

THE COMMUNICATION COMPLEXITY OF SET INTERSECTION AND MULTIPLE EQUALITY TESTING*

DAWEI HUANG[†], SETH PETTIE[‡], YIXIANG ZHANG[§], AND ZHIJUN ZHANG[§]

Abstract. In this paper we explore fundamental problems in randomized communication complexity such as computing `SetIntersection` on sets of size k and `EqualityTesting` between vectors of length k . Sağlam and Tardos [*Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, 2013, pp. 678–687] and Brody et al. [*Algorithmica*, 76 (2016), pp. 796–845] showed that for these types of problems, one can achieve optimal communication volume of $O(k)$ bits, with a randomized protocol that takes $O(\log^* k)$ rounds. They also proved that this is one point along the optimal round-communication trade-off curve. Aside from rounds and communication volume, there is a *third* parameter of interest, namely the *error probability* p_{err} , which we write 2^{-E} . It is straightforward to show that protocols for `SetIntersection` or `EqualityTesting` need to send at least $\Omega(k + E)$ bits, regardless of the number of rounds. Is it possible to simultaneously achieve optimality in all three parameters, namely $O(k + E)$ communication and $O(\log^* k)$ rounds? In this paper we prove that there is no universally optimal algorithm, and we complement the existing round-communication trade-offs [M. Sağlam and G. Tardos, *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, 2013, pp. 678–687; J. Brody et al., *Algorithmica*, 76 (2016), pp. 796–845] with a new trade-off between rounds, communication, and probability of error. In particular, any protocol for solving multiple `EqualityTesting` in r rounds with failure probability $p_{\text{err}} = 2^{-E}$ has communication volume $\Omega(Ek^{1/r})$. We present several algorithms for multiple `EqualityTesting` (and its variants) that match or nearly match our lower bound and the lower bound of [M. Sağlam and G. Tardos, *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, 2013, pp. 678–687; J. Brody et al., *Algorithmica*, 76 (2016), pp. 796–845]. Lower bounds on `EqualityTesting` extend to `SetIntersection` for every r, k , and p_{err} (which is trivial); in the reverse direction, we prove that upper bounds on `EqualityTesting` for r, k, p_{err} imply similar upper bounds on `SetIntersection` with parameters $r + 1, k$, and p_{err} . Our original motivation for considering p_{err} as an independent parameter came from the problem of enumerating triangles in distributed (CONGEST) networks having maximum degree Δ . We prove that this problem can be solved in $O(\Delta/\log n + \log \log \Delta)$ time with high probability $1 - 1/\text{poly}(n)$. This beats the trivial (deterministic) $O(\Delta)$ -time algorithm and is superior to the $\tilde{O}(n^{1/3})$ algorithm of [Y. Chang, S. Pettie, and H. Zhang, *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2019, pp. 821–840; Y. Chang and T. Saranurak, *Proceedings of the ACM Symposium on Principles of Distributed Computing*, 2019, pp. 66–73] when $\Delta = \tilde{O}(n^{1/3})$.

Key words. communication complexity, set disjointness, distributed triangle detection, lower bounds

AMS subject classifications. 94A05, 68R05, 60Cxx

DOI. 10.1137/20M1326040

1. Introduction. Communication complexity was defined by Yao [45] in 1979 and has become an indispensable tool for proving lower bounds in models of computation in which the notions of *parties* and *communication* are not direct. See, e.g., books and monographs [40, 38, 34] and surveys [15, 35] on the subject. In this paper we consider some of the most fundamental and well-studied problems in this model,

*Received by the editors March 18, 2020; accepted for publication (in revised form) January 19, 2021; published electronically April 13, 2021. An extended abstract of this paper was presented at the 31st ACM-SIAM Symposium on Discrete Algorithms (SODA 2020), pp. 1715–1732.

<https://doi.org/10.1137/20M1326040>

Funding: This work was supported by NSF grants CCF-1514383, CCF-1637546, and CCF-1815316.

[†]University of Michigan, Ann Arbor, MI 48109 USA (dwhuang9@gmail.com).

[‡]EECS, University of Michigan, Ann Arbor, MI 48109 USA (pettie@umich.edu).

[§]Tsinghua University, Beijing, China (zhangyix16@mails.tsinghua.edu.cn, zhijunz@princeton.edu).

such as **SetDisjointness**, **SetIntersection**, **ExistsEqual**, and **EqualityTesting**. Let us briefly define these problems formally since the terminology is not completely standard.

SetDisjointness and SetIntersection. In the **SetDisjointness** problem Alice and Bob receive sets $A \subset U$ and $B \subset U$ where $|A|, |B| \leq k$ and must determine whether $A \cap B = \emptyset$. Define $\text{SetDisj}(k, r, p_{\text{err}})$ to be the minimum communication complexity of an r -round randomized protocol for this problem that errs with probability at most p_{err} . We can assume that $|U| = O(k^2/p_{\text{err}})$ without loss of generality.¹ The input to the **SetIntersection** problem is the same, except that the parties must *report the entire set* $A \cap B$. Define $\text{SetInt}(k, r, p_{\text{err}})$ to be the minimum communication complexity of an r -round protocol for **SetIntersection**.

EqualityTesting and ExistsEqual. In the **EqualityTesting** problem Alice and Bob hold vectors $\mathbf{x} \in U^k$ and $\mathbf{y} \in U^k$ and must determine, for each index $i \in [k]$, whether $x_i = y_i$ or $x_i \neq y_i$. A potentially easier version of the problem, **ExistsEqual** (aka **BlockEquality** and **OR-Equality** [11, 26]), is to determine if there *exists* at least one index $i \in [k]$ for which $x_i = y_i$. Define $\text{Eq}(k, r, p_{\text{err}})$ to be the randomized communication complexity of any r -round protocol for **EqualityTesting** that errs with probability p_{err} , and $\exists \text{Eq}(k, r, p_{\text{err}})$ the corresponding complexity of **ExistsEqual**. Once again, we can assume that $|U| = O(k/p_{\text{err}})$ without loss of generality.

Mehlhorn and Schmidt [36] proved that the *deterministic* complexity of **ExistsEqual** is $\Omega(k \log U)$, and in fact the trivial one-round algorithm is optimal for all four problems [34]. Thus, researchers have focused on the *randomized* communication complexity of these problems [28, 39, 31, 21, 18, 37, 12, 41, 11, 6, 27]. Most prior work analyzed the relationship between *round complexity* and *communication volume* and either treated p_{err} as a constant or dependent on k (e.g., $1/\text{poly}(k)$) or considered the regime when $p_{\text{err}} = 1/2 - \epsilon$ is close to chance guessing [6, 27]. In this paper we view $p_{\text{err}} = 2^{-E}$ as independent of k and close to zero.

History. Håstad and Wigderson [28] gave an $O(\log k)$ -round protocol for **SetDisjointness** in which Alice and Bob communicate $O(k)$ bits, which matched an $\Omega(k)$ lower bound of Kalyanasundaram and Schnitger [31]; see also [39, 12, 18]. Feder et al. [21] proved that **EqualityTesting** can be solved with $O(k)$ communication by an $O(\sqrt{k})$ -round protocol that errs with probability $\exp(-\sqrt{k})$. Nikishkin [37] later improved their round complexity and error probability to $\log k$ and $\exp(-k/\text{polylog}(k))$, respectively. Improving [28], Sağlam and Tardos [41] gave an r -round protocol for **SetDisjointness** that uses $O(k \log^{(r)} k)$ communication, where $\log^{(r)}$ is the r -fold iterated logarithm function. For $r = \log^* k$ the error probability of this protocol is $\exp(-\sqrt{k})$, coincidentally matching [21]. In independent work, Brody et al. [11] gave r -round and $O(r)$ -round protocols for **ExistsEqual** and **SetIntersection**, respectively, that use $O(k \log^{(r)} k)$ communication and err with probability $1/\text{poly}(k)$.

Sağlam and Tardos [41] were the first to show that this $O(k \log^{(r)} k)$ round versus communication trade-off is optimal, using a combinatorial round-elimination technique. In particular, this lower bound even applies to an **ExistsEqual** protocol with constant error probability. Independently, Brody et al. [10, 11] established the same lower bound trade-off for **ExistsEqual**, assuming the error probability is at most $1/\text{poly}(k)$. Brody et al. [11] also introduced a *randomized* reduction from **SetIntersection** to

¹Before the first round of communication, pick a pairwise independent $h : U \mapsto [O(k^2/p_{\text{err}})]$ and check whether $h(A) \cap h(B) = \emptyset$ with error probability $p_{\text{err}}/2$. Thus, having **SetDisj** depend additionally on $|U|$ is somewhat redundant, at least when $|U|$ is large.

TABLE 1

Upper and lower bounds on *SetDisjointness*, *SetIntersection*, *EqualityTesting*, and *ExistsEqual*. Via trivial reductions, lower bounds on *ExistsEqual* extend to all four problems, and upper bounds on *SetIntersection* extend to all four problems. From Theorem 1, the upper bounds on *SetIntersection* and *SetDisjointness* follow from those of *EqualityTesting* and *ExistsEqual*, respectively, +1 round of communication. The log-star function is defined as $\log^*(x) = \min\{i : \log^{(i)}(x) \leq 1\}$, e.g., $\log^*(k/E) = 0$ if $E \geq k$.

Problem	Commun.	Rounds	Error Prob.	Notes
EqualityTesting	$O(k)$	$O(\sqrt{k})$	$\exp(-\sqrt{k})$	[21]
EqualityTesting	$O(k)$	$\log k$	$\exp\left(-\frac{k}{\text{polylog}(k)}\right)$	[37]
SetDisjointness	$O(k)$	$O(\log k)$	Constant	[28]
SetDisjointness	$O(k \log^{(r)} k)$	r	$\geq \exp(-\sqrt{k})$	[41]
ExistsEqual	$O(k \log^{(r)} k)$	r	$1/\text{poly}(k)$	[11]
SetIntersection		$O(r)$		
ExistsEqual SetDisjointness	$O(k + Ek^{1/r})$	$r + \log^*(k/E)$	2^{-E}	new
EqualityTesting and [SetIntersection]	$O(k + Ek^{1/r} \cdot \log r + Er \log r)$	$r + \log^*(k/E) + 1$		
	$O(k + E)$	$\log k$		
Lower Bounds				
SetDisjointness	$\Omega(\sqrt{k})$	∞	Constant	[2]
SetDisjointness	$\Omega(k)$	∞	Constant	[31]
SetDisjointness	$\Omega(\epsilon k)$	∞	$1/2 - \epsilon$	[6, 27]
ExistsEqual	$\Omega(k \log^{(r)} k)$	r	Constant	[41]
ExistsEqual	$\Omega(k \log^{(r)} k)$	r	$1/\text{poly}(k)$	[11]
ExistsEqual	$\Omega(Ek^{1/r})$	r	2^{-E}	new

EqualityTesting, which errs with probability $\exp(-\tilde{O}(\sqrt{k}))$, i.e., it cannot be applied when the desired total error probability p_{err} is sufficiently small.

One consequence of [31] is that solving SetDisjointness with success probability $1/2 + \epsilon$ (i.e., slightly better than chance) requires $\Omega(\epsilon^2 n)$ communication. Braverman and Moitra [6] improved this to $\Omega(\epsilon n)$. See Göös and Watson [27] for a generalized form of this lower bound.

Related work. EqualityTesting is the problem of simultaneously solving k independent instances of equality, which fits within the direct sum/direct product framework of many results in communication complexity. However, the existing results on direct sums and amortized communication complexity are largely inapplicable to our setting. Braverman and Rao [7] and Braverman [5] bound the amortized communication cost of a problem in terms of its information cost ($O(1)$ in the case of equality). However, their error parameter ρ is fixed (as $k \rightarrow \infty$) and reflects the probability that the protocol errs on each *individual* coordinate. It does not seem possible to extend these results to variable $\rho = \rho(k)$; see [5, Remark 6.2] and [43, Remark 3.3]. See [21, 3] for general lower bounds on direct sum problems. Direct product theorems are typically of the following form. If a protocol has success probability at most $2/3$ using communication C , then solving n independent instances with Cn communication has success probability exponentially small in n . See Klauck [32] for such a direct product theorem specifically for SetDisjointness. See Braverman et al. [8] for a weaker $\tilde{\Omega}(C\sqrt{n})$ lower bound applicable to any function and Jain, Pereszlényi, and Yao [30] for a bounded-round direct product theorem applicable to any function. These results do not imply any new lower bounds on EqualityTesting or related problems.

1.1. Contributions. First, we observe that a simple *deterministic* reduction shows that, up to one round of communication, `SetIntersection` is equivalent to `EqualityTesting` for any p_{err} , and `SetDisjointness` is equivalent to `ExistsEqual` for any p_{err} . Theorem 1 is proved in Appendix A; it is inspired by the randomized reduction of Brody et al. [11].

THEOREM 1. *For any parameters $k \geq 1, r \geq 1$, and $p_{\text{err}} = 2^{-E}$, it holds that*

$$\begin{aligned} \text{Eq}(k, r, p_{\text{err}}) &\leq \text{SetInt}(k, r, p_{\text{err}}), & \text{SetInt}(k, r + 1, p_{\text{err}}) &\leq \text{Eq}(k, r, p_{\text{err}}) + \zeta, \\ \exists\text{Eq}(k, r, p_{\text{err}}) &\leq \text{SetDisj}(k, r, p_{\text{err}}), & \text{SetDisj}(k, r + 1, p_{\text{err}}) &\leq \exists\text{Eq}(k, r, p_{\text{err}}) + \zeta, \end{aligned}$$

where $\zeta = O(k + \log E)$.

Second, we prove that in any of the four problems, it is impossible to simultaneously achieve communication volume $O(k + E)$ in $O(\log^* k)$ rounds for all $k, p_{\text{err}} = 2^{-E}$. Specifically, any r -round protocol needs $\Omega(Ek^{1/r})$ communication. Whereas the implication of [41, 11] is that optimal $O(k)$ communication is possible only with $\Omega(\log^* k)$ rounds, the implication of our work is that optimal communication $O(k + E)$ is possible only with $\Omega(\log k)$ rounds, whenever $E \geq k$.

We complement our lower bounds with matching or nearly matching upper bounds. First, we show that in any `EqualityTesting/ExistsEqual` instance with $E < k$, one can, with probability $1 - 2^{-\Theta(E)}$, reduce the effective number of coordinates to E using $O(k)$ communication and $\log^*(k/E)$ rounds. Thus, we can simplify the following discussion by assuming that $E \geq k$.

