , and we would like to show that this is optimal. Intuitively, it is hard to get large bias for inputs of type (1) (although query cost is small here) and it is hard to get low query cost for inputs of type (2) (although bias is relatively high here). We first argue that an algorithm that wants to keep  $\text{cost}(R, x)$  small uniformly for all  $x$  (even those  $x$  with high  $\text{bias}(R, x)$ ) cannot afford to solve hard blocks very often. This is formalised by picking an appropriate pair of hard distributions for  $f$  according to the minimax formulation. What remains is the following task: Show that any algorithm that does not solve hard blocks, has large  $\text{cost}/\text{bias}$  ratio relative to a *single* hard distribution, that is, show an LR-style lower bound.

To this end, we develop a suite of techniques to prove lower bounds on the  $\text{cost}/\text{bias}$  trade-off achievable by decision trees in the small-bias expected cost setting, which has not really been studied in the literature before. Consequently, we end up having to re-establish some basic facts in the expected-cost setting that have been long known in the worst-case setting. For example, we show any algorithm for  $\text{GAPMAJ}_n$  (with  $\sqrt{n}$  gap promise) can achieve bias at most  $O(\sqrt{\text{cost}/n})$  (see the full version [BDBGM22]). The proof here exploits the “AND-trick” used by Sherstov [She12] to prove a lower bound on

the (worst-case) randomised communication complexity of the gap-Hamming problem. These techniques also come in handy when we separate LR from  $\bar{\chi}$  for the proof of [Lemma 4](#).

### F. Open questions

The foremost open question is to resolve [Conjecture 1](#). We can equivalently formulate it as follows.

**Open Problem 1** ([Conjecture 1](#) rephrased). *Does  $\text{LR}(f) = \Theta(\text{R}(f))$  for all total functions  $f$ ?*

One intriguing open problem regarding our new-found measure LR is to show that it is lower-bounded by quantum query complexity Q. Indeed, the bias of a quantum algorithm can be amplified linearly in the query cost, so it seems sensible to conjecture this is so. However, quantum query complexity has mostly been studied in the worst-case setting, and it is unclear how one should even define quantum query complexity in expectation (in such a way that it supports linear bias amplification).

**Open Problem 2.** *Does it hold that  $\text{LR}(f) \geq \text{Q}(f)$ ?*

There is a second reason to care about this question, having to do with the *composition limit* of randomised algorithms. Define  $\text{R}^*(f) := \lim_{k \rightarrow \infty} \text{R}(f^{\circ k})^{1/k}$ ; this is the limit of the  $k$ -th root of the randomised query complexity of the  $k$ -fold composition of  $f$ . Our results here imply that  $\text{R}^*(f) \geq \Omega(\text{LR}(f))$  for all (possibly partial) functions  $f$ . Due to the composition theorem for quantum query complexity, it is also known that  $\text{R}^*(f) \geq \Omega(\text{Q}(f))$ . The above open problem asks whether one of these results dominates the other. More generally, it would be nice to characterize  $\text{R}^*(f)$  in terms of a simpler measure (for instance, one which is efficiently computable given the truth table of the function).

Our inner-optimal composition theorem for LR, together with the outer-optimal composition theorem for noisyR [[BB20b](#)] give a relatively satisfying picture of composition in the case of partial functions. However, we can still ask whether there remain other *incomparable* composition theorems.

**Open Problem 3.** *Are there multiplicative composition theorems, stating that  $\text{R}(f \circ g) \geq \Omega(\text{M}_1(f)\text{M}_2(g))$  for all partial  $f, g$ , that can sometimes prove better lower bounds than  $\Omega(\max\{\text{R}(f)\text{LR}(g), \text{noisyR}(f)\text{R}(g)\})$ ?*

Regarding the failure of the distributional characterization of  $\text{R}_\epsilon$  in the low bias regime ([Theorem 5](#)), one may wonder whether the definition of  $\bar{\text{D}}_\epsilon$  should really involve randomized decision trees instead of deterministic ones. As hinted in [Section I-C](#), while considering deterministic trees is the natural choice in the bounded error regime, we feel it might not be in the regime where  $\epsilon \approx 1/2$ . Indeed, while a randomised decision tree can get cost arbitrarily close to zero for  $\epsilon$  approaching 1/2 (by taking an appropriate mixture with the zero-query tree), a deterministic one will get stuck at making one query and thus cost 1. Deciding whether the two versions

are equivalent (up to constant factors and additive terms) is our last open question.

**Open Problem 4.** *Define  $\bar{\text{D}}_\epsilon^*(f)$  for a boolean function  $f$  with*

$$\bar{\text{D}}_\epsilon^*(f) := \max_{\mu} \min_{D \in \mathcal{D}(f, \epsilon, \mu)} \text{cost}(D, \mu)$$

where  $\mathcal{D}(f, \epsilon, \mu)$  is the set of all deterministic decision trees solving  $f$  with error at most  $\epsilon$  relative to inputs sampled from  $\mu$ . For any partial  $f$  and  $\epsilon$ , do we have  $\bar{\text{D}}_\epsilon^*(f) \leq O(\bar{\text{D}}_\epsilon(f)+1)$ ?

## II. PRELIMINARIES

### A. Query complexity notation

Fix a natural number  $n \in \mathbb{N}$ . A total boolean function is a function  $f: \{0, 1\}^n \rightarrow \{0, 1\}$ . We will consider several generalizations of total boolean functions: first, there are *partial* boolean functions, which are defined on a domain which is a subset of  $\{0, 1\}^n$ . We use  $\text{Dom}(f) \subseteq \{0, 1\}^n$  to denote the domain of such a function. A further way to generalize boolean functions is to expand the input and output alphabets; that is, for finite sets  $\Sigma_I$  and  $\Sigma_O$ , we can consider functions  $f: \text{Dom}(f) \rightarrow \Sigma_O$  with  $\text{Dom}(f) \subseteq \Sigma_I^n$ , which take in input strings over the alphabet  $\Sigma_I$  and output a symbol in  $\Sigma_O$ .

A still further way to generalize such functions is to consider *relations* instead of partial functions. A relation is a subset of  $\Sigma_I^n \times \Sigma_O$ , or alternatively, it is a function that maps  $\Sigma_I^n$  to a subset of  $\Sigma_O$ . Any partial function can be viewed as a (total) relation, where on an input  $x$  which is not in the domain of the partial function, the corresponding relation relates all output symbols to  $x$  (meaning that if  $x$  is the input, any output symbol is considered valid).

Given a boolean function  $f$  (or, more generally, a relation), we will denote its *deterministic query complexity* by  $\text{D}(f)$ . This is the minimum height of a *decision tree*  $D$  which correctly computes  $f(x)$  on any  $x \in \text{Dom}(f)$ ; in other words, it is the minimum number of worst-case adaptive queries required by a deterministic algorithm computing  $f$ . For a formal definition, see [[BdW02](#)].

In this work we will mostly be dealing with randomised algorithms rather than deterministic ones, so let us more carefully define those. A *randomised query algorithm* or *randomised decision tree* will be a probability distribution over deterministic decision trees. Such deterministic decision trees will have internal nodes labeled by  $[n] := \{1, 2, \dots, n\}$  (representing the index of the input to query), arcs labeled by  $\Sigma_I$  (representing the symbol we might see after querying an index), and leaves labeled by  $\Sigma_O$  (representing output symbols to return at the end of the algorithm). We will assume that no internal node shares a label with an ancestor, meaning that a deterministic algorithm does not query the same index twice.