We give *four* distinct protocols, the first of which solves `EqualityTesting` with $O(rEk^{1/r})$ communication, which is optimal whenever $r = O(1)$. The remaining three protocols attempt to get rid of this extraneous r factor in different situations. Our second protocol shows that it *is* possible to achieve $O(Ek^{1/r})$ complexity, but for the slightly simpler problem of `ExistsEqual`. Our third protocol shows that with $O(r)$ rounds (instead of r rounds) it is possible to achieve $O(Ek^{1/r})$ communication. In particular, absolutely optimal communication $O(E)$ is possible with $\log k = O(r)$ rounds.

Our first three protocols show that the optimal round-communication-error trade-off for `EqualityTesting` can be achieved whenever $r = O(1)$ or $r = \Omega(\log k)$, or for any r in the case of `ExistsEqual`. The remaining problem (`EqualityTesting` in r rounds, r between $\omega(1)$ and $o(\log k)$) seems to be quite difficult. Our fourth protocol solves `EqualityTesting` with $O(Ek^{1/r} \log r + Er \log r)$ communication, which for $r \in [1, \log k / \log \log k]$ is dominated by the first term and therefore within a $\log r \leq \log \log k$ factor of optimal. A close study of our second and fourth protocols reveals a key distinction between `EqualityTesting` from `ExistsEqual`, which is only relevant when the probability of error is quite small (e.g., $E \geq k$). It is *plausible* that `EqualityTesting` is asymptotically harder than `ExistsEqual` for many values of r and, for similar reasons, that `SetIntersection` is asymptotically harder than `SetDisjointness`.

Our original interest in `SetIntersection` came from distributed subgraph detection in `CONGEST`² networks, which has garnered significant interest in recent years [14, 13, 29, 1, 19, 33, 23, 17, 25]. Izumi and LeGall [29] proved that triangle enumeration³

²In the `CONGEST` model there is a graph $G = (V, E)$ whose vertices are identified with processors and whose edges represent bidirectional communication links. Each vertex v does not know G and is only initially aware of an $O(\log n)$ -bit $\text{ID}(v)$, $\text{deg}(v)$, and global parameters $n \geq |V|$ and $\Delta \geq \max_{u \in V} \text{deg}(u)$. Communication proceeds in synchronized rounds; in each round, each processor can send a (different) $O(\log n)$ -bit message to each of its neighbors.

³Every triangle (3-cycle) in G must be reported by *some* vertex.

requires $\Omega(n^{1/3}/\log n)$ rounds in the CONGEST model and further showed that *local* triangle enumeration⁴ requires $\Omega(\Delta/\log n)$ rounds in CONGEST, which can be as large as $\Omega(n/\log n)$.

The most natural way to solve (local) triangle enumeration is, for every edge $\{u, v\} \in E(G)$, to have u and v run a two-party `SetIntersection` protocol in which they compute $N(u) \cap N(v)$, where $N(u) = \{\text{ID}(x) \mid \{u, x\} \in E(G)\}$ and $\text{ID}(x) \in \{0, 1\}^{O(\log n)}$ is x 's unique identifier. Any r -round protocol with communication volume $O(\Delta)$ can be simulated in CONGEST in $O(\Delta/\log n + r)$ rounds since the message size is $O(\log n)$ bits. However, to guarantee a *global* probability of success at least $1 - 1/\text{poly}(n)$, the failure probability of each `SetIntersection` instance must be $p_{\text{err}} = 2^{-E}$, $E = \Theta(\log n)$, which is independent of Δ . Our communication complexity lower bound suggests that to achieve this error probability, we would need $\Omega((\Delta + E\Delta^{1/r})/\log n + r)$ CONGEST rounds, i.e., with $r = \log \Delta$ we should not be able to do better than $O(\Delta/\log n + \log \Delta)$. We prove that (local) triangle enumeration can actually be solved *exponentially* faster, in $O(\Delta/\log n + \log \log \Delta)$ CONGEST rounds, without *necessarily* solving every `SetIntersection` instance.

Organization. The proof of Theorem 1 on the near-equivalence of `SetIntersection/``SetDisjointness` and `EqualityTesting/ExistsEqual` appears in Appendix A. Section 2 reviews concepts from information theory and communication complexity. In section 3 we present new lower bounds for both `EqualityTesting` and `ExistsEqual` that incorporate rounds, communication, and error probability. Section 4 presents nearly matching upper bounds for `EqualityTesting` and `ExistsEqual`, and section 5 applies them to the distributed triangle enumeration problem. We conclude with some open problems in section 6.

2. Preliminaries.

2.1. Notational conventions. The set of positive integers at most t is denoted $[t]$. Random variables are typically written as capital letters (X, Y, M , etc.) and the values they take on are lower case (x, y, m , etc.). The letters p, q, μ, \mathcal{D} are reserved for probability mass functions (*p.m.f.*). For example, $\mathcal{D}(x)$ denotes the probability that $X = x$ whenever $X \sim \mathcal{D}$. The *support* $\text{supp}(\mathcal{D})$ of a distribution \mathcal{D} is the set of all x for which $\mathcal{D}(x) > 0$. If $\mathcal{X} \subseteq \text{supp}(\mathcal{D})$, $\mathcal{D}(\mathcal{X}) = \sum_{x \in \mathcal{X}} \mathcal{D}(x)$.

Many of our random variables are vectors. If x is a k -dimensional vector and $I \subseteq [k]$, x_I is the projection of x onto the coordinates in I and x_i is short for $x_{\{i\}}$. Similarly, if \mathcal{D} is the p.m.f. of a k -dimensional random variable, \mathcal{D}_I is the marginal distribution of \mathcal{D} on the index set $I \subseteq [k]$.

Throughout the paper, \log and \exp are the base-2 logarithm and exponential functions, and $\log^{(r)}$ and $\exp^{(r)}$ their r -fold iterated versions:

$$\begin{aligned}\log^{(0)}(x) &= \exp^{(0)}(x) = x, \\ \log^{(r)}(x) &= \log(\log^{(r-1)}(x)), \\ \exp^{(r)}(x) &= \exp(\exp^{(r-1)}(x)).\end{aligned}$$

The log-star function is defined to be $\log^*(x) = \min\{r \mid \log^{(r)}(x) \leq 1\}$. In particular, $\log^*(x) = 0$ if $x \leq 1$.

⁴Every triangle in G must be reported by at least one of the three constituent vertices. Izumi and LeGall [29] only stated the $\Omega(n/\log n)$ lower bound but it can also be expressed in terms of Δ .

2.2. Information theory. The most fundamental concept in information theory is *Shannon entropy*. The Shannon entropy of a discrete random variable X is defined as

$$H(X) = - \sum_{x \in \text{supp}(X)} \Pr[X = x] \log \Pr[X = x].$$

Since there may be cases in which different distributions are defined for the “same” random variable, we use $H(p)$ in place of $H(X)$ if X is drawn from a p.m.f. p . We also write $H(\alpha)$, $\alpha \in (0, 1)$, to be the entropy of a Bernoulli random variable with success probability α . In general, we freely use a random variable and its p.m.f. interchangeably.

The *joint entropy* $H(X, Y)$ of two random variables X and Y is simply

$$H(X, Y) = - \sum_{x \in \text{supp}(X)} \sum_{y \in \text{supp}(Y)} \Pr[X = x \wedge Y = y] \log \Pr[X = x \wedge Y = y].$$

This notion can be easily extended to cases of more than two random variables. Here, we state a well-known fact about joint entropy.

FACT 2.1. *For any random variables X_1, X_2, \dots, X_n , their joint entropy is at most the sum of their individual entropies, i.e., $H(X_1, X_2, \dots, X_n) \leq \sum_{i=1}^n H(X_i)$.*

The *conditional entropy* of Y conditioned on another random variable X , denoted $H(Y | X)$, measures the expected amount of extra information required to fully describe Y if X is known. It is defined to be

$$\begin{aligned} H(Y | X) &= H(X, Y) - H(X) \\ &= - \sum_{x \in \text{supp}(X)} \Pr[X = x] \sum_{y \in \text{supp}(Y)} \Pr[Y = y | X = x] \log \Pr[Y = y | X = x] \\ &\geq 0, \end{aligned}$$

which can be viewed as a weighted sum of entropies of a number of conditional distributions.

Finally, the *mutual information* $I(X ; Y)$ between two random variables X and Y quantifies the amount of information that is revealed about one random variable through knowing the other one:

$$\begin{aligned} I(X ; Y) &= H(X) - H(X | Y) \\ &= H(X) + \sum_{y \in \text{supp}(Y)} \Pr[Y = y] \sum_{x \in \text{supp}(X)} \Pr[X = x | Y = y] \log \Pr[X = x | Y = y]. \end{aligned}$$

2.3. Communication complexity. Let $f(x, y)$ be a function over domain $\mathcal{X} \times \mathcal{Y}$, and consider any two-party communication protocol $Q(x, y)$ that computes $f(x, y)$, where one party holds x and the other holds y . The *transcript* of Q on (x, y) is defined to be the concatenation of all messages exchanged by the two parties, in order, as they execute on input (x, y) . The receiver of the last message outputs the answer, which is denoted $Q(x, y)$. The *communication cost* of Q is the maximum transcript length produced by Q over all possible inputs.

Let Q_d be a deterministic protocol for f and suppose μ is a distribution over $\mathcal{X} \times \mathcal{Y}$. The *distributional error probability* of Q_d with respect to μ is the probability $\Pr_{(x,y) \sim \mu}[Q_d(x, y) \neq f(x, y)]$. For any $0 < \epsilon < 1$, the (μ, ϵ) -*distributional deterministic*

communication complexity of the function f is the minimum communication cost of any protocol Q_d that has distributional error probability at most ϵ with respect to μ .

A randomized protocol $Q_r(x, y, w)$ also takes a public random string $w \sim \mathcal{W}$ as input. The error probability of Q_r is calculated as $\max_{(x,y) \in \mathcal{X} \times \mathcal{Y}} \Pr_{w \sim \mathcal{W}}[Q_r(x, y, w) \neq f(x, y)]$. The ϵ -*randomized communication complexity* of f is the minimum communication cost of Q_r over all protocols Q_r with error probability at most ϵ .

Yao's *minimax principle* [44] is a common starting point for lower bound proofs in randomized communication complexity. The easy direction of Yao's minimax principle states that the communication cost of the best deterministic protocol specific to any particular distribution is at most the communication cost of any randomized protocol on its worst case input.

LEMMA 2.2 (Yao's minimax principle [44]). *Let $f : \mathcal{X} \times \mathcal{Y} \mapsto \mathcal{Z}$ be the function to be computed. Let $D_{\mu, \epsilon}(f)$ be the (μ, ϵ) -distributional deterministic communication complexity of f , and let $R_\epsilon(f)$ be the ϵ -randomized communication complexity of f . Then for any $0 < \epsilon < 1/2$,*

$$\max_{\mu} D_{\mu, \epsilon}(f) \leq R_\epsilon(f).$$

Therefore, to show a lower bound on the ϵ -randomized communication complexity of a function f , it suffices to find a hard distribution μ on the input set and prove a lower bound for the communication cost of any deterministic protocol that has distributional error probability at most ϵ with respect to μ .

In this paper we consider randomized protocols for `EqualityTesting`, `ExistsEqual`, `SetIntersection`, and `SetDisjointness`. The receiver of the last message issues the output, which is a single bit in the case of the second and fourth problems, a vector in $\{=, \neq\}^k$ in the case of `EqualityTesting`, and the set $A \cap B$ in the case of `SetIntersection`. A protocol *errs* if any portion of its output is incorrect. (This is in contrast with general direct sum theorems [7, 5], where, for example, an `EqualityTesting` protocol would guarantee only that each *individual* component of the output vector is correct with probability $1 - p_{\text{err}}$.) Our lower bounds work for two-sided errors but our upper bounds have only one-sided errors, e.g., in the case of `EqualityTesting` we may only claim that a coordinate is “=” when the true answer is “≠” and never vice versa.

3. Lower bounds on `ExistsEqual` and `EqualityTesting`. In this section we prove lower bounds on `EqualityTesting` and `ExistsEqual`. Theorem 2 obviously follows directly from Theorem 3, but we prove them in that order nonetheless because Theorem 2 is a bit simpler.

THEOREM 2. *Any r -round randomized protocol for `EqualityTesting` on vectors of length k that errs with probability $p_{\text{err}} = 2^{-E}$ requires at least $\Omega(Ek^{1/r})$ bits of communication.*

THEOREM 3. *Any r -round randomized protocol for `ExistsEqual` on vectors of length k that errs with probability $p_{\text{err}} = 2^{-E}$ requires at least $\Omega(Ek^{1/r})$ bits of communication.*

Without any constraint on the number of rounds, `EqualityTesting` trivially requires $\Omega(k)$ communication. `ExistsEqual` also requires $\Omega(k)$ communication, through a small modification to the `SetDisjointness` lower bounds [31, 39]. Even when $k = 1$, we need at least $\Omega(E)$ communication to solve `EqualityTesting/ExistsEqual` with error probability 2^{-E} [34]. Thus, we can assume that $E = \Omega(k^{1-1/r})$, $k^{1/r} = \Omega(1)$, and hence $r = O(\log k)$. For example, some calculations later in our proof hold when $r \leq (\log k)/6$. When proving Theorem 3, we will further assume $E = \Omega(\log k)$ when

$r = 1$, which is reasonable because of Sağlam and Tardos' $\Omega(k \log^{(r)} k) = \Omega(k \log k)$ lower bound [41].

3.1. Structure of the proof. We consider deterministic strategies for ExistsEqual/EqualityTesting when Alice and Bob pick their input vectors independently from the uniform distribution on $[t]^k$, where $t = 2^{cE}$ and $c = 1/2$. Although the probability of seeing a collision in any particular coordinate is small, it is still much larger than the tolerable error probability (since $c < 1$), so it is incorrect to declare “ \neq ” in every coordinate without performing any communication.