For such a randomised algorithm  $R$  and for an input  $x \in \Sigma_I^n$ , we denote by  $R(x)$  the random variable we get by sampling a deterministic tree  $D$  from  $R$ , and returning  $D(x)$  (the label of the leaf of  $D$  reached after starting from the root and taking the path determined by  $x$ ). For a function  $f$ , we write  $\text{err}_f(R, x) := \Pr_R[R(x) \neq f(x)]$  (or  $\Pr_R[R(x) \notin f(x)]$  if  $f$

is a relation), and we write  $\text{bias}_f^\pm(R, x) := 1 - 2 \text{err}_f(R, x)$ ,  $\text{bias}_f(R, x) := \max\{\text{bias}_f^\pm(R, x), 0\}$ ; we omit the subscript  $f$  when it is clear from context.

For a deterministic tree  $D$ , let  $\text{cost}(D, x)$  be the number of queries  $D$  makes on input  $x$ ; this is the height of the leaf of  $D$  that is reached when  $D$  is run on  $x$ . For a randomised algorithm  $R$ , we then define  $\text{cost}(R, x) := \mathbb{E}_{D \sim R}[\text{cost}(D, x)]$  (this is the expected number of queries  $R$  makes when run on  $x$ ).

We extend both of the above to distributions  $\mu$  over  $\Sigma_I^n$  instead of just inputs  $x$ ; that is, define

$$\begin{aligned} \text{bias}_f^\pm(R, \mu) &:= \mathbb{E}_{x \sim \mu}[\text{bias}_f^\pm(R, x)] \\ &= \mathbb{E}_{x \sim \mu} \mathbb{E}_{D \sim R}[\text{bias}_f^\pm(D, x)] \end{aligned}$$

$$\text{cost}(R, \mu) := \mathbb{E}_{x \sim \mu}[\text{cost}(R, x)] = \mathbb{E}_{x \sim \mu} \mathbb{E}_{D \sim R}[\text{cost}(D, x)],$$

with  $\text{bias}_f(R, \mu) := \max\{\text{bias}_f^\pm(R, \mu), 0\}$ . We also define  $\text{tran}(R, \mu)$  to be the random variable we get by sampling a decision tree  $D$  from  $R$ , a string  $x$  from  $\mu$ , and returning the pair  $(D, \ell)$ , where  $\ell$  is the leaf of  $D$  reached when  $D$  is run on  $x$ . Intuitively,  $\text{tran}(R, \mu)$  is the ‘‘transcript’’ when  $R$  is run on an input sampled from  $\mu$ , and such a transcript records all information that an agent running  $R$  knows about the input  $x$  at the end of the algorithm. We will use  $\text{TV}(\mu, \nu) := \frac{1}{2} \sum_{x \in \mathcal{X}} |\mu[x] - \nu[x]|$  to denote the total variation distance between distributions  $\mu$  and  $\nu$  over set  $\mathcal{X}$ . Most often, we will employ it with respect to the transcript of  $R$  on two different distributions as a way to quantify the extent to which  $R$  can tell these distributions apart.

We say that a randomised algorithm  $R$  computes  $f$  to error  $\epsilon$  if  $\text{err}_f(R, x) \leq \epsilon$  for all  $x \in \text{Dom}(f)$ . We then let  $\mathbf{R}_\epsilon(f)$  the minimum possible value of  $\max_x \text{cost}(R, x)$  over randomised algorithms  $R$  satisfying  $\text{err}_f(R, x) \leq \epsilon$  for all  $x \in \text{Dom}(f)$ . We also use  $\mathbf{R}_\epsilon(f)$  to denote the minimum number  $T$  such that there is a randomised algorithm  $R$  with  $\text{err}_f(R, x) \leq \epsilon$  for all  $x \in \text{Dom}(f)$  such that all decision trees in the support of  $R$  have height at most  $T$ . The difference between  $\mathbf{R}_\epsilon(f)$  and  $\overline{\mathbf{R}}_\epsilon(f)$  is that the former measures the worst-case cost of an algorithm computing  $f$  to error  $\epsilon$  (maximizing over both the input string and the internal randomness), while the latter measures the expected worst-case cost of the algorithm computing  $f$  to error  $\epsilon$  (this still maximizes over the input strings  $x$ , but takes an expectation over the internal randomness of the algorithm  $R$ ).

It is easy to see that  $\overline{\mathbf{R}}_\epsilon(f) \leq \mathbf{R}_\epsilon(f)$  for all  $f$ . The other direction also holds if we tolerate a constant-factor loss, as well as an additive constant loss in  $\epsilon$ ; to see this, note that if we cut off the  $\overline{\mathbf{R}}_\epsilon(f)$  algorithm after it makes 10 times more queries than it is expected to, then the probability of reaching such a cutoff is at most  $1/10$  by Markov’s inequality, and hence the error probability of the algorithm increases by at most  $1/10$ ; this converts an  $\overline{\mathbf{R}}_\epsilon(f)$  algorithm into a  $\mathbf{R}_\epsilon(f)$  algorithm.

Standard error reduction techniques imply that for a boolean function  $f$ ,  $\mathbf{R}_\epsilon(f)$  is related to  $\mathbf{R}_{\epsilon'}(f)$  by a constant factor that depends only on  $\epsilon$  and  $\epsilon'$ , so long as both are in  $(0, 1/2)$ . For

this reason, the value of  $\epsilon$  does not matter when  $\epsilon$  is a constant in  $(0, 1/2)$  (so long as we ignore constant factors and so long as the function is boolean), so we omit  $\epsilon$  when  $\epsilon = 1/3$ . The same error reduction property holds for  $\overline{\mathbf{R}}_\epsilon(f)$ . Combined with the Markov inequality argument above, both  $\mathbf{R}(f)$  and  $\overline{\mathbf{R}}(f)$  are the same measure (up to constant factors) for a boolean function and for constant values of  $\epsilon$ .

We warn that these equivalences break if  $f$  is not boolean (especially if  $f$  is a relation) or if the value of  $\epsilon$  is not constant; in particular, when  $\epsilon = 1/n$  or when  $\epsilon = 1/2 - 1/n$ , the values of  $\mathbf{R}_\epsilon(f)$  and  $\overline{\mathbf{R}}_\epsilon(f)$  may differ by more than a constant factor.

## B. Linearised $\mathbf{R}$

For a (possibly partial) boolean function  $f$  on  $n$  bits, we define

$$\text{LR}(f) := \min_R \max_x \frac{\text{cost}(R, x)}{\text{bias}(R, x)}.$$

Here  $R$  ranges over randomised decision trees and  $x$  ranges over the domain of  $f$ , and we treat  $0/0$  as  $\infty$ .