We suppose, for the purpose of obtaining a contradiction, that there is a protocol for EqualityTesting with error probability 2^{-E} and communication complexity $c'Ek^{1/r}$, where $c' = c/100$. The length of the j th message is l_j , which could depend on the parameters $(E, r, k, \text{etc.})$ and possibly in some complicated way on the transcript of the protocol before round j .⁵

Our proof must necessarily consider transcripts of the protocol that are extremely unlikely (occurring with probability close to 2^{-E}) and also maintain a high level of uncertainty about *which* coordinates of Alice's and Bob's vectors might be equal. Consider the first message. Alice picks her input vector $x \in [t]^k$, which dictates the first message m_1 . Suppose, for simplicity, that it betrays exactly $l_1/k < c'Ek^{1/r-1}$ bits of information per coordinate of x . Before Bob can respond with a message m_2 he must commit to his input, say, y . Most values of y result in “good” outcomes: nearly all nonequal coordinates get detected immediately and the effective size of the problem is dramatically reduced. We are not interested in these values of y , only very “bad” values. Let I_1 be the first $k^{1-1/r}$ coordinates (or, more generally, $k^{1-1/r}$ coordinates that m_1 revealed below-average information about). With probability about $(2^{-c'Ek^{1/r-1}})^{|I_1|} = 2^{-c'E}$, Bob picks an input y that is *completely consistent* with Alice's on I_1 , i.e., as far as he can tell $y_i = x_i$ for every $i \in I_1$. Rather than sample y uniformly from $[t]^k$, we sample it from a “hybrid” distribution: y_{I_1} is sampled from the same distribution that m_1 revealed about x_{I_1} (forcing the above event to happen with probability 1), and $y_{[k] \setminus I_1}$ is sampled from Bob's former distribution (in this case, the uniform distribution on $[t]^{k-|I_1|}$), conditioned on the value of y_{I_1} .

This process continues round by round. Bob's message m_2 betrays at most $l_2/|I_1| < c'Ek^{2/r-1}$ bits of information on each coordinate of y_{I_1} , and there must be an index set $I_2 \subset I_1$ with $|I_2| = k^{1-2/r}$ such that, with probability around $2^{-c'E}$, it is completely consistent that $x_{I_2} = y_{I_2}$. Alice *resamples* her input so that this (rare) event occurs with probability 1, generates m_3 , and continues.

At the end of this process $|I_r| = k^{1-r/r} = 1$, and yet Alice and Bob have revealed less than the full cE bits of entropy about x_{I_r} and y_{I_r} . Regardless of whether they report “=” or “ \neq ” (on I_r), they are wrong with probability greater than 2^{-E} . Are we done? Absolutely not! The problem is that this strange process for sampling a possible transcript of the protocol might itself only find transcripts that occur with probability $\ll 2^{-E}$, making any conclusions we make about its (probability of) correctness moot. Generally speaking, we need to show that Alice's and Bob's actions are consistent with events that occur with probability $\gg 2^{-E}$.

Let us first make every step of the above process a bit more formal. In reality the inputs to Alice and Bob are *fixed* before the first round. However, it is much more

⁵In the context of ExistsEqual/EqualityTesting, it is natural to think about uniform-length messages, $l_j = c'Ek^{1/r}/r$, or lengths that decay according to some convergent series, e.g., $l_j \propto c'Ek^{1/r}/2^j$ or $l_j \propto c'Ek^{1/r}/j^2$.

useful to take the point of view of an external observer, who treats the inputs to Alice and Bob as *distributions* over vectors, that evolve as messages are transmitted.

- Before the j th round of communication, the sender of the j th message's input is drawn from a discrete distribution $\widehat{\mathcal{D}}^{(j-1)}$ over $[t]^k$. The receiver of the j th message's input is drawn from the distribution $\mathcal{D}^{(j-1)}$. For example, when $j = 1$, if Alice speaks first, then her initial distribution, $\widehat{\mathcal{D}}^{(0)}$, and Bob's initial distribution, $\mathcal{D}^{(0)}$, are both uniform over $[t]^k$.
- Before the j th round of communication both parties are aware of an index set I_{j-1} such that, informally, (i) the distributions $\mathcal{D}_{I_{j-1}}^{(j-1)}$ and $\widehat{\mathcal{D}}_{I_{j-1}}^{(j-1)}$ are very similar, and in particular, it is consistent that their inputs are identical on I_{j-1} , and (ii) the messages transmitted so far reveal "average" or below-average information about coordinates I_{j-1} . For example, $I_0 = [k]$ and it is consistent with the empty transcript that Alice's and Bob's inputs are identical on every coordinate.
- The j th message is a random variable $M_j \in \{0, 1\}^{l_j}$. In order to pick an m_j according to the right distribution, the sender picks an input $x \sim \widehat{\mathcal{D}}^{(j-1)}$ which, together with the history m_1, \dots, m_{j-1} , determines m_j . The sender transmits m_j to the receiver and promptly forgets x . The sender's new distribution (i.e., $\widehat{\mathcal{D}}^{(j-1)}$, conditioned on $M_j = m_j$) is called $\mathcal{D}^{(j)}$.
- The distribution $\mathcal{D}^{(j)}$ may reveal information about the coordinates I_{j-1} in an irregular fashion. We find a subset $I_j \subset I_{j-1}$ of coordinates, $|I_j| = k^{1-j/r}$, for which the amount of information revealed by $\mathcal{D}_{I_j}^{(j)}$ is at most average. The receiver of m_j changes his input distribution to $\widehat{\mathcal{D}}^{(j)}$, which is defined so that it basically agrees with $\mathcal{D}_{I_j}^{(j)}$ and the marginal distribution $\widehat{\mathcal{D}}_{[k] \setminus I_j}^{(j)}$, conditioned on the value selected by $\mathcal{D}_{I_j}^{(j)}$, is identical to $\mathcal{D}_{[k] \setminus I_j}^{(j-1)}$.
- The reason $\mathcal{D}_{I_j}^{(j)}$ and $\widehat{\mathcal{D}}_{I_j}^{(j)}$ are not *identical* is due to two filtering steps. To generate $\widehat{\mathcal{D}}^{(j)}$, we remove points from the support that have tiny (but non-zero) probability, which may be too close to the error probability. Intuitively these rare events necessarily represent a small fraction of the probability mass. Second, we remove points from the support if the ratio of their probability occurring under $\mathcal{D}^{(j)}$ over $\mathcal{D}^{(j-1)}$ is too high. Intuitively, we want to conclude that if there is a high probability of an error occurring under $\mathcal{D}^{(j)}$, then the probability is also high under $\mathcal{D}^{(j-1)}$ (and by unrolling this further, under $\mathcal{D}^{(0)}$). This argument works only if the ratios are what we would expect, given how much information is being revealed about these coordinates by m_j . As a result of these two filtering steps, $\mathcal{D}_{I_j}^{(j)}(x_{I_j})$ and $\widehat{\mathcal{D}}_{I_j}^{(j)}(x_{I_j})$ differ by at most a constant factor, for any particular vector $x_{I_j} \in [t]^{|I_j|}$.

3.2. A lower bound on Equality Testing. We begin with two general lemmas about discrete probability distributions that play an important role in our proof.

Roughly speaking, Lemma 3.1 captures and generalizes the following intuition: Suppose p is a *high entropy* distribution on some universe U and q is obtained from p by conditioning on an event $\mathcal{X} \subseteq U$ such that $p(\mathcal{X})$ is large, say, some constant like $1/4$. If p 's entropy is close to $\log|U|$, then q 's entropy should not be much smaller than that of p . As our proof goes on round by round, we will constantly throw away part of the input distribution's support to meet certain conditions. It is Lemma 3.1 that guarantees that the input distributions continue to have relatively high entropy.

Lemma 3.2 comes into play because the error probability will be calculated backward in a round-by-round manner. Suppose the old distribution (p) has no extremely low probability point and the new distribution (q) has almost full entropy. Lemma 3.2 provides us with a useful tool to transfer a lower bound on the probability of any event w.r.t. q to a lower bound on the same event w.r.t. p . It can be seen as a version of Markov's inequality for Kullback–Leibler divergences.

Variants of Lemmas 3.1 and 3.2 have appeared before. See, for example, [9, Claim 3.3] and [11, Lemma 3.8]. For completeness we include proofs in Appendix B.

LEMMA 3.1. *Let p and q be distributions defined on a universe of size 2^s . Suppose both of the following properties are satisfied:*

- (1) *The entropy of p is $H(p) \geq s - g$, where $g \in [0, s]$.*
- (2) *There exists $\alpha \in (0, 1)$ such that $q(x) \leq p(x)/\alpha$ holds for every value $x \in \text{supp}(q)$.*

The entropy of q is lower bounded by

$$H(q) \geq s - g/\alpha - H(\alpha)/\alpha.$$

LEMMA 3.2. *Let p and q be distributions defined on a universe of size 2^s . Suppose both of the following properties are satisfied:*

- (1) *The entropy of q is $H(q) \geq s - g_1$, where $g_1 \in [0, s]$.*
- (2) *There exists $g_2 \geq 0$ such that $p(x) \geq 2^{-s-g_2}$ holds for every value $x \in \text{supp}(q)$.*

Then, for any $\alpha \in (0, 1)$,

$$\Pr_{x \sim q} \left[\frac{q(x)}{p(x)} > 2^{g_1/\alpha + g_2 - (1-\alpha) \log(1-\alpha)/\alpha} \right] \leq \alpha.$$

We are now ready to begin the proof of Theorem 2 proper. We sample a transcript m_1, m_2, \dots, m_r one message at a time. Before m_j is transmitted the sender's input is drawn from $\widehat{\mathcal{D}}^{(j-1)}$ and the receiver's input is drawn from $\mathcal{D}^{(j-1)}$. We assume that at this moment the transcript satisfies Invariant 3.3 (below), which says, informally, that the distributions $\widehat{\mathcal{D}}^{(j-1)}$ and $\mathcal{D}^{(j-1)}$ are very similar and that the parties do not know much about each other's inputs on coordinates I_{j-1} , where $|I_{j-1}| = k^{1-(j-1)/r}$.

Through round j the parties' input distributions evolve as follows. Let $\mu(m_j)$ be the probability that m_j is transmitted next, conditioned on m_1, \dots, m_{j-1} . Conditioned on m_j , the sender's distribution becomes $\mathcal{D}^{(j)}[m_j]$, denoted $\mathcal{D}^{(j)}$ if m_j is clear from context. The receiver examines m_j and then *changes* their distribution to be $\widehat{\mathcal{D}}^{(j)}$ so that it basically matches $\mathcal{D}^{(j)}$ when projected onto I_j and has few very low probability points. This will be explained in detail later.

Roadmap. Our goal is to show that there is a large set \mathcal{M}_j of messages where \mathcal{M}_j depends on the transcript m_1, \dots, m_{j-1} up to round j , such that after appending any $m_j \in \mathcal{M}_j$ to the transcript, $\mathcal{D}^{(j)}$ and $\widehat{\mathcal{D}}^{(j)}$ continue to satisfy Invariant 3.3. Lemma 3.4 defines a set of messages \mathcal{M}'_j with $\mu(\mathcal{M}'_j) \geq 1/2$ whose entropy-reduction on indices in I_{j-1} is about what we would expect, and Lemma 3.5 shows that a subset $\mathcal{M}_j \subset \mathcal{M}'_j$ of those messages with $\mu(\mathcal{M}_j) \geq 1/4$ allows us to find an index set $I_j \subseteq I_{j-1}$ with the right properties. After describing the procedure for constructing $\widehat{\mathcal{D}}^{(j)}$, Lemma 3.6 proves that Invariant 3.3 is satisfied after round j if $m_j \in \mathcal{M}_j$. Lemma 3.7 bounds the ratio $\widehat{\mathcal{D}}^{(j)}(x)/\mathcal{D}^{(j-1)}(x)$, repeated applications of which allow us to deduce lower bounds on the error probability before round 1 (when the inputs are uniform in $[t]^k$) as a function on the error probability after round r . Putting it all together, Lemmas 3.8 and 3.9 derive lower bounds on the error probability of EqualityTesting and ExistsEqual protocols, respectively, that use too little communication.

INVARIANT 3.3. After round $j \in [0, r]$ the partial transcript is m_1, \dots, m_j , which determines the values $\{l_{j'}, \widehat{\mathcal{D}}_{I_{j'}}^{(j')}, \mathcal{D}_{I_{j'}}^{(j')}, I_{j'}\}_{j' \leq j}$. The index set $I_j \subseteq [k]$ satisfies all of the following:

- (1) $|I_j| = k^{1-j/r}$.
- (2) Each value $x_{I_j} \in [t]^{|I_j|}$ satisfies $\widehat{\mathcal{D}}_{I_j}^{(j)}(x_{I_j}) \leq 4 \cdot \mathcal{D}_{I_j}^{(j)}(x_{I_j})$.
- (3) Each nonempty subset $I' \subseteq I_j$ satisfies

$$\mathbb{H}(\widehat{\mathcal{D}}_{I'}^{(j)}) \geq \left(cE - \sum_{u=1}^j \frac{16^{j-u+1} l_u}{k^{1-(u-1)/r}} - 22^j \right) |I'|.$$

In accordance with our informal discussion in section 3.1, I_j is a subset of indices on which both parties have learned little information about each other from the partial transcript m_1, \dots, m_j . Invariant 3.3(2) ensures that the two parties draw their inputs after the j th round from similar distributions. Invariant 3.3(3) is the most important property. It says that the information revealed by $\widehat{\mathcal{D}}^{(j)}$ about I' is roughly what one would expect, given the message lengths l_1, \dots, l_j . Note that the u th message conveys information about $|I_{u-1}| = k^{1-(u-1)/r}$ indices so the average information-per-index should be $l_u/k^{1-(u-1)/r}$. The factor 16^{j-u+1} and the extra term 22^j come from Lemma 3.1, which throws away part of the input distribution in each round, progressively distorting the distributions in minor ways.

To begin our induction, at round j we find a large fraction \mathcal{M}'_j of possible messages m_j that reveal little information about the sender's input, projected onto I_{j-1} . This is possible because the length of the message $l_j = |m_j|$ reflects an upper bound on the expected information gain. This idea is formalized in Lemma 3.4.