We call this measure *linearised randomised query complexity*. The name comes from the linear dependence on the bias achieved by the algorithm. Note that if we wanted to amplify bias  $\gamma$  to constant bias, we would, in general, have to repeat the algorithm  $\Theta(1/\gamma^2)$  times to do so. In some sense, then, the measure  $\mathbf{R}(f)$  charges  $1/\gamma^2$  for an algorithm that achieves bias  $\gamma$  instead of achieving constant bias. The measure  $\text{LR}(f)$ , in contrast, charges only  $1/\gamma$  for such an algorithm, so it can be up to quadratically smaller than  $\mathbf{R}(f)$ .

A minimax theorem for ratios such as [BB20a] (Theorem 2.18) can show that

$$\text{LR}(f) = \max_\mu \min_D \frac{\text{cost}(D, \mu)}{\text{bias}(D, \mu)}, \quad (5)$$

where  $D$  ranges over deterministic decision trees and  $\mu$  ranges over probability distributions over  $\text{Dom}(f)$ .

It is not hard to see that the maximizing distribution  $\mu$  above will place equal weight on 0 and 1 inputs. This is because otherwise, we could take  $D$  to be a decision tree that makes 0 queries, and then  $\text{cost}(D, \mu)$  would be 0 while  $\text{bias}(D, \mu)$  would be positive.

If  $\mu$  is balanced over 0 and 1 inputs, we may express it as  $\mu := \mu_0/2 + \mu_1/2$  and it is not hard to show that for the best possible choice of leaf labels for an unlabeled decision tree  $D$ , we have

$$\text{bias}(D, \mu)^\pm = \text{bias}(D, \mu) = \text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1)).$$

This follows, for example, from [BB20a] (Lemma 3.9); to see this intuitively, recall that  $\text{tran}(D, \mu)$  is the random variable for the leaf of  $D$  reached when  $D$  is run on  $\mu$ , and note that the best choice of leaf label if  $D$  reaches a leaf  $\ell$  is 0 if the probability of  $D$  reaching  $\ell$  is higher when run on  $\mu_0$  than on  $\mu_1$ , and it is 1 otherwise. Therefore, the bias for the best choice of leaf labels is the sum, over leaves  $\ell$  of  $D$ , of  $2 \max\{\Pr_{\mu_0}[\ell], \Pr_{\mu_1}[\ell]\} - 1$ , which is easily seen to be the total variation distance between the two distributions over leaves.

Given (??), we can also write

$$\text{LR}(f) = \max_{\mu_0, \mu_1} \min_D \frac{\text{cost}(D, \frac{\mu_0 + \mu_1}{2})}{\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))},$$

where  $\mu_0$  ranges over probability distributions with support  $f^{-1}(0)$  and  $\mu_1$  ranges over probability distributions with support  $f^{-1}(1)$ . Observe that neither the top nor the bottom depend on the leaf labels of  $D$ , so we can now assume  $D$  is an unlabeled decision tree if we wish. Note also that  $\text{cost}(D, \mu)$  is linear in the second argument, so we can write

$$\text{LR}(f) = \max_{\mu_0, \mu_1} \min_D \frac{\text{cost}(D, \mu_0) + \text{cost}(D, \mu_1)}{2\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))}.$$

We clearly have

$$\text{LR}(f) \geq \max_{\mu_0, \mu_1} \min_D \frac{\min\{\text{cost}(D, \mu_0), \text{cost}(D, \mu_1)\}}{\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))}.$$

**Lemma 7.** For any fixed  $\mu_0$  and  $\mu_1$ , we have

$$\begin{aligned} & \min_D \frac{\text{cost}(D, \mu_1)}{\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))} \\ & \leq 6 \min_D \frac{\text{cost}(D, \mu_0)}{\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))}. \end{aligned}$$

*Proof.* See the full version of this article [BDBGM22].  $\square$

**Corollary 8.**

$$\begin{aligned} \text{LR}(f) & \geq \max_{\mu_0, \mu_1} \min_D \frac{\min\{\text{cost}(D, \mu_0), \text{cost}(D, \mu_1)\}}{\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))} \text{ and} \\ \text{LR}(f) & \leq 6 \max_{\mu_0, \mu_1} \min_D \frac{\min\{\text{cost}(D, \mu_0), \text{cost}(D, \mu_1)\}}{\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))} \end{aligned}$$

One useful property of LR complexity is that up to a multiplicative factor of 2, we can consider only randomised decision trees that always query at least one bit of their input.

**Lemma 9.** For every non-constant partial function  $g$ , there is a randomised decision tree  $A$  that always queries at least one bit of  $g$ 's input and satisfies, for every  $x$ ,

$$\frac{\text{cost}(A, x)}{\text{bias}(A, x)} \leq 2 \cdot \text{LR}(g).$$

*Proof.* See the full version of this article [BDBGM22].  $\square$

As a corollary, we obtain a universal lower bound on the LR complexity of every non-constant function.

**Corollary 10.** For every non-constant partial function  $g$ ,  $\text{LR}(g) \geq \frac{1}{2}$ .

*Proof.* By Lemma 9, there exists a randomised decision tree  $A$  that always queries at least one bit of its input and satisfies  $\text{cost}(A, x)/\text{bias}(A, x) \leq 2 \cdot \text{LR}(g)$  for all  $x$  in the domain of  $g$ . But since  $A$  always makes at least one query,  $\text{cost}(A, x) \geq 1$ . And by definition,  $\text{bias}(A, x) \leq 1$ , so the cost-bias ratio of  $A$  is always bounded below by 1.  $\square$

### III. AN OVERVIEW OF THE COMPOSITION THEOREM

In the following sections, we prove our inner-optimal composition theorem, Theorems 1 and 2 restated below, along with related results.

**Theorem 1.**  $\text{R}(f \circ g) \geq \Omega(\text{R}(f)\text{LR}(g))$  for all partial boolean functions  $f, g$ .

**Theorem 2.** Theorem 1 is optimal: If  $\mathbb{M}$  is any complexity measure such that  $\text{R}(f \circ g) \geq \Omega(\text{R}(f)\mathbb{M}(g))$  for all partial  $f, g$ , then  $\text{LR}(g) \geq \Omega(\mathbb{M}(g))$  for all partial  $g$ .

The heart of the proof of Theorem 1 is a simulation theorem showing that for any two distributions  $\mu_0$  and  $\mu_1$  and any decision tree  $T$ , it is possible to simulate  $T$  on inputs drawn from  $\mu_b$  for some initially unknown  $b \in \{0, 1\}$  while querying the actual value of  $b$  with probability bounded by the total variation distance between the two distributions  $\mu_0$  and  $\mu_1$ . This result, Theorem 11, is established in Section IV.

In Section V, we use the simulation theorem to complete the proof of the main composition theorem, Theorem 13, a slightly more general version of Theorem 1. We also use the simulation theorem to establish the perfect composition for LR complexity, Theorem 3, in this section.

The proof of Theorem 2 is completed in Section VI. Finally, in Section VII, we establish the separation between LR complexity and max-conflict complexity of Lemma 4.

### IV. DECISION TREE SIMULATION THEOREM

An *online decision tree simulator* is a randomised algorithm that is given two distributions  $\mu_0$  and  $\mu_1$  on inputs  $\{0, 1\}^n$ , oracle access to a bit  $b \in \{0, 1\}$ , and a stream of queries  $i_1, \dots, i_k \in [n]$  that represent the queries made by a decision tree  $T$  that is not known to the algorithm. The goal of an online decision tree simulator is to answer the queries according to the distribution  $\mu_b$  while querying the value of  $b$  itself with as small probability as possible. We think of this protocol as having a big red button that gives  $b$ , and it tries to pretend to have a sample from  $\mu_b$  without pressing the button for as long as possible.

**Theorem 11.** There exists an online decision tree simulator that simulates the queries of  $T$  on  $\mu_b$  while querying the value of  $b$  with probability  $\text{TV}(\text{tran}(T, \mu_0), \text{tran}(T, \mu_1))$ .

The algorithm that satisfies the theorem is stated below. In the algorithm,  $x \in \{0, *, 1\}^n$  is a partially defined boolean string: the coordinates labelled with  $*$  are undefined. Given a string  $x \in \{0, *, 1\}^n$ , an index  $i \in [n]$ , and a value  $a \in \{0, 1\}$ , the notation  $x^{(i \leftarrow a)}$  denotes the string  $y$  which equals  $x$  on all coordinates except  $i$ , where it takes the value  $y_i = a$ .

Note that each vertex in a decision tree  $T$  corresponds to the partial string  $x \in \{0, 1, *\}^n$  of the values revealed on the path to that vertex in  $T$ . Our main task is to show that each vertex in  $T$  (including each leaf) is reached with probability  $\mu_b(x)$  in the algorithm and that the probability that we reach  $x$  and don't reveal  $b$  along the way is  $\mu_{\min}(x)$ .

---

**Algorithm 1:** ONLINEQUERYSIMULATOR( $\mu_0, \mu_1$ )

---

**for all**  $x \in \{0, *, 1\}^n$  **do**  
     $\mu_{\min}(x) \leftarrow \min\{\mu_0(x), \mu_1(x)\};$   
     $x \leftarrow *^n;$   
     $b \leftarrow *;$   
**while** *more queries remain* **do**  
     $i \leftarrow \text{NEXTQUERY};$   
     $u \leftarrow \mu_{\min}(x^{(i \leftarrow 0)}) + \mu_{\min}(x^{(i \leftarrow 1)});$   
    **if**  $b = *$  **then**  
        With probability  $1 - u/\mu_{\min}(x)$ , query the value of  $b$ ;  
    **if**  $b = *$  **then**  
         $x_i \leftarrow \text{Ber}(\mu_{\min}(x^{(i \leftarrow 1)})/u);$   
    **else**  
         $x_i \leftarrow \text{Ber}\left(\frac{\mu_b(x^{(i \leftarrow 1)}) - \mu_{\min}(x^{(i \leftarrow 1)})}{\mu_b(x) - u}\right);$

---

**Lemma 12.** *For every  $x \in \{0, 1, *\}^n$ , when we run the ONLINEQUERYSIMULATOR, then*

- 1) *We reach the vertex  $x$  with probability  $\mu_b(x)$ , and*
- 2) *We reach the vertex  $x$  and don't query the value  $b$  on the way to  $x$  with probability  $\mu_{\min}(x)$ .*

*Proof.* We prove the claim by induction on the number of defined coordinates on  $x$ . The base case corresponds to  $x = *^n$ , which trivially satisfies both conditions of the claim.

Consider now any  $x \neq *^n$ . Let  $z$  be the parent of  $x$  in the decision tree  $T$ , and let  $i$  denote the coordinate where  $z_i = *$  and  $x_i \neq *$ . Define also  $y$  to be  $x$ 's sibling in  $T$ . Let us assume that  $x_i = 1$ . (The case where  $x_i = 0$  is essentially identical.)

By the induction hypothesis, the probability that we reach  $z$  and don't query the value  $b$  is  $\mu_{\min}(z)$ . With probability  $(\mu_{\min}(x) + \mu_{\min}(y))/\mu_{\min}(z)$ , we don't query the value of  $b$  while processing the query  $i$  either. And when this occurs the algorithm next reaches  $x$  with probability  $\mu_{\min}(x)/(\mu_{\min}(x) + \mu_{\min}(y))$ . So the overall probability that we reach  $x$  without querying  $b$  along the way is

$$\mu_{\min}(z) \cdot \frac{\mu_{\min}(x) + \mu_{\min}(y)}{\mu_{\min}(z)} \cdot \frac{\mu_{\min}(x)}{\mu_{\min}(x) + \mu_{\min}(y)} = \mu_{\min}(x).$$

Next, by the induction hypothesis again the probability that we query the value of  $b$  either on the way to  $z$  or while processing the query  $i$  is

$$\begin{aligned} & (\mu_b(z) - \mu_{\min}(z)) + \mu_{\min}(z) \cdot \left(1 - \frac{\mu_{\min}(x) + \mu_{\min}(y)}{\mu_{\min}(z)}\right) \\ &= \mu_b(z) - (\mu_{\min}(x) + \mu_{\min}(y)). \end{aligned}$$

Then the probability we output  $x$  conditioned on having revealed  $b$  is

$$\frac{\mu_b(x) - \mu_{\min}(x)}{\mu_b(z) - (\mu_{\min}(x) + \mu_{\min}(y))},$$

so that the overall probability that we reach  $x$  and reveal  $b$  along the way is  $\mu_b(x) - \mu_{\min}(x)$ . Therefore, the overall probability that we reach  $x$  is  $\mu_b(x)$ .  $\square$

The proof of [Theorem 11](#) is now essentially complete, as it just requires combining the lemma with a simple identity on total variation distance.

*Proof of Theorem 11.* [Lemma 12](#) implies that the Oracle-QuerySimulator indeed reaches each leaf with the correct probability  $\mu_b(x)$ . And the probability that it queries the value of  $b$  is  $1 - \sum_{\ell \in T} \min\{\mu_0(\ell), \mu_1(\ell)\}$ , which is the total variation distance between  $\text{tran}(T, \mu_0)$  and  $\text{tran}(T, \mu_1)$ .  $\square$

## V. COMPOSITION THEOREMS

The inner-optimal composition theorem, [Theorem 1](#), is established in [Section V-A](#). In fact, we establish a slight generalization of that theorem, stated below in [Theorem 13](#). Then the perfect composition theorem for LR complexity, [Theorem 3](#), is established in [Section V-B](#).

### A. Composition for randomised query complexity

For a boolean string  $y \in \{0, 1\}^n$  and a pair of distributions  $\mu_0, \mu_1$ , we define  $y \circ (\mu_0, \mu_1)$  to be the product distribution  $\bigotimes_{i=1}^n \mu_{y_i}$ . In particular, if  $\mu_0$  and  $\mu_1$  are hard distributions for the 0- and 1-inputs of  $g$  respectively, and if  $y$  is an input to  $f$ , then  $y \circ (\mu_0, \mu_1)$  will give a distribution over the inputs to the composed  $f \circ g$  (all of which correspond to the same  $f$ -input  $y$ ).

We prove the following composition theorem, which is a slightly more general version of [Theorem 1](#).