LEMMA 3.4. Fix $j \in [1, r]$ and suppose Invariant 3.3 holds for m_1, \dots, m_{j-1} . Then there exists a subset of messages \mathcal{M}'_j with $\mu_j(\mathcal{M}'_j) \geq 1/2$ such that each message $m_j \in \mathcal{M}'_j$ satisfies

$$\mathbb{H}(\mathcal{D}_{I_{j-1}}^{(j)}[m_j]) \geq \left(cE - 2 \sum_{u=1}^j \frac{16^{j-u} l_u}{k^{1-(u-1)/r}} - 2 \cdot 22^{j-1} \right) |I_{j-1}|.$$

Proof. Let \mathcal{M}'_j contain all messages m_j satisfying the above inequality and $\overline{\mathcal{M}'_j}$ be its complement. Suppose, for the purpose of obtaining a contradiction, that the conclusion of the lemma is not true, i.e., $\mu_j(\overline{\mathcal{M}'_j}) = \alpha > 1/2$. Then the entropy of $\widehat{\mathcal{D}}_{I_{j-1}}^{(j-1)}$ can be upper bounded as follows:

$$\begin{aligned} & \mathbb{H}(\widehat{\mathcal{D}}_{I_{j-1}}^{(j-1)}) \\ &= \mathbb{I}(\widehat{\mathcal{D}}_{I_{j-1}}^{(j-1)}; M_j) + \sum_{m_j \in (\mathcal{M}'_j \cup \overline{\mathcal{M}'_j})} \mu_j(m_j) \mathbb{H}(\mathcal{D}_{I_{j-1}}^{(j)}[m_j]) \\ &\leq \mathbb{H}(M_j) + \sum_{m_j \in (\mathcal{M}'_j \cup \overline{\mathcal{M}'_j})} \mu_j(m_j) \mathbb{H}(\mathcal{D}_{I_{j-1}}^{(j)}[m_j]). \end{aligned}$$

The steps above follow from the definition of $\mathbb{I}(\cdot; \cdot)$ and the fact that $\mathbb{I}(X; \cdot) \leq \mathbb{H}(X)$. Since $\mathbb{H}(M_j) \leq |M_j| = l_j$, we have

$$\leq l_j + \sum_{m_j \in \mathcal{M}'_j} \mu_j(m_j) \mathbb{H}(\mathcal{D}_{I_{j-1}}^{(j)}[m_j]) + \sum_{m_j \in \overline{\mathcal{M}'_j}} \mu_j(m_j) \mathbb{H}(\mathcal{D}_{I_{j-1}}^{(j)}[m_j])$$

and from the definition of $\overline{\mathcal{M}}'_j$,

$$\begin{aligned} &< l_j + (1 - \alpha)cE|I_{j-1}| + \alpha \left(cE - 2 \sum_{u=1}^j \frac{16^{j-u}l_u}{k^{1-(u-1)/r}} - 2 \cdot 22^{j-1} \right) |I_{j-1}| \\ &= l_j + \left(cE - 2\alpha \sum_{u=1}^j \frac{16^{j-u}l_u}{k^{1-(u-1)/r}} - 2\alpha \cdot 22^{j-1} \right) |I_{j-1}| \\ &< \left(cE - \sum_{u=1}^{j-1} \frac{16^{j-u}l_u}{k^{1-(u-1)/r}} - 22^{j-1} \right) |I_{j-1}|. \end{aligned}$$

The last line follows because $\alpha > 1/2$, which contradicts Invariant 3.3(3) at index $j - 1$. \square

After the j th message m_j is sent, the next step is to identify a set of coordinates I_j such that $\mathcal{D}^{(j)}$ still reveals little information about I_j and every subset of I_j , since we need this property to hold for I_{j+1}, \dots, I_r in the future, all of which are subsets of I_j . We also want I_j not to contain many low probability points w.r.t. $\mathcal{D}^{(j-1)}$, since this may stop us from applying Lemma 3.2 later on. These two constraints are captured by parts (2) and (1), respectively, of Lemma 3.5.

LEMMA 3.5. Fix $j \in [1, r]$ and suppose Invariant 3.3 holds for m_1, \dots, m_{j-1} . Then there exists a subset of messages $\mathcal{M}_j \subseteq \mathcal{M}'_j$ (from Lemma 3.4) with $\mu_j(\mathcal{M}_j) \geq 1/4$ such that for each message $m_j \in \mathcal{M}_j$, there exists a subset $I_j \subseteq I_{j-1}$ of size $|I_j| = k^{1-j/r}$ satisfying both of the following properties:

1. $\Pr_{x_{I_j} \sim \mathcal{D}^{(j)}_{I_j}}[\mathcal{D}^{(j-1)}_{I_j}(x_{I_j}) < (4t)^{-|I_j|}/32] \leq 1/2$.
2. Each nonempty subset $I' \subseteq I_j$ satisfies

$$H(\mathcal{D}^{(j)}_{I'}) \geq \left(cE - 4 \sum_{u=1}^j \frac{16^{j-u}l_u}{k^{1-(u-1)/r}} - 4 \cdot 22^{j-1} \right) |I'|.$$

Proof. We first prove that for each message $m_j \in \mathcal{M}'_j$ (from Lemma 3.4), there exists a subset $J_0 \subseteq I_{j-1}$ of size $|J_0| \geq |I_{j-1}|/2$ such that each nonempty subset $I' \subseteq J_0$ satisfies part (2) of the lemma. Suppose J_1, J_2, \dots, J_w are disjoint subsets of I_{j-1} , each of which violates the inequality of part (2), whereas none of the subsets of $J_0 = I_{j-1} \setminus (\bigcup_{v=1}^w J_v)$ do. Using Fact 2.1 and the definition of J_v , we upper bound the entropy of $\mathcal{D}^{(j)}_{I_{j-1}}$ as follows:

$$\begin{aligned} H(\mathcal{D}^{(j)}_{I_{j-1}}) &\leq \sum_{v=0}^w H(\mathcal{D}^{(j)}_{J_v}) \\ &< cE|J_0| + \sum_{v=1}^w \left(cE - 4 \sum_{u=1}^j \frac{16^{j-u}l_u}{k^{1-(u-1)/r}} - 4 \cdot 22^{j-1} \right) |J_v| \\ &= cE|I_{j-1}| - 4|I_{j-1} \setminus J_0| \left(\sum_{u=1}^j \frac{16^{j-u}l_u}{k^{1-(u-1)/r}} + 22^{j-1} \right). \end{aligned}$$

On the other hand, from Lemma 3.4, having $m_j \in \mathcal{M}'_j$ guarantees that

$$H(\mathcal{D}^{(j)}_{I_{j-1}}) \geq \left(cE - 2 \sum_{u=1}^j \frac{16^{j-u}l_u}{k^{1-(u-1)/r}} - 2 \cdot 22^{j-1} \right) |I_{j-1}|.$$

The two inequalities above are consistent only if $|I_{j-1} \setminus J_0| \leq |I_{j-1}|/2$, or equivalently $|J_0| \geq |I_{j-1}|/2$. Thus, J_0 exists with the right cardinality, as claimed.

Now suppose, for the purpose of obtaining a contradiction, that the lemma is false. For every $m_j \in \mathcal{M}'_j$ there is a corresponding index set J_0 whose subsets satisfy part (2) of the lemma. If the lemma is false, that means there is a subset $\mathcal{M}''_j \subseteq \mathcal{M}'_j$ of “bad” messages with $\mu_j(\mathcal{M}''_j) > 1/4$ such that, for each $m_j \in \mathcal{M}''_j$, none of the $\binom{|J_0|}{|I_j|}$ choices for $I_j \subseteq J_0$ satisfy part (1) of the lemma. (Remember that J_0 depends on m_j but the lower bound on $|J_0| \geq |I_{j-1}|/2$ is independent of m_j .) Consider the following summation:

$$Z = \sum_{\substack{I_j \subseteq I_{j-1} : \\ |I_j| = k^{1-j/r}}} \sum_{\substack{x_{I_j} \in [t]^{|I_j|} : \\ \mathcal{D}_{I_j}^{(j-1)}(x_{I_j}) < (4t)^{-|I_j|}/32}} \mathcal{D}_{I_j}^{(j-1)}(x_{I_j}).$$

We can easily upper bound Z as follows:

$$Z < \binom{|I_{j-1}|}{|I_j|} \cdot t^{|I_j|} \cdot \frac{(4t)^{-|I_j|}}{32} = \binom{|I_{j-1}|}{|I_j|} 2^{-2|I_j|-5}.$$

Invariant 3.3(2) relates $\mathcal{D}^{(j-1)}$ and $\widehat{\mathcal{D}}^{(j-1)}$, which lets us lower bound Z :

$$Z \geq \frac{1}{4} \sum_{\substack{I_j \subseteq I_{j-1} : \\ |I_j| = k^{1-j/r}}} \sum_{\substack{x_{I_j} \in [t]^{|I_j|} : \\ \mathcal{D}_{I_j}^{(j-1)}(x_{I_j}) < (4t)^{-|I_j|}/32}} \widehat{\mathcal{D}}_{I_j}^{(j-1)}(x_{I_j}).$$

By definition, $\widehat{\mathcal{D}}^{(j-1)}$ is a convex combination of the $\mathcal{D}^{(j)}[m_j]$ distributions, weighted according to $\mu_j(\cdot)$. Hence, the expression above is lower bounded by

$$\geq \frac{1}{4} \sum_{\substack{I_j \subseteq I_{j-1} : \\ |I_j| = k^{1-j/r}}} \sum_{\substack{x_{I_j} \in [t]^{|I_j|} : \\ \mathcal{D}_{I_j}^{(j-1)}(x_{I_j}) < (4t)^{-|I_j|}/32}} \sum_{m_j \in \mathcal{M}''_j} \mu_j(m_j) \cdot \mathcal{D}_{I_j}^{(j)}[m_j](x_{I_j})$$

and rearranging sums,

$$\geq \frac{1}{4} \sum_{m_j \in \mathcal{M}''_j} \mu_j(m_j) \sum_{\substack{I_j \subseteq J_0 : \\ |I_j| = k^{1-j/r}}} \sum_{\substack{x_{I_j} \in [t]^{|I_j|} : \\ \mathcal{D}_{I_j}^{(j-1)}(x_{I_j}) < (4t)^{-|I_j|}/32}} \mathcal{D}_{I_j}^{(j)}[m_j](x_{I_j}).$$

By definition, for every $m_j \in \mathcal{M}''_j$ and every choice of $I_j \subseteq J_0$, part (1) of the lemma is violated. Continuing with the inequalities,

$$\begin{aligned} &> \frac{1}{4} \sum_{m_j \in \mathcal{M}''_j} \mu_j(m_j) \cdot \binom{|J_0|}{|I_j|} \cdot \frac{1}{2} \\ &> \frac{1}{32} \binom{|I_{j-1}|/2}{|I_j|}. \end{aligned}$$

The last line follows since $\mu_j(\mathcal{M}''_j) > 1/4$. This contradicts the upper bound on Z whenever $k^{1/r}$ is at least some sufficiently large constant. \square

The receiver of m_j constructs a new distribution $\widehat{\mathcal{D}}^{(j)}$ in two steps. After fixing I_j , we construct $\widetilde{\mathcal{D}}^{(j)}$ by combining $\mathcal{D}^{(j-1)}$ and $\mathcal{D}^{(j)}$, filtering out some points in the space whose probability mass is too low. We then construct $\widehat{\mathcal{D}}^{(j)}$ from $\widetilde{\mathcal{D}}^{(j)}$ and $\mathcal{D}^{(j-1)}$ by filtering out points that occur under $\widetilde{\mathcal{D}}^{(j)}$ with substantially larger probability than they do under $\mathcal{D}^{(j-1)}$.

Formally, suppose Invariant 3.3 holds for m_1, \dots, m_{j-1} . For each message $m_j \in \mathcal{M}_j$ (from Lemma 3.5), let I_j be selected to satisfy both properties of Lemma 3.5. Define the probability mass of a vector $x \in [t]^k$ under $\widetilde{\mathcal{D}}^{(j)}$ as follows:

$$\widetilde{\mathcal{D}}^{(j)}(x) = \begin{cases} 0 & \text{if } \mathcal{D}_{I_j}^{(j-1)}(x_{I_j}) < \frac{(4t)^{-|I_j|}}{32}; \\ \frac{\mathcal{D}_{I_j}^{(j)}(x_{I_j})}{\beta_1} \cdot \frac{\mathcal{D}^{(j-1)}(x)}{\mathcal{D}_{I_j}^{(j-1)}(x_{I_j})} & \text{otherwise,} \end{cases}$$

where β_1 is

$$\beta_1 = \Pr_{x_{I_j} \sim \mathcal{D}_{I_j}^{(j)}} \left[\mathcal{D}_{I_j}^{(j-1)}(x_{I_j}) \geq \frac{(4t)^{-|I_j|}}{32} \right].$$

In other words, we discard a $1 - \beta_1$ fraction of the distribution $\mathcal{D}^{(j)}$, but ignoring this effect, the projection of $\widetilde{\mathcal{D}}^{(j)}$ onto I_j has the same distribution as $\mathcal{D}^{(j)}$ onto I_j , and conditioned on the value of x_{I_j} , the distribution $\widetilde{\mathcal{D}}^{(j)}$ (projected onto $[k] \setminus I_j$) is identical to $\mathcal{D}^{(j-1)}$. We derive $\widehat{\mathcal{D}}^{(j)}$ from $\widetilde{\mathcal{D}}^{(j)}$ with a similar transformation.

$$\widehat{\mathcal{D}}^{(j)}(x) = \begin{cases} 0 & \text{if } \frac{\widetilde{\mathcal{D}}_{I_j}^{(j)}(x_{I_j})}{\mathcal{D}_{I_j}^{(j-1)}(x_{I_j})} > 2^{\gamma_j}; \\ \frac{\widetilde{\mathcal{D}}_{I_j}^{(j)}(x_{I_j})}{\beta_2} \cdot \frac{\mathcal{D}^{(j-1)}(x)}{\mathcal{D}_{I_j}^{(j-1)}(x_{I_j})} & \text{otherwise,} \end{cases}$$

where β_2 and γ_j are defined to be

$$\begin{aligned} \beta_2 &= \Pr_{x_{I_j} \sim \widetilde{\mathcal{D}}_{I_j}^{(j)}} \left[\frac{\widetilde{\mathcal{D}}_{I_j}^{(j)}(x_{I_j})}{\mathcal{D}_{I_j}^{(j-1)}(x_{I_j})} \leq 2^{\gamma_j} \right], \\ \gamma_j &= \left(\sum_{u=1}^j l_u \left(\frac{16^{j-u+1}}{k^{1-(u-1)/r}} \right) + (16 \cdot 22^{j-1} + 6) \right) |I_j| + 6 \\ &\leq \sum_{u=1}^j l_u \left(\frac{16}{k^{1/r}} \right)^{j-u+1} + 22^j \cdot |I_j| + 6. \end{aligned}$$

The proofs of Lemmas 3.6 and 3.7 use several simple observations about $\widetilde{\mathcal{D}}^{(j)}$ and $\widehat{\mathcal{D}}^{(j)}$:

1. Lemma 3.5(1) states that $\beta_1 \geq 1/2$. Lemma 3.5(2) lower bounds the entropy of $\mathcal{D}_{I_j}^{(j)}$. We apply Lemma 3.1 to $\mathcal{D}_{I_j}^{(j)}$ and $\widetilde{\mathcal{D}}_{I_j}^{(j)}$ (taking the roles of p and q , respectively) with parameter $\alpha = 1/2 \leq \beta_1$ and obtain the following lower bound on the entropy of $\widetilde{\mathcal{D}}_{I_j}^{(j)}$:

$$H(\widetilde{\mathcal{D}}_{I_j}^{(j)}) \geq \left(cE - 8 \sum_{u=1}^j \frac{16^{j-u} l_u}{k^{1-(u-1)/r}} - 8 \cdot 22^{j-1} - 2 \right) |I_j|.$$

2. We can then apply Lemma 3.2 to $\mathcal{D}_{I_j}^{(j-1)}$ and $\tilde{\mathcal{D}}_{I_j}^{(j)}$ (taking the roles of p and q , respectively) with parameters

$$\begin{aligned} g_1 &= \left(8 \sum_{u=1}^j \frac{16^{j-u} l_u}{k^{1-(u-1)/r}} + (8 \cdot 22^{j-1} + 2) \right) |I_j|, \\ g_2 &= 2|I_j| + 5, \\ \text{and } \alpha &= 1/2. \end{aligned}$$

Since $g_1/\alpha + g_2 - (1-\alpha) \log(1-\alpha)/\alpha = \gamma_j$, we conclude that $\beta_2 \geq 1 - \alpha = 1/2$. Thus, for each value $x_{I_j} \in \text{supp}(\hat{\mathcal{D}}_{I_j}^{(j)})$,

$$(3.1) \quad \hat{\mathcal{D}}_{I_j}^{(j)}(x_{I_j}) = \frac{\tilde{\mathcal{D}}_{I_j}^{(j)}(x_{I_j})}{\beta_2} = \frac{\mathcal{D}_{I_j}^{(j)}(x_{I_j})}{\beta_1 \beta_2} \leq 4 \cdot \mathcal{D}_{I_j}^{(j)}(x_{I_j}).$$

Lemma 3.6 completes the inductive step by lower bounding the entropy of $\hat{\mathcal{D}}_{I'}^{(j)}$ for every nonempty subset $I' \subseteq I_j$. To put it another way, it ensures that the coordinates in I_j remain almost completely unknown to both parties.