**Theorem 13.** *Let  $\Sigma_I$  and  $\Sigma_O$  be finite alphabets, and let  $n, m \in \mathbb{N}$ . Let  $f \subseteq \{0, 1\}^n \times \Sigma_O$  be a (possibly partial) relation on  $n$  bits, and let  $g: \text{Dom}(g) \rightarrow \{0, 1\}$  be a (possibly partial) boolean function, with  $\text{Dom}(g) \subseteq \Sigma_I^m$ . Let  $\epsilon \in [0, 1/2)$ . Then*

$$\bar{R}_\epsilon(f \circ g) \geq \bar{R}_\epsilon(f)\text{LR}(g)/6.$$

*Proof.* Let  $\mu_0$  and  $\mu_1$  be distributions over the 0-inputs and 1-inputs to  $g$ , respectively, that maximize the expression in the right-hand side of [Corollary 8](#). Let  $\Pi$  be the online decision tree simulator from [Theorem 11](#). Let  $R$  be a randomised algorithm that computes  $f \circ g$  to error  $\epsilon$  using  $\bar{R}_\epsilon(f \circ g)$  expected queries. We describe a randomised algorithm  $R'$  for computing  $f$  on worst-case inputs.

Given input  $y \in \{0, 1\}^n$ , the algorithm  $R'$  will instantiate  $n$  copies of  $\Pi$ , which we denote  $\Pi_1, \Pi_2, \dots, \Pi_n$ , one for each bit of the input; if protocol  $\Pi_i$  presses the button, it gets  $y_i$  (and this causes  $R'$  to make a real query to the real input). Each of these copies of  $\Pi$  will assume the distributions to be simulated are  $\mu_0$  and  $\mu_1$ . Then  $R'$  will run  $R$ , and whenever  $R$  makes a query  $(i, j)$  (corresponding to querying bit  $j$  inside of the  $i$ -th copy of  $g$ ), the algorithm  $R'$  will ask  $\Pi_i$  to give an answer to query  $j$ , and it will use that answer to determine the next query of  $R$ .

Note that since the protocols  $\Pi_i$  are guaranteed to be sound, the outcome of the simulation of  $R$  made by  $R'$  is precisely the same (in distribution) as the outcome of running  $R$  on an input sampled from  $y \circ (\mu_0, \mu_1)$ . Therefore, by the correctness guarantee of  $R$ , the output will be a valid output for  $f(y)$ .



except with error probability  $\epsilon$ . It remains to show that for each  $y \in \text{Dom}(f)$ , the expected number of real queries  $R'$  makes when run on  $y$  is at most  $6\bar{R}_\epsilon(f \circ g)/\text{LR}(g)$ .

Fix any  $y \in \text{Dom}(f)$ . Now, when  $R'$  is run on  $y$ , let  $T$  be the expected number of fake queries it makes; in other words, let  $T = \text{cost}(R, y \circ (\mu_0, \mu_1)) \leq \bar{R}_\epsilon(f \circ g)$ . For each  $i$ , let  $T_i$  be the expected number of queries to  $\Pi_i$  that  $R'$  makes when run on  $y$ , so that  $T_1 + T_2 + \dots + T_n = T$ . Let  $p_i$  the overall probability that  $\Pi_i$  presses the button when  $R'$  runs on  $y$ ; the sum  $q = p_1 + p_2 + \dots + p_n$  is therefore the expected number of real queries made by  $R'$  on  $y$ . We would like to show that  $q \leq 6T/\text{LR}(g)$ , or equivalently,  $T/q \geq \text{LR}(g)/6$ .

Since  $T/q = (T_1 + \dots + T_n)/(p_1 + \dots + p_n)$ , there must be some  $i$  such that  $T/q \geq T_i/p_i$ . It will therefore suffice to show that  $T_i/p_i \geq \text{LR}(g)/6$  for all  $i \in [n]$ . Fix such  $i$ , and recall that  $T_i$  is the number of (fake) queries  $R'$  makes to  $\Pi_i$  when run on  $y$ , and  $p_i$  is the probability that  $\Pi_i$  presses the button when  $R'$  is run on  $y$ . Consider the algorithm  $R_{y,i}$  which takes in an input  $x$  in  $\text{Dom}(g)$ , generates  $n-1$  additional fake inputs to  $g$  from the distributions  $\mu_{y_\ell}$  for  $\ell \neq i$ , places the real input  $x$  as the  $i$ -th input among the  $n$  inputs to  $g$ , and runs  $R$  on this tuple (treating it as an input to  $f \circ g$ ). Note that when  $R_{y,i}$  is run on an input from  $\mu_{y_i}$ , its behavior is exactly the same as the behavior of  $R$  when run on  $y \circ (\mu_0, \mu_1)$ ; therefore, it makes  $T_i$  expected queries. Consider running  $R_{y,i}$  with query answers generated by  $\Pi$  instead of by making real queries; then when  $\Pi$  uses the hidden bit  $y_i$  and simulates the distributions  $\mu_0, \mu_1$ , the behavior of  $R_{y,i}$  is the same as when we run it on  $\mu_{y_i}$ , and hence the expected number of queries it makes to  $\Pi$  is  $T_i$  and the probability that  $\Pi$  presses the button is exactly  $p_i$ .

Now, by [Theorem 11](#), we know that  $\Pi$  presses the button with probability  $\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))$  when simulating a deterministic decision tree  $D$ . For a random decision tree such as the one given by  $R_{y,i}$ , the probability  $p_i$  of the button being pressed will be the mixture of the values  $\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))$  for the deterministic decision trees  $D$  in the support of  $R_{y,i}$ . Also, the expected number of queries  $T_i$  that  $R_{y,i}$  makes is a matching mixture of the expected number of queries made by the decision trees  $D$  in the support of  $R_{y,i}$ ; the latter is  $\text{cost}(D, \mu_{y_i})$ . Hence to lower bound  $T_i/p_i$ , it will suffice to lower bound  $\frac{\text{cost}(D, \mu_{y_i})}{\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))}$  for all deterministic decision trees  $D$  acting on inputs in  $\text{Dom}(g)$ . We now write

$$\begin{aligned} & \frac{\text{cost}(D, \mu_{y_i})}{\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))} \\ & \geq \frac{\min\{\text{cost}(D, \mu_0), \text{cost}(D, \mu_1)\}}{\text{TV}(\text{tran}(D, \mu_0), \text{tran}(D, \mu_1))} \\ & \geq \text{LR}(g)/6 \end{aligned}$$

(using [Corollary 8](#)). The desired result follows.  $\square$

### B. Composition for LR complexity

**Theorem 3.**  $\text{LR}(f \circ g) \geq \Omega(\text{LR}(f)\text{LR}(g))$  for all partial boolean functions  $f, g$ .

We actually prove the more explicit result  $\text{LR}(f \circ g) \geq \text{LR}(f)\text{LR}(g)/6$ .

*Proof.* The proof is similar to that of [Theorem 13](#). We fix hard distributions  $\mu_0$  and  $\mu_1$  for  $\text{LR}(g)$ , and we fix a randomised algorithm  $R$  for  $f \circ g$  such that  $\max_z \text{cost}(R, z)/\text{bias}(R, z) \leq \text{LR}(f \circ g)$ . We then define a randomised algorithm  $R'$  for  $f$ ; this time, unlike in the proof of [Theorem 13](#), we want  $R'$  to solve  $f$  in the  $\text{LR}(f)$  sense instead of being a randomised algorithm that solves  $f$  to error  $\epsilon$ . We define  $R'$  as before: on input  $y \in \text{Dom}(f)$ ,  $R'$  instantiates  $n$  protocols  $\Pi_i$ , one for each bit of  $y$ ; it instantiates each with the distributions  $(\mu_0, \mu_1)$ , and gives  $\Pi_i$  the hidden bit  $y_i$  if it presses the button. Then  $R'$  will run  $R$ , and whenever  $R$  makes a query  $(i, j)$  (to the bit  $j$  inside the  $i$ -th input to  $g$ ),  $R'$  will ask  $\Pi_i$  for bit  $j$ .

Note that by the soundness of the protocols  $\Pi_i$ , we have  $\text{bias}(R', y) = \text{bias}(R, y \circ (\mu_0, \mu_1))$ . We will next show that  $\text{cost}(R', y) \leq 6 \text{cost}(R, y \circ (\mu_0, \mu_1))/\text{LR}(g)$ ; This way, we will have

$$\begin{aligned} \text{LR}(f) &= \max_y \frac{\text{cost}(R', y)}{\text{bias}(R', y)} \\ &\leq \frac{6}{\text{LR}(g)} \max_y \frac{\text{cost}(R, y \circ (\mu_0, \mu_1))}{\text{bias}(R, y \circ (\mu_0, \mu_1))} \\ &\leq \frac{6}{\text{LR}(g)} \max_z \frac{\text{cost}(R, z)}{\text{bias}(R, z)} \\ &\leq \frac{6\text{LR}(f \circ g)}{\text{LR}(g)}. \end{aligned}$$

Fix any  $y \in \text{Dom}(f)$ ; it remains to show that  $\text{cost}(R, y \circ (\mu_0, \mu_1))/\text{cost}(R', y) \geq \text{LR}(g)/6$ . For every  $i \in [n]$ , let  $T_i$  be the expected number of queries  $R$  makes to the  $i$ -th input on  $y \circ (\mu_0, \mu_1)$ , and let  $p_i$  be the probability that  $R'$  queries the  $i$ -th bit when run on input  $y$ . Let  $T = T_1 + \dots + T_n$ , and let  $q = p_1 + \dots + p_n$ . We wish to show  $T/q \geq \text{LR}(g)/6$ . This precise statement was shown in the proof of [Theorem 13](#), which completes this proof as well.  $\square$

## VI. OPTIMALITY OF THE COMPOSITION THEOREM

We complete the proof of [Theorem 2](#) in this section.

**Theorem 2.** *Theorem 1 is optimal: If  $\mathbf{M}$  is any complexity measure such that  $\mathbf{R}(f \circ g) \geq \Omega(\mathbf{R}(f)\mathbf{M}(g))$  for all partial  $f, g$ , then  $\text{LR}(g) \geq \Omega(\mathbf{M}(g))$  for all partial  $g$ .*

The proof of [Theorem 2](#) is obtained by a characterization of LR complexity in terms of the complexity of functions composed with the *approximate index* partial function  $\text{APPROXINDEX}_k : \{0, 1\}^k \times \{0, 1, 2\}^{2^k} \rightarrow \{0, 1, *\}$  where  $\text{APPROXINDEX}_k(a, y) = y_a$  if  $y_a \in \{0, 1\}$ ,  $y_b = y_a$  for all  $|b - a| \leq \frac{k}{2} - 2\sqrt{k \log k}$  and  $y_b = 2$  for all other  $b$  and  $\text{APPROXINDEX}_k(a, y) = *$  else. The randomised query complexity of the approximate index function is as follows.

**Lemma 14** ([\[BB20b, Lemma 27\]](#)).  $\mathbf{R}(\text{APPROXINDEX}_k) = \Theta(\sqrt{k \log k})$ .

The key to the proof of [Theorem 2](#) is the following characterization of LR complexity in terms of composition with the approximate index function.

**Lemma 15.** For every partial boolean function  $g : \Sigma^m \rightarrow \{0, 1, *\}$ , when  $k \in \mathbb{N}$  satisfies  $\frac{k}{\log k} \geq (36m)^2$  then

$$\text{LR}(g) = \Theta\left(\frac{\mathbb{R}(\text{APPROXINDEX}_k \circ g)}{\mathbb{R}(\text{APPROXINDEX}_k)}\right).$$

*Proof.* The lemma trivially holds when  $g$  is a constant function. For the rest of the proof, fix  $g$  to be any non-constant partial function. [Theorem 1](#) implies the upper bound

$$\text{LR}(g) = O\left(\frac{\mathbb{R}(\text{APPROXINDEX}_k \circ g)}{\mathbb{R}(\text{APPROXINDEX}_k)}\right).$$

The goal of the remainder of the proof is to establish a matching lower bound by showing that

$$\bar{\mathbb{R}}(\text{APPROXINDEX}_k \circ g) = O\left(\sqrt{k \log k} \cdot \text{LR}(g)\right).$$

This bound suffices to complete the proof because  $\bar{\mathbb{R}}(f) = \Theta(\mathbb{R}(f))$  for every partial function  $f$ .

Let  $R$  denote a randomised algorithm that satisfies

$$\text{cost}(R, x) \leq 2 \cdot \text{LR}(g) \cdot \text{bias}(R, x)$$

for all  $x$  in the domain of  $g$  and always queries at least one bit of its input. Such an algorithm is guaranteed to exist by [Lemma 9](#). We define a new randomised algorithm  $A$  that proceeds as follows: it runs the algorithm  $R$  sequentially on the first instances  $x_1, x_2, \dots, x_\ell$  of  $g$  which correspond to the initial address bits of the input to  $\text{APPROXINDEX}_k$ . It continues this process until the total number of queries made to the underlying inputs exceeds  $36\sqrt{k \log k} \cdot \text{LR}(g)$ . By the choice of  $k$  and the trivial bound  $\text{LR}(g) \leq m$ , this process terminates when  $R$  has computed the first  $\ell$  instances of  $g$  with some biases  $b_1, \dots, b_\ell$  for some  $\ell \leq k$ . The algorithm  $A$  then guesses the value of the remaining  $k - \ell$  bits of the address. It finally computes the value of  $g$  on the instance corresponding to the address obtained with error probability at most  $\frac{1}{9}$  and returns that value.

Let  $c_1, \dots, c_\ell$  denote the query cost incurred by  $R$  when running on the  $\ell$  computed instances of  $g$ . The random variables  $(X_i)_{i \leq k}$  defined by  $X_i = \sum_{j \leq i} c_j - \text{cost}(R, x_j)$  form a discrete-time martingale and  $\ell$  is the stopping time of this martingale. By the optional stopping theorem,  $\mathbb{E}[X_\ell] = 0$ . So  $\mathbb{E}[\sum_{i \leq \ell} c_i] = \sum_{i \leq \ell} \text{cost}(R, x_i)$ . By Markov's inequality, the probability that the total cost exceeds 6 times the expected cost on the same inputs is at most  $1/6$ ; let us consider from now on only the case when this does not occur. In this case,

$$\sum_{i=1}^{\ell} \text{cost}(R, x_i) \geq \frac{1}{6} \sum_{i=1}^{\ell} c_i \geq 6\sqrt{k \log k} \cdot \text{LR}(g).$$

By our choice of  $R$ , the biases  $\beta_1, \dots, \beta_\ell$  on the values  $g(x_1), \dots, g(x_\ell)$  satisfy  $\sum_{i=1}^{\ell} \beta_i \geq 3\sqrt{k \log k}$  and so if we let  $b \in \{0, 1\}^k$  denote the address computed by the algorithm, we observe that

$$\mathbb{E}[|b - a|] = \sum_{i=1}^k \Pr[b_i \neq g(x_i)] \leq \frac{k}{2} - 3\sqrt{k \log k}.$$

Furthermore, each of the  $k$  events  $b_i \neq g(x_i)$  are independent. So by Hoeffding's bound the probability that more than  $\frac{k}{2} - 2\sqrt{k \log k}$  of these events occur is at most  $e^{-2 \log^2 k}$ , which is less than  $\frac{1}{9}$  when  $k \geq 3$ . When this event does not occur, the address  $b$  computed by the algorithm satisfies  $x_b = x_a$ . Since  $A$  lastly computes  $g(x_b)$  with error at most  $\frac{1}{9}$ , in total it computes  $\text{APPROXINDEX}_k \circ g$  with error at most  $\frac{1}{3}$ .