LEMMA 3.6. *Fix $j \in [1, r]$ and suppose Invariant 3.3 holds for m_1, \dots, m_{j-1} . Then, for each message $m_j \in \mathcal{M}_j$ (from Lemma 3.5), Invariant 3.3 also holds for m_1, \dots, m_j .*

Proof. Due to Lemma 3.5 and (3.1), the first two properties of Invariant 3.3 are satisfied. For each nonempty subset $I' \subseteq I_j$, the third property of Invariant 3.3 can be derived from the second property of Lemma 3.5 and an application of Lemma 3.1 to $\mathcal{D}_{I'}^{(j)}$ and $\hat{\mathcal{D}}_{I'}^{(j)}$ (taking the roles of p and q , respectively) with parameter $\alpha = 1/4$ as follows:

$$\begin{aligned} \mathbb{H}(\hat{\mathcal{D}}_{I'}^{(j)}) &\geq \left(cE - 16 \sum_{u=1}^j \frac{16^{j-u} l_u}{k^{1-(u-1)/r}} - 16 \cdot 22^{j-1} - 4 \right) |I'| \\ &\geq \left(cE - \sum_{u=1}^j \frac{16^{j-u+1} l_u}{k^{1-(u-1)/r}} - 22^j \right) |I'|. \quad \square \end{aligned}$$

Aside from maintaining Invariant 3.3 round by round, another important part of our proof is to compute the error probability. Lemma 3.7 shows how the error probabilities of two consecutive rounds are related after our modification to the protocol. More importantly, it also illustrates the reason to bound the *pointwise* ratio between $\tilde{\mathcal{D}}_{I_j}^{(j)}$ and $\mathcal{D}_{I_j}^{(j-1)}$.

LEMMA 3.7. *Fix a round $j \in [1, r]$ and suppose Invariant 3.3 holds for m_1, \dots, m_{j-1} . Fix any specific message $m_j \in \mathcal{M}_j$ (from Lemma 3.5). Define p to be the probability of error, when the protocol begins after round j with the inputs drawn from $\mathcal{D}^{(j)}$ and $\hat{\mathcal{D}}^{(j)}$, respectively. Then the probability of error is at least $2^{-\gamma_j - 1} p$ when the inputs are instead drawn from $\mathcal{D}^{(j)}$ and $\mathcal{D}^{(j-1)}$, respectively.*

Proof. From the definition of $\hat{\mathcal{D}}^{(j)}$, for each value $x \in \text{supp}(\hat{\mathcal{D}}^{(j)})$, we have

$$(3.2) \quad \frac{\hat{\mathcal{D}}^{(j)}(x)}{\mathcal{D}^{(j-1)}(x)} = \frac{\tilde{\mathcal{D}}_{I_j}^{(j)}(x_{I_j})}{\beta_2 \mathcal{D}_{I_j}^{(j-1)}(x_{I_j})} \leq \frac{2^{\gamma_j}}{\beta_2} \leq 2^{\gamma_j + 1}.$$

This concludes the proof. □

Finally, with all lemmas proved above, we have reached the point to calculate the initial error probability.

LEMMA 3.8. *Recall that $c = 1/2, c' = c/100$. Fix any $r \in [1, (\log k)/6]$ and $E \geq 100k^{1-1/r}/c$. Suppose the initial input vectors are drawn independently and uniformly from $[t]^k$, where $t = 2^{cE}$. Then the error probability of the EqualityTesting protocol, p_{err} , is greater than 2^{-E} .*

Proof. First suppose Invariant 3.3 holds for m_1, \dots, m_r and consider the situation after the final round, where the inputs are drawn from $\mathcal{D}^{(r)}$ and $\widehat{\mathcal{D}}^{(r)}$, respectively. Notice that I_r is a singleton set, so the entropy of $\widehat{\mathcal{D}}_{I_r}^{(r)}$ can be lower bounded as follows.

$$\begin{aligned} \text{H}(\widehat{\mathcal{D}}_{I_r}^{(r)}) &\geq cE - \sum_{u=1}^r \frac{16^{r-u+1}l_u}{k^{1-(u-1)/r}} - 22^r && \text{Invariant 3.3(3)} \\ &= cE - \frac{16}{k^{1/r}} \sum_{u=1}^r l_u \left(\frac{16}{k^{1/r}}\right)^{r-u} - 22^r \\ &\geq cE - \frac{16}{k^{1/r}} \sum_{u=1}^r l_u - 22k^{1-1/r} && k^{1/r} \geq 2^6 \text{ due to } r \leq (\log k)/6 \\ &\geq cE - 16c'E - 22k^{1-1/r} > \frac{cE}{2} && \text{because } \sum_{u=1}^r l_u \leq c'E k^{1/r}. \end{aligned}$$

From the lower bound on the entropy of $\widehat{\mathcal{D}}_{I_r}^{(r)}$, we can easily show that there exists no value x_{I_r} such that $\widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r}) = \alpha > 3/4$. If there were such a value, then the entropy of $\widehat{\mathcal{D}}_{I_r}^{(r)}$ can also be upper bounded as

$$\begin{aligned} \text{H}(\widehat{\mathcal{D}}_{I_r}^{(r)}) &\leq \alpha \log \frac{1}{\alpha} + (1 - \alpha) \log \frac{t}{1 - \alpha} \\ &< \frac{cE}{4} + \alpha \log \frac{1}{\alpha} + (1 - \alpha) \log \frac{1}{1 - \alpha} < \frac{cE}{2}, \end{aligned}$$

contradicting the lower bound on $\text{H}(\widehat{\mathcal{D}}_{I_r}^{(r)})$.

After all r rounds of communication, the receiver of the last message has to make the decision on I_r , either “=” or “≠,” depending only on his own input on x_{I_r} . Define $\mathcal{X}_0 \subseteq [t]$ to be the subset of values x_{I_r} such that the protocol outputs “≠” on I_r after r rounds of communication, and let $\mathcal{X}_1 = [t] \setminus \mathcal{X}_0$ and $\beta = \widehat{\mathcal{D}}_{I_r}^{(r)}(\mathcal{X}_0)$. Let x'_{I_r} be the other party’s value at index I_r . The protocol errs if $x_{I_r} = x'_{I_r} \in \mathcal{X}_0$ (a false “≠” on I_r) or if $x_{I_r} \neq x'_{I_r}$ and $x_{I_r} \in \mathcal{X}_1$ (a false “=” on I_r). Thus, the final error probability is at least

$$\begin{aligned} &\sum_{x_{I_r} \in \mathcal{X}_0} \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r}) \mathcal{D}_{I_r}^{(r)}(x_{I_r}) + \sum_{x_{I_r} \in \mathcal{X}_1} \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r}) (1 - \mathcal{D}_{I_r}^{(r)}(x_{I_r})) \\ &= \sum_{x_{I_r} \in \mathcal{X}_0} \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r}) \mathcal{D}_{I_r}^{(r)}(x_{I_r}) + \sum_{x_{I_r} \in \mathcal{X}_1} \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r}) \sum_{x'_{I_r} \neq x_{I_r}} \mathcal{D}_{I_r}^{(r)}(x'_{I_r}). \end{aligned}$$

By Invariant 3.3(2), this is

$$\begin{aligned} &\geq \frac{1}{4} \sum_{x_{I_r} \in \mathcal{X}_0} \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r})^2 + \frac{1}{4} \sum_{x_{I_r} \in \mathcal{X}_1} \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r}) \sum_{x'_{I_r} \neq x_{I_r}} \widehat{\mathcal{D}}_{I_r}^{(r)}(x'_{I_r}) \\ &= \frac{1}{4} \sum_{x_{I_r} \in \mathcal{X}_0} \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r})^2 + \frac{1}{4} \sum_{x_{I_r} \in \mathcal{X}_1} \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r}) \left(1 - \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r})\right) \\ &\geq \frac{1}{4} \sum_{x_{I_r} \in \mathcal{X}_0} \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r})^2 + \frac{1}{16} \sum_{x_{I_r} \in \mathcal{X}_1} \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r}). \end{aligned}$$

The previous line follows from $\widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r}) \leq 3/4$. By the convexity of x^2 we have

$$\geq \frac{\beta^2}{4t} + \frac{1-\beta}{16} \geq \frac{1}{4t}.$$

This result also meets the simple intuition that when the inputs to the two parties are almost uniformly random and no communication is allowed, the best strategy would be guessing “not equal” regardless of the actual input.

Finally, we are ready to transfer the error probability back round by round. From Lemmas 3.5 through 3.7, the error probability w.r.t. $\mathcal{D}^{(j)}$ and $\widehat{\mathcal{D}}^{(j)}$ differs from the error probability w.r.t. $\mathcal{D}^{(j-1)}$ and $\widehat{\mathcal{D}}^{(j-1)}$ by at most a $4 \cdot 2^{\gamma_j+1} = 2^{\gamma_j+3}$ factor. In particular, Lemmas 3.5 and 3.6 say that the j th message m_j satisfies Invariant 3.3 with respect to transcript m_1, \dots, m_j with probability at least $1/4$, provided Invariant 3.3 holds for m_1, \dots, m_{j-1} , and Lemma 3.7 says the error probabilities under the two measures differ by a 2^{γ_j+1} factor for any such m_j . Repeating this for each $j \in [1, r]$, we conclude that the initial error probability p_{err} is lower bounded by

$$p_{\text{err}} \geq \frac{1}{4t} \cdot \exp\left(-3r - \sum_{j=1}^r \gamma_j\right) = \exp\left(-cE - 2 - 3r - \sum_{j=1}^r \gamma_j\right) > 2^{-E}$$

since

$$\begin{aligned} &cE + 2 + 3r + \sum_{j=1}^r \gamma_j \\ &\leq cE + 2 + 3r + 6r + \sum_{j=1}^r \sum_{u=1}^j l_u \left(\frac{16}{k^{1/r}}\right)^{j-u+1} + \sum_{j=1}^r 22^j |I_j|. \end{aligned}$$

Rearranging sums, we have

$$\begin{aligned} &\leq cE + 11r + \sum_{u=1}^r \frac{16l_u}{k^{1/r}} \sum_{j=u}^r \left(\frac{16}{k^{1/r}}\right)^{j-u} + 22k^{1-1/r} \sum_{j=1}^r \left(\frac{22}{k^{1/r}}\right)^{j-1} \\ &\leq cE + 11r + \frac{32}{k^{1/r}} \sum_{u=1}^r l_u + 44k^{1-1/r}. \end{aligned}$$

The last step follows from $k^{1/r} \geq 2^6$ since $r \leq (\log k)/6$. Because $\sum_{u=1}^r l_u \leq c'Ek^{1/r}$, we have

$$\leq cE + \frac{11cE}{100} + \frac{32cE}{100} + \frac{44cE}{100} < E. \quad \square$$

Proof of Theorem 2. Lemma 3.8 actually shows that given integers $k \geq 1$ and $r \leq (\log k)/6$, any r -round deterministic protocol for EqualityTesting on vectors of length k that has distributional error probability $p_{\text{err}} = 2^{-E}$ with respect to the uniform input distribution on $[t]^k$, where $t = 2^{cE}$, requires at least $\Omega(Ek^{1/r})$ bits of communication. Notice that the additional assumption $E \geq 100k^{1-1/r}/c$ always makes sense since there is a trivial $\Omega(k)$ lower bound on the communication complexity of EqualityTesting, regardless of r . Thus, Theorem 2 follows directly from Yao’s minimax principle. \square

3.3. A lower bound on ExistsEqual. The proof of Theorem 3 is almost the same as that of Theorem 2, except for the final step, namely Lemma 3.8, in which we first compute the final error probability after all r rounds of communication and then transfer it backward round by round using Lemma 3.7. The problem with applying the same argument to ExistsEqual protocols is that the receiver of the last message may be able to announce the correct answer, even though it knows little information about the inputs on the single coordinate I_r .

In order to prove Theorem 3, first notice that Lemma 3.4 through Lemma 3.7 also hold perfectly well for ExistsEqual protocols as no modification is required in their proofs. Therefore, it is sufficient to prove the following Lemma 3.9, which is an analogue of Lemma 3.8 for ExistsEqual. It is based mainly on Markov’s inequality.