It remains to show that the expected query cost of the algorithm  $A$  satisfies the desired bound. The first round of the algorithm uses at most  $36\sqrt{k \log k} \cdot \text{LR}(g)$  queries plus the number of queries of the instance of  $R$  run on  $x_\ell$ . In expectation, this additional number of queries is at most  $\text{cost}(R, x_\ell) \leq \text{LR}(g)$ . And then computing  $g(x_b)$  requires another  $\mathbb{R}(g) \leq m < \sqrt{k}$  queries. So the overall expected query complexity of  $A$  is at most  $(36\sqrt{k \log k} + 1) \cdot \text{LR}(g) + \sqrt{k}$ . By [Corollary 10](#),  $\text{LR}(g) \geq \frac{1}{2}$  for every non-constant function  $g$  so this query complexity is bounded above by  $O(\sqrt{k \log k} \cdot \text{LR}(g))$ , as required.  $\square$

The proof of [Theorem 2](#) now follows easily from [Lemma 15](#).

*Proof of Theorem 2.* Let  $M$  be a measure that satisfies the condition of the theorem. Then, choosing  $f$  to be the  $\text{APPROXINDEX}_k$  function for a large enough value of  $k$  and applying [Lemma 15](#), we obtain

$$M(g) = O\left(\frac{\mathbb{R}(\text{APPROXINDEX}_k \circ g)}{\mathbb{R}(\text{APPROXINDEX}_k)}\right) = O(\text{LR}(g)). \quad \square$$

## VII. SEPARATION FROM $\bar{\chi}$

In this section, we exhibit a polynomial separation between  $\text{LR}$  and  $\bar{\chi}$ , the max conflict complexity introduced by Gavinsky, Lee, Santha and Sanyal in [\[GLSS19\]](#) (see [Section VII-A](#) for a formal definition of  $\bar{\chi}$ ).

**Lemma 4.** There exists a partial  $f$  such that  $\text{LR}(f) \geq \Omega(\bar{\chi}(f)^{1.5})$ .

*Proof.* The function  $f$  we build takes input of size  $n^2 + \sqrt{n}$  with format  $(x_1, \dots, x_{n^2}, a_1, \dots, a_{\sqrt{n}})$ . The function value is given as the parity of  $\text{GAPMAJ}(x)$  and  $\text{XOR}(a)$ , i.e.:

$$f(x_1, \dots, x_{n^2}, a_1, \dots, a_{\sqrt{n}}) = \text{GAPMAJ}_{n-1/2}^{n^2}(x_1, \dots, x_{n^2}) \oplus \text{XOR}_{\sqrt{n}}(a_1, \dots, a_{\sqrt{n}})$$

$\text{GAPMAJ}_{n-1/2}^{n^2}(x)$  is the majority function on  $n^2$  bits with promise that  $|x| \notin [n^2/2 - n^{3/2}, n^2/2 + n^{3/2}]$  so that returning the value of a random index holds bias at least  $n^{-1/2}$ . Thus,  $f$  is a partial function whose domain is constrained by the gap majority instance. [Lemma 16](#) shows that  $\text{LR}(f) \geq \Omega(n^{3/4})$  and [Lemma 17](#) that  $\bar{\chi}(f) \leq O(n^{1/2})$ , as desired.  $\square$

**Lemma 16.**  $\text{LR}(F) \geq \Omega(n^{3/4})$

*Proof.* See the full version of this article [\[BDBGM22\]](#).  $\square$

### A. An upper bound for $\bar{\chi}$

We recall here the definition of max conflict complexity (but see [GLSS19] for an in-depth treatment of the measure). Let  $f$  be a fixed boolean function,  $\mu^0, \mu^1$  a pair of distribution over  $f^{-1}(0)$  and  $f^{-1}(1)$  respectively and  $D$  a deterministic decision tree solving  $f$ . For each node  $v$  in  $D$ , we let  $\mu^0|_v, \mu^1|_v$  be the distributions conditioned on reaching  $v$  and  $q(v)$  be the index queried at node  $v$ . Furthermore, we associate to each  $v \in \mathcal{N}(D)$  a number  $R_\mu^D(v)$  inductively. If  $v$  is the root of  $D$ , we let  $R_\mu^D(v) = 1$  and if  $v$  is the child of  $w$  which is reached when the query answer to  $q(w)$  is  $b \in \{0, 1\}$ :

$$R_\mu^D(v) = R_\mu^D(w) \cdot \min \left\{ \Pr_{x \sim \mu^0|_w} [x_{q(w)} = b], \Pr_{x \sim \mu^1|_w} [x_{q(w)} = b] \right\}$$

Finally, we define  $\Delta_\mu^D(v)$  for each  $v \in \mathcal{N}(D)$  with:

$$\Delta_\mu^D(v) := |\Pr_{x \sim \mu^0|_w} [x_{q(v)} = 0] - \Pr_{x \sim \mu^1|_w} [x_{q(v)} = 0]|$$

$R_\mu^D(v)$  can be interpreted as the probability of reaching node  $v$  in a random walk that starts at the root and with probability  $\min\{\Pr_{x \sim \mu^0|_v} [x_i = 0], \Pr_{x \sim \mu^1|_v} [x_i = 0]\}$  moves left, with probability  $\min\{\Pr_{x \sim \mu^0|_v} [x_i = 1], \Pr_{x \sim \mu^1|_v} [x_i = 1]\}$  moves right and with remaining probability  $\Delta_\mu^D(v)$  stops. As such, it holds that  $\sum_{v \in \mathcal{N}(D)} \Delta_\mu^D(v) R_\mu^D(v) = 1$  and that for any partition  $\Gamma$  of  $\{0, 1\}^n$  we have  $\sum_{\gamma \in \Gamma} R_\mu^D(\gamma) \leq 1$ . The max conflict complexity  $\bar{\chi}(f)$  is defined as:

$$\bar{\chi}(f) := \max_{\mathcal{Q}} \min_{D \in \mathcal{D}(f)} \mathbb{E}_{\mu \sim \mathcal{Q}} \left[ \sum_{v \in \mathcal{N}(D)} |v| \Delta_\mu^D(v) R_\mu^D(v) \right]$$

Where  $\mathcal{Q}$  ranges over distributions of pairs of distributions over  $f^{-1}(0)$  and  $f^{-1}(1)$  and  $\mathcal{D}(f)$  is the set of all decision tree solving  $f$  correctly.