LEMMA 3.9. *Recall that $c = 1/2, c' = c/100$. Consider an execution of a deterministic r -round ExistsEqual protocol, $r \in [1, (\log k)/6]$, on input vectors drawn independently and uniformly from $[t]^k$, where $t = 2^{cE}$. Here $E \geq 100k^{1-1/r}/c$ if $r > 1$ and $E \geq (100 \log k)/c$ otherwise. Then the protocol errs with probability $p_{\text{err}} > 2^{-E}$.*

Proof. Similarly to the proof of Lemma 3.8, we first consider the situation after the final round; fix the transcript m_1, \dots, m_r . The receiver of m_r now makes a Boolean decision for ExistsEqual based on every coordinate of his input x . Let $\mathcal{X}_0 \subseteq [t]^k$ be the subset of values x such that the protocol outputs “no” (there exists no coordinate where the two vectors are equal) after seeing transcript m_1, \dots, m_r , and let $\mathcal{X}_1 = [t]^k \setminus \mathcal{X}_0$. The protocol may have false negatives on \mathcal{X}_0 and false positives on \mathcal{X}_1 . We only consider false negatives in which the vectors agree on coordinate I_r . Thus, the error probability at this moment is at least

$$(3.3) \quad \sum_{x \in \mathcal{X}_0} \widehat{\mathcal{D}}^{(r)}(x) \mathcal{D}_{I_r}^{(r)}(x_{I_r}) + \sum_{x \in \mathcal{X}_1} \widehat{\mathcal{D}}^{(r)}(x) \left(1 - \sum_{y \in \mathcal{N}(x)} \mathcal{D}^{(r)}(y) \right),$$

where $\mathcal{N}(x) = \{y \in [t]^k \mid \text{there exists some } i \in [k] \text{ such that } x_i = y_i\}$ is the subset of input vectors that agree with x on at least one coordinate. There are two difficulties with (3.3). The first is that we do not know anything about \mathcal{X}_0 and \mathcal{X}_1 . The second is that we are really interested in the initial probability of error, before m_1 is transmitted, not the error under $\widehat{\mathcal{D}}^{(r)}$ and $\mathcal{D}^{(r)}$.

Our first task is to lower bound the second term of (3.3). Consider the following summation Z_0 over all transcripts m_1, \dots, m_r in which $m_j \in \mathcal{M}_j$ (from Lemma 3.5) and all $x \in [t]^k$. Recall that \mathcal{M}_j and μ_j depend on m_1, \dots, m_{j-1} .

$$Z_0 = \sum_{m_1 \in \mathcal{M}_1} \mu_1(m_1) \sum_{m_2 \in \mathcal{M}_2} \mu_2(m_2) \cdots \sum_{m_r \in \mathcal{M}_r} \mu_r(m_r) \sum_{x \in [t]^k} \widehat{\mathcal{D}}^{(r)}(x) \sum_{y \in \mathcal{N}(x)} \mathcal{D}^{(r)}(y).$$

From the proof of Lemma 3.7 (equation (3.2)), we can upper bound Z_0 as follows:

$$Z_0 \leq \sum_{m_1 \in \mathcal{M}_1} \mu_1(m_1) \cdots \sum_{m_r \in \mathcal{M}_r} \mu_r(m_r) \sum_{\substack{x \in [t]^k, \\ y \in \mathcal{N}(x)}} 2^{\gamma_r+1} \cdot \mathcal{D}^{(r-1)}(x) \cdot \mathcal{D}^{(r)}(y).$$

Notice that γ_r and $\mathcal{D}^{(r-1)}$ are independent of the choice of m_r , hence by rearranging sums, this is equal to

$$= \sum_{m_1 \in \mathcal{M}_1} \mu_1(m_1) \cdots \sum_{m_{r-1} \in \mathcal{M}_{r-1}} \mu_{r-1}(m_{r-1}) \sum_{\substack{x \in [t]^k, \\ y \in \mathcal{N}(x)}} 2^{\gamma_r+1} \cdot \mathcal{D}^{(r-1)}(x) \sum_{m_r \in \mathcal{M}_r} \mu_r(m_r) \cdot \mathcal{D}^{(r)}(y).$$

By definition, $\widehat{\mathcal{D}}^{(r-1)}$ is a convex combination of the $\mathcal{D}^{(r)}[m_r]$ distributions, weighted according to $\mu_r(\cdot)$. Hence, the expression above is upper bounded by

$$\leq \sum_{m_1 \in \mathcal{M}_1} \mu_1(m_1) \cdots \sum_{m_{r-1} \in \mathcal{M}_{r-1}} \mu_{r-1}(m_{r-1}) \sum_{\substack{x \in [t]^k, \\ y \in \mathcal{N}(x)}} 2^{\gamma_r+1} \cdot \mathcal{D}^{(r-1)}(x) \cdot \widehat{\mathcal{D}}^{(r-1)}(y).$$

By the symmetry of x and y , this is equal to

$$= \sum_{m_1 \in \mathcal{M}_1} \mu_1(m_1) \cdots \sum_{m_{r-1} \in \mathcal{M}_{r-1}} \mu_{r-1}(m_{r-1}) \sum_{\substack{x \in [t]^k, \\ y \in \mathcal{N}(x)}} 2^{\gamma_r+1} \cdot \widehat{\mathcal{D}}^{(r-1)}(x) \cdot \mathcal{D}^{(r-1)}(y).$$

We repeat the same argument for rounds $r - 1$ down to 1, upper bounding Z_0 by

$$\leq \exp \left(r + \sum_{j=1}^r \gamma_j \right) \cdot \sum_{\substack{x \in [t]^k, \\ y \in \mathcal{N}(x)}} \widehat{\mathcal{D}}^{(0)}(x) \cdot \mathcal{D}^{(0)}(y) \leq \exp \left(r + \sum_{j=1}^r \gamma_j \right) \cdot \frac{k}{t}.$$

The last inequality above follows from a union bound since, under the initial distributions $\widehat{\mathcal{D}}^{(0)}, \mathcal{D}^{(0)}$, each of the k coordinates is equal with probability $1/t$. Recall that $E \geq 100k^{1-1/r}/c$ when $r > 1$ and $E \geq (100 \log k)/c$ otherwise. Hence, using the same argument as that in the proof of Lemma 3.8, we can further bound this as

$$\leq 2^{0.83cE} \cdot 2^{0.02cE} \cdot 2^{-cE} = 2^{-0.15cE}$$

since

$$r + \sum_{j=1}^r \gamma_j \leq 7r + \sum_{j=1}^r \sum_{u=1}^j l_u \left(\frac{16}{k^{1/r}} \right)^{j-u+1} + \sum_{j=1}^r 22^j |I_j| \leq \frac{7cE}{100} + \frac{32cE}{100} + \frac{44cE}{100} = \frac{83cE}{100},$$

and $k \leq (cE/100)^{r/(r-1)} \leq (cE/100)^2 \leq 2^{0.02cE}$ when $r > 1$ and $k \leq 2^{0.01cE}$ otherwise.

We now want to narrow our attention to transcripts that both satisfy Invariant 3.3 and terminate with $\sum_{x \in [t]^k} \widehat{\mathcal{D}}^{(r)}(x) \sum_{y \in \mathcal{N}(x)} \mathcal{D}^{(r)}(y)$ not much larger than its expectation. Fix a partial transcript (m_1, \dots, m_j) through round j such that $m_{j'} \in \mathcal{M}_{j'}$ holds for every $j' \leq j$. Define Z_j as follows:

$$Z_j = \sum_{m_{j+1} \in \mathcal{M}_{j+1}} \mu_{j+1}(m_{j+1}) \cdots \sum_{m_r \in \mathcal{M}_r} \mu_r(m_r) \sum_{x \in [t]^k} \widehat{\mathcal{D}}^{(r)}(x) \sum_{y \in \mathcal{N}(x)} \mathcal{D}^{(r)}(y).$$

By Markov's inequality, there exists a subset of messages $\widehat{\mathcal{M}}_1 \subseteq \mathcal{M}_1$ with $\mu_1(\widehat{\mathcal{M}}_1) \geq \mu_1(\mathcal{M}_1)/2 \geq 1/8$ such that each message $m_1 \in \widehat{\mathcal{M}}_1$ satisfies $Z_1 \leq 2Z_0/\mu_1(\mathcal{M}_1) \leq 8Z_0$ since $\mu_1(\mathcal{M}_1) \geq 1/4$ from Lemma 3.5. Similarly, conditioned on any specific $m_1 \in \widehat{\mathcal{M}}_1$, by Markov's inequality, there exists a subset of messages $\widehat{\mathcal{M}}_2 \subseteq \mathcal{M}_2$ with $\mu_2(\widehat{\mathcal{M}}_2) \geq \mu_2(\mathcal{M}_2)/2 \geq 1/8$ such that each message $m_2 \in \widehat{\mathcal{M}}_2$ satisfies $Z_2 \leq 2Z_1/\mu_2(\mathcal{M}_2) \leq 8^2Z_0$. In general, conditioned on any specific partial transcript m_1, \dots, m_{j-1} such that $m_{j'} \in \widehat{\mathcal{M}}_{j'}$ holds for every $j' < j$, there exists a subset of messages $\widehat{\mathcal{M}}_j \subseteq \mathcal{M}_j$ with $\mu_j(\widehat{\mathcal{M}}_j) \geq \mu_j(\mathcal{M}_j)/2 \geq 1/8$ such that each message $m_j \in \widehat{\mathcal{M}}_j$ satisfies $Z_j \leq 8^j Z_0$.

After repeating the same argument r times, we get $\widehat{\mathcal{M}}_1, \dots, \widehat{\mathcal{M}}_r$ in sequence. For any sampled transcript m_1, \dots, m_r such that $m_j \in \widehat{\mathcal{M}}_j$ for all $j \leq r$, we have

$$Z_r = \sum_{x \in [t]^k} \widehat{\mathcal{D}}^{(r)}(x) \sum_{y \in \mathcal{N}(x)} \mathcal{D}^{(r)}(y) \leq 8^r Z_0 \leq 2^{3r} \cdot 2^{-0.15cE} \leq 2^{-0.12cE} \leq \frac{1}{4}$$

as $r \leq cE/100$ and $cE \geq 100$. Further, one more application of Markov's inequality shows that there exists a subset of values $\mathcal{X}' \subseteq [t]^k$ with $\widehat{\mathcal{D}}^{(r)}(\mathcal{X}') = \alpha \geq 1/2$ such that $\sum_{y \in \mathcal{N}(x)} \mathcal{D}^{(r)}(y) \leq 1/2$ holds for every $x \in \mathcal{X}'$.

As a result, we can now lower bound the final error probability as follows. Define $\beta = \widehat{\mathcal{D}}^{(r)}(\mathcal{X}_0 \cap \mathcal{X}')$.

$$\begin{aligned} & \sum_{x \in \mathcal{X}_0} \widehat{\mathcal{D}}^{(r)}(x) \mathcal{D}_{I_r}^{(r)}(x_{I_r}) + \sum_{x \in \mathcal{X}_1} \widehat{\mathcal{D}}^{(r)}(x) \left(1 - \sum_{y \in \mathcal{N}(x)} \mathcal{D}^{(r)}(y) \right) \\ & \geq \sum_{x \in (\mathcal{X}_0 \cap \mathcal{X}')} \widehat{\mathcal{D}}^{(r)}(x) \mathcal{D}_{I_r}^{(r)}(x_{I_r}) + \sum_{x \in (\mathcal{X}_1 \cap \mathcal{X}')} \widehat{\mathcal{D}}^{(r)}(x) \left(1 - \sum_{y \in \mathcal{N}(x)} \mathcal{D}^{(r)}(y) \right) \\ & \geq \frac{1}{4} \sum_{x \in (\mathcal{X}_0 \cap \mathcal{X}')} \widehat{\mathcal{D}}^{(r)}(x) \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r}) + \sum_{x \in (\mathcal{X}_1 \cap \mathcal{X}')} \widehat{\mathcal{D}}^{(r)}(x) \left(1 - \sum_{y \in \mathcal{N}(x)} \mathcal{D}^{(r)}(y) \right). \end{aligned}$$

The previous step follows from Invariant 3.3(2). By the definition of \mathcal{X}' , we have

$$\geq \frac{1}{4} \sum_{x \in (\mathcal{X}_0 \cap \mathcal{X}')} \widehat{\mathcal{D}}^{(r)}(x) \widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r}) + \frac{1}{2} \sum_{x \in (\mathcal{X}_1 \cap \mathcal{X}')} \widehat{\mathcal{D}}^{(r)}(x).$$

In order to minimize the above expression, we can now assume without loss of generality that the partition between $\mathcal{X}_0 \cap \mathcal{X}'$ and $\mathcal{X}_1 \cap \mathcal{X}'$ depends solely on x_{I_r} , as only the relative magnitude of $\widehat{\mathcal{D}}_{I_r}^{(r)}(x_{I_r})/4$ and $1/2$ matters. Continuing, from the convexity of x^2 ,

$$\geq \frac{\beta^2}{4t} + \frac{\alpha - \beta}{2} \geq \frac{\alpha^2}{4t} \geq \frac{1}{16t}.$$

Finally, we are ready to transfer the error probability back in exactly the same manner as we did in the proof of Lemma 3.8. Using a similar argument, the existence of $\widehat{\mathcal{M}}_j$ guarantees that

$$p_{\text{err}} \geq \frac{1}{16t} \cdot \exp\left(-4r - \sum_{j=1}^r \gamma_j\right) = \exp\left(-cE - 4 - 4r - \sum_{j=1}^r \gamma_j\right) > 2^{-E}$$

since

$$cE + 4 + 4r + \sum_{j=1}^r \gamma_j \leq cE + \frac{14cE}{100} + \frac{32cE}{100} + \frac{44cE}{100} < E. \quad \square$$

Proof of Theorem 3. Similarly to the proof of Theorem 2, Theorem 3 follows from Lemma 3.9 and a direct application of Yao's minimax principle. \square

4. New protocols for EqualityTesting and ExistsEqual. In this section, we attempt to prove that our $\Omega(Ek^{1/r})$ lower bound is *tight* for EqualityTesting. We manage to attain this bound in several situations but fail to achieve it for every value of E, k, r .

First of all, the $\Omega(Ek^{1/r})$ bound is binding only when it is at least $\Omega(k)$, which is necessary even when E is constant [31, 39, 18]. In Theorem 8 we give a $\log^*(k/E)$ -round protocol that reduces the effective dimension of the problem from k to at most E with $O(k)$ communication and basically lets us proceed under the assumption that $E \geq k$. (Note that if $E \geq k$ initially, $\log^*(k/E) = 0$.)

In Theorem 9 we give a simple protocol for EqualityTesting with communication $O(rEk^{1/r})$ when $E \geq k$. According to Theorem 2 this is optimal when $r = O(1)$. All of our remaining protocols aim to eliminate or reduce this seemingly unnecessary factor of r . In Theorem 10 we prove that ExistsEqual can be solved with $O(Ek^{1/r})$ communication, for any r and $E \geq k$, and Theorem 11 shows the same communication can be attained for EqualityTesting, but with $O(r)$ rounds rather than r . In particular, Theorems 8, 10, and 11 imply that EqualityTesting/ExistsEqual can be solved with absolutely optimal communication $O(k + E)$ in $\log k$ rounds, which is also round-optimal according to Theorems 2 and 3. However, Theorems 2, 9, and 11 leave the precise complexity of EqualityTesting open when $E \geq k$ and r is between $\omega(1)$ and $o(\log k)$.

Theorem 12 is our most sophisticated upper bound, in many ways. It proves that EqualityTesting can be solved using $O(Ek^{1/r} \log r + Er \log r)$ communication when $E \geq k$. When $r \geq \log k / \log \log k$ the first term is dominant, and the protocol comes within a $\log r \leq \log \log k$ factor of Theorem 2's lower bound. Taken together, these theorems highlight a potential complexity separation between ExistsEqual and EqualityTesting and between SetDisjointness and SetIntersection in the low error probability regime.

Theorems 4–7 follow by combining the dimension reduction of Theorem 8 with Theorems 9–12.

THEOREM 4. *There exists a $(\log^*(k/E) + r)$ -round randomized protocol for EqualityTesting on vectors of length k that errs with probability $p_{\text{err}} = 2^{-E}$, using $O(k + rEk^{1/r})$ bits of communication.*

THEOREM 5. *There exists a $(\log^*(k/E) + r)$ -round randomized protocol for ExistsEqual on vectors of length k that errs with probability $p_{\text{err}} = 2^{-E}$, using $O(k + Ek^{1/r})$ bits of communication.*

THEOREM 6. *There exists a $(\log^*(k/E) + O(r))$ -round randomized protocol for `EqualityTesting` on vectors of length k that errs with probability $p_{\text{err}} = 2^{-E}$, using $O(k + Ek^{1/r})$ bits of communication.*

THEOREM 7. *There exists a $(\log^*(k/E) + r)$ -round randomized protocol for `EqualityTesting` on vectors of length k that errs with probability $p_{\text{err}} = 2^{-E}$, using $O(k + Ek^{1/r} \log r + Er \log r)$ bits of communication.*

Remark 1. The $\log^*(k/E)$ terms in the round complexity of Theorems 4–7 are not absolute. They can each be replaced with $\max\{0, \log^*(k/E) - \log^*(C)\}$, at the cost of increasing the communication by $O(Ck)$.

Remark 2. By applying Theorem 1 to Theorems 4–7 we obtain `SetDisjointness/ SetIntersection` protocols with the same communication complexity, but with one more round of communication. In the case of `SetDisjointness` (Theorem 1 + Theorem 5), it is straightforward to skip the reduction of Theorem 1 and solve the problem directly with $O(k + Ek^{1/r})$ communication in $(r + \log^*(k/E))$ rounds. However, we do not see how to avoid Theorem 1’s extra round of communication when solving `SetIntersection`. That is, the `SetIntersection` protocols implied by Theorems 1, 4, and 7 use $(\log^*(k/E) + r + 1)$ rounds.

4.1. Overview and preliminaries. We start by giving a generic protocol for `EqualityTesting`. The protocol uses a simple subroutine for `ExistsEqual/EqualityTesting` when $k = 1$. Suppose Alice and Bob hold $x, y \in U = \{0, 1\}^l$, respectively. Alice picks a random $w \in \{0, 1\}^l$ from the shared random source and sends Bob $\tilde{x} = \langle x, w \rangle \bmod 2$, where $\langle \cdot, \cdot \rangle$ is the inner product operator. Bob computes $\tilde{y} = \langle y, w \rangle \bmod 2$ and declares “ $x = y$ ” iff $\tilde{x} = \tilde{y}$. Clearly, Bob never errs if $x = y$; it is straightforward to show that the probability of error is exactly $1/2$ when $x \neq y$. We call this protocol an *inner product test* and \tilde{x}, \tilde{y} *test bits*. A b -bit inner product test on x and y refers to b independent inner product tests on x and y .

At the beginning of phase j , $j \geq 1$, Alice and Bob agree on a subset I_{j-1} of coordinates on which all previous inner product tests have passed. In other words, they have *refuted* the potential equality $x_i \stackrel{?}{=} y_i$ for all $i \in [k] \setminus I_{j-1}$. Each coordinate $i \in I_{j-1}$ represents either an actual equality ($x_i = y_i$) or a *false positive* ($x_i \neq y_i$). At the beginning of the protocol, $I_0 = [k]$. In phase j , we perform l_j independent inner product tests on each coordinate in I_{j-1} and let $I_j \subseteq I_{j-1}$ be the remaining coordinates that pass all their respective inner product tests. Notice that each coordinate in I_{j-1} corresponding to equality will always pass all the tests and enter I_j , while those corresponding to inequalities will only enter I_j with probability 2^{-l_j} . At the end of the protocol, we declare all coordinates in I_r *equal* and all other coordinates *not equal*.

This finishes the description of the generic `EqualityTesting` protocol. Theorems 4–12 all build on the framework of the generic protocol, instantiating its steps in different ways.

4.1.1. A protocol for exchanging test bits. For `EqualityTesting`, it is possible that a constant fraction of the coordinates are actually equalities, which makes $|I_j| = \Theta(k)$ for every j . The naive implementation would explicitly exchange all $l_j |I_{j-1}|$ test bits and use $\Omega(kE)$ bits of communication in total. All the test bits corresponding to equalities are “wasted” in a sense.

For our application, it is important that the communication volume that Alice and Bob use to exchange their test bits in phase j be proportional to the number of

false positives in I_{j-1} , instead of the size of I_{j-1} . We will use a slightly improved version of a protocol of Feder et al. [21] for exchanging the test bits.

Imagine packing the test bits into vectors $\hat{x}, \hat{y} \in B^{|I_{j-1}|}$ where $B = \{0, 1\}^L$. Lemma 4.1 shows that Alice can transmit \hat{x} to Bob, at a cost that depends on an a priori upper bound on the Hamming distance $\text{dist}(\hat{x}, \hat{y})$, i.e., the number of the coordinates in I_{j-1} where they differ.

LEMMA 4.1 (cf. Feder et al. [21]). *Suppose Alice and Bob hold length- K vectors $x, y \in B^K$, where $B = \{0, 1\}^L$. Alice can send one $O(dL + d \log(K/d))$ -bit message to Bob, who generates a string $x' \in B^K$ such that the following holds. If the Hamming distance $\text{dist}(x, y) \leq d$, then $x = x'$; if $\text{dist}(x, y) > d$, then there is no guarantee.*

Proof. Define $G = (V, E)$ to be the graph on $V = B^K$ such that $\{u, v\} \in E$ iff $\text{dist}(u, v) \leq 2d$. The maximum degree in G is clearly at most $\Delta = \binom{K}{2d} \cdot 2^{2Ld}$ since there are $\binom{K}{2d}$ ways to select the $2d$ indices and 2^{2Ld} ways to change the coordinates at those indices so that there are at most $2d$ different coordinates. Let $\phi : V \mapsto [\Delta]$ be a proper Δ -coloring of G . Alice sends $\phi(x)$ to Bob, which requires $\log \Delta = O(dL + d \log(K/d))$ bits. Every string in the radius- d ball around y (w.r.t. dist) is colored differently since they are all at distance at most $2d$; hence if $\text{dist}(x, y) \leq d$, Bob can reconstruct x without error. \square

COROLLARY 4.2. *Suppose at phase j , it is guaranteed that the number of false positives in I_{j-1} is at most k_{j-1} . Then phase j can be implemented with $O(k_{j-1}l_j + k_{j-1} \log(k/k_{j-1}))$ bits in two rounds.*

Finally, a naive implementation of the protocol requires $2r$ rounds if the generic protocol has r phases. In fact, the protocol can be compressed into exactly r rounds in the following way. At the beginning, both parties agree that $I_0 = [k]$. Alice generates her $l_1|I_0|$ test bits $\hat{x}^{(1)}$ for phase 1 and communicates them to Bob; Bob first generates his own test bits $\hat{y}^{(1)}$ for phase 1 and determines I_1 , then generates $l_2|I_1|$ test bits $\hat{y}^{(2)}$ for phase 2 and transmits both $\hat{y}^{(1)}$ and $\hat{y}^{(2)}$ to Alice. Alice computes I_1 , generates $\hat{x}^{(2)}$, computes I_2 , generates $\hat{x}^{(3)}$, and then sends $\hat{x}^{(2)}$ and $\hat{x}^{(3)}$ to Bob, and so on. There is no asymptotic increase in the communication volume.

4.1.2. Reducing the number of false positives. Our protocols for EqualityTesting and ExistsEqual are divided into two parts. The goal of the first part is to reduce the number of false positives from at most k to at most E ; if $E \geq k$, we can skip this part. The details of this part are very similar to Sağlam and Tardos's SetDisjointness protocol [41].

THEOREM 8. *Let (x, y) be an instance of ExistsEqual with $|x| = |y| = k$. In $\log^*(k/E)$ rounds, we can reduce this to a new instance (x', y') of ExistsEqual where $|x'| = |y'| \leq E$, using $O(k)$ communication. The failure probability of this protocol is at most $2^{-(E+1)}$.*

For EqualityTesting, we can reduce the initial instance to a new instance (x', y') such that the Hamming distance $\text{dist}(x', y') \leq E$, with the same round complexity, communication volume, and error probability.

Proof. We first give the protocol for ExistsEqual, then apply the necessary changes to make it work for EqualityTesting.

The protocol for ExistsEqual uses our generic protocol and imposes a strict upper bound k_j on $|I_j|$. Whenever $|I_j|$ exceeds this upper bound, we halt the entire protocol and answer *yes* (there exists a coordinate where the input vectors are equal). We start by setting the parameters k_j and l_j for any $j \in [1, \log^*(k/E)]$ as follows:

$$\begin{aligned} k_0 &= k, \\ k_j &= \max \left\{ \frac{k}{2^{j-1} \exp^{(j)}(2)}, E \right\}, \\ l_j &= 3 + \exp^{(j-1)}(2). \end{aligned}$$

Note that it is reasonable to assume $k_j > E$ before the last phase, since whenever we find $k_j \leq E$, we can simply terminate the protocol prematurely after phase j , and our goal would be achieved.

Now suppose the input vectors share no equal coordinates. We know that $|I_{j-1}| \leq k_{j-1}$ at the beginning of phase j . The probability of any particular coordinate in I_{j-1} passing all tests in phase j is exactly $p_j = \exp(-l_j)$. Thus, the expected size of I_j is at most

$$k_{j-1}p_j = \frac{k}{2^{j-2} \exp^{(j-1)}(2)} \cdot \frac{1}{2^3 \exp^{(j)}(2)} \leq \frac{k}{2^{j+2} \exp^{(j)}(2)} \leq \frac{k_j}{8}.$$

Recall the statement of the usual Chernoff bound.

FACT 4.3 (see [20]). *Let $X = \sum_{i=1}^n X_i$, where each X_i is an independent and identically distributed Bernoulli random variable. Letting $\mu = \mathbb{E}[X]$, the following inequality holds for any $\delta > 0$:*

$$\Pr[X \geq (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu.$$

In our case $X_i = 1$ iff the i th coordinate in I_{j-1} survives to I_j . By linearity of expectation, $\mu \leq k_j/8$. Setting $\delta = k_j/\mu - 1 \geq 7$, we have

$$\Pr[X \geq k_j] = \Pr[X \geq (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\frac{k_j}{1+\delta}} < \left(\frac{e^7}{8^8} \right)^{k_j/8} < 2^{-1.7k_j}.$$

Hence, the probability that there are at least k_j coordinates remaining after phase j is at most $2^{-1.7k_j} \leq 2^{-1.7E}$, and the probability this happens in any phase is at most $\sum_j 2^{-1.7k_j} \leq 2^{-(E+1)}$. Notice that when x and y share at least one equal coordinate, the error probability of this protocol is 0 because if it fails to reduce the number of coordinates to E it (correctly) answers *yes*. The communication volume of the protocol is asymptotic to

$$\sum_j l_j |I_{j-1}| \leq \sum_j l_j k_{j-1} = \sum_j O(k/2^j) = O(k).$$

For `EqualityTesting`, we use the same k_j as an upper bound on the number of false positives in I_j , instead of the size of I_j . Since the number of false positives is at most k at the beginning, we can still use the same argument to show that with the same choice of k_j and l_j , after $\log^*(k/E)$ phases, the number of false positives is at most E with error probability $2^{-(E+1)}$. By Lemma 4.1, the number of bits we need to exchange in phase j is $O(k_{j-1}l_j + k_{j-1} \log(k/k_{j-1}))$. Notice that $\log(k/k_{j-1}) = j - 2 + \exp^{(j-2)}(2) = O(\log l_j)$, so the total communication volume is still $O(k)$. \square

In all of our protocols, we first apply Theorem 8 to reduce the number of coordinates (in the case of `ExistsEqual`) or false positives (in the case of `EqualityTesting`) to be at most E . This requires no communication if $E \geq k$ to begin with. Hence, with $\log^*(k/E)$ extra rounds and $O(k)$ communication, we will assume henceforth that all instances of `ExistsEqual` have $k \leq E$ and all instances of `EqualityTesting` have $\text{dist}(x, y) \leq E$.

4.2. A simple EqualityTesting protocol. In light of Theorem 8, we can assume that the input vectors to EqualityTesting are guaranteed to differ in at most $k_0 = \min\{k, E\}$ coordinates.

THEOREM 9. Fix any $k \geq 1$, $E \geq 1$, and $r \in [1, (\log k_0)/2]$, where $k_0 = \min\{k, E\}$. There exists a randomized protocol for EqualityTesting length- k vectors x, y with Hamming distance $\text{dist}(x, y) \leq k_0$ that uses r rounds, $O(k + rEk_0^{1/r})$ bits of communication, and errs with probability $p_{\text{err}} = 2^{-(E+1)}$.

Proof. We instantiate the generic protocol. The parameter l_j is the number of test bits generated per coordinate of I_{j-1} in phase j . The parameter k_j is an upper bound on the number of false positives surviving in I_j (with high probability $1 - 2^{-\Theta(E)}$).

$$\begin{aligned} k_j &= k_0^{1-j/r}, \\ l_j &= 4Ek_0^{j/r-1}. \end{aligned}$$

Now fix a phase $j \in [1, r]$ and suppose at the beginning of phase j that the number of false positives in I_{j-1} is at most k_{j-1} . By assumption this holds for $j = 1$. The probability that at least k_j false positives survive phase j is upper bounded by

$$\begin{aligned} \binom{k_{j-1}}{k_j} 2^{-k_j l_j} &\leq \left(\frac{ek_{j-1}}{k_j}\right)^{k_j} 2^{-k_j l_j} && \binom{n}{k} \leq \left(\frac{en}{k}\right)^k \\ &\leq 2^{2k_j \log(k_{j-1}/k_j) - k_j l_j} && e \leq k_0^{1/r} = \frac{k_{j-1}}{k_j} \text{ due to } r \leq \frac{\log k_0}{2} \\ &\leq 2^{-2E} && \log \frac{k_{j-1}}{k_j} \leq \frac{k_{j-1}}{k_j} = k_0^{1/r} \leq \frac{l_j}{4}. \end{aligned}$$

Thus, by a union bound, the number of false positives surviving phase j is *strictly less than* k_j , for all $j \in [1, r]$, with probability at least $1 - 2^{-(E+1)}$. In particular, there are no false positives at the end since $k_r = 1$.