**Lemma 17.**  $\bar{\chi}(F) \leq O(n^{1/2})$

*Proof.* See the full version of this article [BDBGM22].  $\square$

*Acknowledgements:* We thank anonymous FOCS reviewers for their comments.

### REFERENCES

- [ABK16] Scott Aaronson, Shalev Ben-David, and Robin Kothari. Separations in query complexity using cheat sheets. In *Proceedings of the 48th Symposium on Theory of Computing (STOC)*, pages 863–876. ACM, 2016. doi:10.1145/2897518.2897644.
- [AGJ<sup>+</sup>18] 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 *Proceedings of the 37th Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 93, pages 10:1–10:13, 2018. doi:10.4230/LIPIcs.FSTTCS.2017.10.
- [AKK16] Andris Ambainis, Martins Kokainis, and Robin Kothari. Nearly optimal separations between communication (or query) complexity and partitions. In *Proceedings of the 31st Computational Complexity Conference (CCC)*, pages 4:1–4:14. Schloss Dagstuhl, 2016. doi:10.4230/LIPIcs.CCC.2016.4.
- [BB19] Eric Blais and Joshua Brody. Optimal separation and strong direct sum for randomized query complexity. In *Proceedings of the 34th Computational Complexity Conference (CCC)*, pages 29:1–29:17. Schloss Dagstuhl, 2019. doi:10.4230/LIPIcs.CCC.2019.29.
- [BB20a] Shalev Ben-David and Eric Blais. A new minimax theorem for randomized algorithms. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*. IEEE, nov 2020. doi:10.1109/focs46700.2020.00045.
- [BB20b] Shalev Ben-David and Eric Blais. A tight composition theorem for the randomized query complexity of partial functions. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*. IEEE, nov 2020. doi:10.1109/focs46700.2020.00031.
- [BDBGM22] Shalev Ben-David, Eric Blais, Mika Göös, and Gilbert Maystre. Randomised composition and small-bias minimax. 2022. URL: <https://eccc.weizmann.ac.il/report/2022/112/>.
- [BDG<sup>+</sup>20] Andrew Basilakis, Andrew Drucker, Mika Göös, Lunjia Hu, Weiyun Ma, and Li-Yang Tan. The power of many samples in query complexity. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 168, pages 9:1–9:18. Schloss Dagstuhl, 2020. doi:10.4230/LIPIcs.ICALP.2020.9.
- [BdW02] Harry Buhrman and Ronald de Wolf. Complexity measures and decision tree complexity: a survey. *Theoretical Computer Science*, 288(1):21–43, 2002. Complexity and Logic. doi:10.1016/S0304-3975(01)00144-X.
- [BGKW20] Shalev Ben-David, Mika Göös, Robin Kothari, and Thomas Watson. When is amplification necessary for composition in randomized query complexity? In *Proceedings of the 24th International Conference on Randomization and Computation (RANDOM)*, volume 176, pages 28:1–28:16. Schloss Dagstuhl, 2020. doi:10.4230/LIPIcs.APPROX/RANDOM.2020.28.
- [BK16] Shalev Ben-David and Robin Kothari. Randomized query complexity of sabotaged and composed functions. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 55,

- pages 60:1–60:14, 2016. doi:10.4230/LIPIcs.ICALP.2016.60.
- [BKLS20] Joshua Brody, Jae Tak Kim, Peem Lerdputtipongporn, and Hariharan Srinivasulu. A strong XOR lemma for randomized query complexity, 2020. arXiv:2007.05580.
- [DM21] Yogesh Dahiya and Meena Mahajan. On (simple) decision tree rank. In *Proceedings of the 41st Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 213, pages 15:1–15:16. Schloss Dagstuhl, 2021. doi:10.4230/LIPIcs.FSTTCS.2021.15.
- [GJPW18] Mika Göös, T. S. Jayram, Toniann Pitassi, and Thomas Watson. Randomized communication versus partition number. *ACM Transactions on Computation Theory*, 10(1), 2018. doi:10.1145/3170711.
- [GLSS19] Dmitry Gavinsky, Troy Lee, Miklos Santha, and Swagato Sanyal. A composition theorem for randomized query complexity via max-conflict complexity. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 132, pages 64:1–64:13, 2019. doi:10.4230/LIPIcs.ICALP.2019.64.
- [GM21] Mika Göös and Gilbert Maystre. A majority lemma for randomised query complexity. In *Proceedings of the 36th Computational Complexity Conference (CCC)*, volume 200, pages 18:1–18:15. Schloss Dagstuhl, 2021. doi:10.4230/LIPIcs.CCC.2021.18.
- [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. doi:10.1007/s00493-014-3189-x.
- [GTW21] Uma Girish, Avishay Tal, and Kewen Wu. Fourier Growth of Parity Decision Trees. In *Proceedings of the 36th Computational Complexity Conference (CCC)*, volume 200, pages 39:1–39:36. Schloss Dagstuhl, 2021. doi:10.4230/LIPIcs.CCC.2021.39.
- [JKS10] Rahul Jain, Hartmut Klauck, and Miklos Santha. Optimal direct sum results for deterministic and randomized decision tree complexity. *Information Processing Letters*, 110(20):893–897, 2010. doi:10.1016/j.ipl.2010.07.020.
- [Li21] Yaqiao Li. Conflict complexity is lower bounded by block sensitivity. *Theoretical Computer Science*, 856:169–172, feb 2021. doi:10.1016/j.tcs.2020.12.038.
- [LMR<sup>+</sup>11] Troy Lee, Rajat Mittal, Ben Reichardt, Robert Spalek, and Mario Szegedy. Quantum query complexity of state conversion. In *Proceedings of the 52nd Symposium on Foundations of Computer Science (FOCS)*, pages 344–353. IEEE, 2011. doi:10.1109/FOCS.2011.75.
- [Rei11] Ben Reichardt. Reflections for quantum query algorithms. In *Proceedings of the 22nd Symposium on Discrete Algorithms (SODA)*, pages 560–569, 2011.
- [Sav02] Petr Savický. On determinism versus unambiguous nondeterminism for decision trees. Technical Report TR02-009, Electronic Colloquium on Computational Complexity (ECCC), 2002. URL: <http://eccc.hpi-web.de/report/2002/009/>.
- [She12] Alexander A. Sherstov. The communication complexity of gap hamming distance. *Theory of Computing*, 8(8):197–208, 2012. URL: <http://www.theoryofcomputing.org/articles/v008a008>, doi:10.4086/toc.2012.v008a008.
- [Tal13] Avishay Tal. Properties and applications of boolean function composition. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 441–454, 2013. doi:10.1145/2422436.2422485.
- [Ver98] Nikolai Vereshchagin. Randomized boolean decision trees: Several remarks. *Theoretical Computer Science*, 207(2):329–342, nov 1998. doi:10.1016/s0304-3975(98)00071-1.
- [Yao77] Andrew Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of the 18th Symposium on Foundations of Computer Science (FOCS)*, pages 222–227, Oct 1977. doi:10.1109/SFCS.1977.24.