Meanwhile, by Lemma 4.1, the total communication volume is $O(k + rEk_0^{1/r})$ since

$$\sum_{j=1}^r k_{j-1} l_j = 4rEk_0^{1/r},$$

and

$$\begin{aligned} \sum_{j=1}^r k_{j-1} \log \frac{k}{k_{j-1}} &= k_0 \sum_{j=0}^{r-1} \frac{1}{k_0^{j/r}} \left(\log \frac{k}{k_0} + \log k_0^{j/r} \right) \\ &\leq 2k_0 \log \frac{k}{k_0} + k_0 \sum_{j=0}^{r-1} \frac{\log k_0^{j/r}}{k_0^{j/r}} && k_0^{1/r} \geq 2^2 \text{ due to } r \leq \frac{\log k_0}{2} \\ &= O(k) && k_0^{1/r} \geq 2^2 \text{ and } k_0 \leq k. \quad \square \end{aligned}$$

Proof of Theorem 4. Applying Theorems 8 and 9 in sequence, we obtain a $(\log^*(k/E) + r)$ -round randomized protocol for EqualityTesting on vectors of length k that errs with probability $p_{\text{err}} = 2^{-E}$ and uses $O(k + rE \min\{k, E\}^{1/r})$ bits of communication. When $E \geq k$ the protocol is obtained directly from Theorem 9 and uses

$O(rEk^{1/r})$ communication. When $E < k$ the communication implied by Theorems 8 and 9 is $O(k + rE^{1+1/r}) = O(k + rEk^{1/r})$.⁶ \square

4.3. An optimal ExistsEqual protocol.

4.3.1. Overview of the protocol. In this section, we show that we can obtain a $(\log^*(k/E) + r)$ -round, $O(k + Ek^{1/r})$ -bit protocol for ExistsEqual. This matches the lower bound of Theorem 3, asymptotically, when $E \geq k$. Theorem 8 covers dimension reduction in $\log^*(k/E)$ rounds, so we assume without loss of generality that $E \geq k$ and we have exactly r rounds.

Suppose the inputs x and y share no equal coordinates. Imagine writing down all the possible results of the inner product tests in a matrix A of dimension $(E + \log k) \times k$, where $A_{q,i}$ is “=” if x_i, y_i pass the q th inner product test and “ \neq ” otherwise. By a union bound, with probability $1 - 2^{-E}$, each column contains at least one “ \neq .” Now consider the area above the first “ \neq ” in each column. The probability that this area is at least E' is, by a union bound, at most

$$(4.1) \quad \binom{E' + k - 1}{k - 1} 2^{-E'} < \exp(k \log(e(E' + k)/k) - E').$$

For $E' = E + O(k \log(E/k)) = O(E)$, this probability is $\ll 2^{-E}$. In our analysis it suffices to consider a situation where an *adversary* can decide the contents of A , subject to the constraint that its *error budget* (the area above the curve defined by the first “ \neq ” in each column) never exceeds $E' = O(E)$. The notion of an error budget is also essential for analyzing the EqualityTesting protocols of Sections 4.4 and 4.5.

In the j th phase, $j \geq 1$, our protocol exposes the fragment of A consisting of the next l_j rows of columns in I_{j-1} . The set I_j consists of those columns without any “ \neq ” exposed so far. The *communication budget* for phase j is equal to $l_j |I_{j-1}|$. In the worst case, the first exposed value in each column of $I_{j-1} \setminus I_j$ is “ \neq ,” so the adversary spends at least $l_j |I_j|$ of its *error budget* in phase j .

If we witness at least one “ \neq ” in every column, we can correctly declare there does not exist an equal coordinate and answer *no*. Otherwise, if the adversary has not exceeded his error budget but there is some column without any “ \neq ,” we answer *yes*. If the adversary ever exhausts his error budget, we terminate the protocol and answer *yes*. Recall that the notion of an error budget tacitly assumed that x and y differ in every coordinate. It is important to note that if they do *not* differ in every coordinate, the protocol answers correctly with probability 1, regardless of whether the protocol halts prematurely or not. Thus, there is nothing to prove in this case and it is fine to measure the error budget expended *as if* Alice’s and Bob’s inputs differ in all coordinates. The probability that the error budget is exhausted when x and y differ in all coordinates (causing the algorithm to incorrectly answer *yes*) is $\ll 2^{-E}$, according to (4.1).

4.3.2. Analysis. In this section we give a formal proof to the following theorem.

THEOREM 10. *Fix any $k \geq 1$, $E \geq k$, and $r \in [1, (\log k)/2]$. There exists an r -round randomized protocol for ExistsEqual on vectors of length k that errs with probability $p_{\text{err}} = 2^{-(E+1)}$, using $O(Ek^{1/r})$ bits of communication.*

⁶It appears as if $rE^{1+1/r}$ is an improvement over $rEk^{1/r}$ when $E < k$, but this is basically an illusion. In light of Remark 1, we can always dedicate $\log^*(k/E) - 2$ rounds to the first part and $r + 2$ rounds to the second part while increasing the communication by $O(k)$. When $E \geq k^{1-1/r}$, $rE^{1+1/r} = \Omega((r+2)Ek^{1/(r+2)})$, meaning there is no clear benefit to using the $rE^{1+1/r}$ expression.

Proof. The number of tests per coordinate in phase j is l_j :

$$l_j = 2Ek^{j/r-1}.$$

Define $E_j = \sum_{j'=1}^j l_{j'}|I_{j'}|$ to be the portion of the error budget spent in phases 1 through j . We can express the asymptotic communication cost of the protocol in terms of the error budget as follows:

$$\begin{aligned} \sum_{j=1}^r l_j |I_{j-1}| &\leq l_1 |I_0| + k^{1/r} \sum_{j=2}^r l_{j-1} |I_{j-1}| & l_j &= k^{1/r} l_{j-1}. \\ &\leq 2Ek^{1/r} + E_{r-1} k^{1/r} & &\text{definition of } E_{r-1}. \end{aligned}$$

Recall that the protocol terminates immediately after phase j if $E_j \geq E'$, which indicates $E_{r-1} < E'$. Hence, the total cost is bounded by

$$\leq (2E + E')k^{1/r} = O(Ek^{1/r}).$$

The protocol can err only if x and y differ in every coordinate. In this case, there are two possible sources of error. The first possibility is that the protocol answers *yes* because $|I_r| \geq 1$. By a union bound, this happens with probability at most

$$k2^{-\sum_{j=1}^r l_j} \leq k2^{-l_r} = k2^{-2E}.$$

The second possibility is that the protocol terminates prematurely and answers *yes* if $E_j \geq E'$ for some $j \in [1, r]$. The probability of this event occurring is also $\ll 2^{-E}$; see (4.1). This concludes the proof. \square

Proof of Theorem 5. Theorem 5 follows directly by combining Theorems 8 and 10. \square

4.4. A log k -round-communication optimal EqualityTesting protocol.

Suppose we want a communication optimal EqualityTesting protocol using $O(k + E)$ bits. When $E \geq k$ we need $r = \Omega(\log k)$ rounds, by Theorem 2. In this section, we give a protocol for EqualityTesting that uses $O(r)$ rounds (rather than r) and $O(Ek^{1/r})$ bits of communication, assuming $E \geq k$. Observe that when $r = \Theta(\log k)$, there is no (asymptotic) difference between r rounds and $O(r)$ rounds as this only influences the leading constant in the communication volume.

4.4.1. Overview of the protocol. The protocol uses the concept of an *error budget* introduced in section 4.3. To shave the factor r off the communication volume, we cannot afford to use $Ek^{j/r-1}$ test bits for each coordinate that participates in phase j . Consequently, we *cannot guarantee* with high probability (say, $1 - 2^{-\Theta(E)}$) that the number of false positives in I_j is less than $k^{1-j/r}$. The protocol uses a “backtracking” strategy that has been applied elsewhere, e.g., in the Feige et al. [22] binary search algorithm using unreliable comparisons.

Our protocol needs to be able to respond to the rare event that the number of false positives in I_j is larger than k_j . Notice that this type of error cannot be detected in the first j phases and is not easily detectable in the following phases. The danger in the number of false positives in I_j exceeding k_j is that when the test bits for phase $j + 1$ are exchanged using Lemma 4.1, the protocol may silently fail, with all test bits potentially corrupted.

To address these challenges, Alice and Bob each keep a *history* of all the test bits they have generated so far. They also keep a history of the test bits they have received from the other party, *which may have been corrupted*. Define T_A and T_B to be the

true history of the test bits generated by Alice and Bob, respectively. Define $T_B^{(A)}$ to be what Alice believes Bob's history to be, and define $T_A^{(B)}$ analogously. Observe that if every invocation of Lemma 4.1 succeeds, then $T_A = T_A^{(B)}$ and $T_B = T_B^{(A)}$.

To detect inconsistencies, after Alice and Bob generate and exchange their test bits for phase j , they accumulate their views of the history into strings $T^{(A)} = T_A \circ T_B^{(A)}$ and $T^{(B)} = T_A^{(B)} \circ T_B$, respectively, where \circ is the concatenation operator, and verify that $T^{(A)} = T^{(B)}$ with a certain number of inner product tests. This is called a *history check*. If the history check passes, they can proceed to phase $j + 1$. If the history check fails, then the results of phase j are junk, and we can infer that one of two types of low probability events occurred in phase $j - 1$. The first possibility is that the test bits at phase $j - 1$ were exchanged successfully (and consequently, the history check succeeded), but I_{j-1} contains more than k_{j-1} false positives. The second possibility is that Alice's and Bob's histories were already inconsistent at phase $j - 1$, but the phase- $(j - 1)$ history check failed to detect this. Notice that Alice and Bob cannot detect which of these types of errors occurred. In either case, we must undo the effects of phases j and $j - 1$ and restart the protocol at the beginning of phase $j - 1$. It may be that the history check then fails at the re-execution of phase $j - 1$, in which case we would continue to rewind to the beginning of phase $j - 2$, and so on. Being able to rewind multiple phases is important because we do not know which phase suffered the *first* error.

Both parties maintain an empirical *error meter* E'' that measures the sum of logarithms of probabilities of low probability (error) events that have been detected. If the error meter ever exceeds the error budget $E' = \Theta(E)$ we terminate the protocol, which we show occurs with probability $\ll 2^{-E}$. Thus, the process above (proceeding iteratively with phases, undoing and redoing them when errors are detected) must end by either successfully completing phase r or exceeding the error budget.

If Alice and Bob successfully finish phase r , we are still not done. This is because an error can happen in the later phases but we do not have sufficiently high $(1 - 2^{-E})$ confidence that they all succeeded. To build this confidence, Alice and Bob do inner product tests on the whole history, gradually increasing their number until $\Theta(E)$ tests have been done. If one of these history checks fails, we increase the error meter E'' appropriately and rewind the protocol to a suitable phase j in the first stage of the protocol.

Let us make every step of this protocol more quantitatively precise.

- The protocol has two stages, the *refutation stage* (in which potential equalities are refuted) and the *verification stage*, each consisting of a series of *phases*. Although the refutation stage logically precedes the verification stage, because phases can be undone, an execution of the protocol may oscillate between refutation and verification multiple times.
- The refutation stage is similar to the protocol in section 4.2 except Alice and Bob will verify whether the messages conveyed by Lemma 4.1 are successfully received with further inner product tests. The *budget* of phase j is

$$B_j = \frac{Ek_0^{1/r}}{\min\{j^2, r\}}.$$

Observe that the sum of budgets, $\sum_{j'=1}^r B_{j'}$, is $O(Ek_0^{1/r})$. Thus, in phase j , we perform $l_j = B_j/k_{j-1}$ independent inner product tests on each coordinate in I_{j-1} and exchange test bits with a Hamming distance of k_{j-1} . That is, we

are working under the assumption (perhaps false) that there are k_{j-1} false positives still in I_{j-1} . As usual, I_0 is initially $[k]$ and

$$\begin{aligned} k_0 &\leq E, \\ k_j &\leq k_0^{1-j/r}. \end{aligned}$$

All histories $T_A, T_B, T_B^{(A)}, T_A^{(B)}$ are initially empty, and the error meter E'' is initially zero.

- Phase j has two steps, the *test step* and the *history check step*. In the test step, Alice and Bob conduct inner product tests as in section 4.2, i.e., they generate l_j test bits for each coordinate in I_{j-1} and exchange them using Lemma 4.1, assuming their Hamming distance is at most k_{j-1} . Alice appends the test bits she generates onto the history T_A and appends the test bits she receives from Bob onto $T_B^{(A)}$. Bob does likewise. In the history check step, they use B_j independent inner product tests to check whether $T^{(A)} = T^{(B)}$, where $T^{(A)} = T_A \circ T_B^{(A)}$, and $T^{(B)} = T_A^{(B)} \circ T_B$. The history check *fails* if they detect inequality and *passes* otherwise. Since B_j is, in general, less than E , we are still skeptical of history checks that pass.
- If the history check for phase j passes, Alice and Bob proceed to phase $j + 1$, or proceed to the verification stage if $j = r$. Otherwise, an error has been detected: either the number of false positives in I_{j-1} is at least k_{j-1} , or the history check at phase $j - 1$ *mistakenly* passed. The latter occurs with probability $\exp(-B_{j-1})$ and we show the former occurs with probability $\exp(-3k_0^{-1/r} B_{j-1}/4)$. Not knowing which occurred, we increment the error meter E'' by $k_0^{-1/r} B_{j-1}/2$ due to a union bound. If E'' exceeds the *error budget* $E' = cE$, then we halt, where $c \geq 2$ is a suitable constant. Otherwise we retract the effects of phases j and $j - 1$ and continue the protocol at the beginning of phase $j - 1$, with “fresh” random bits so as not to recreate previous errors.
- Observe that after phase r of the refutation stage, each coordinate in I_r has passed only about $B_r/k_{r-1} = E/r$ inner product tests, which is not high enough. Before the verification stage begins, Alice and Bob each generate E' test bits for each coordinate in I_r and append them to $T^{(A)}$ and $T^{(B)}$. (This can be viewed as a degenerate instantiation of Lemma 4.1 with $d = 0$, which requires no communication.) If there are no false positives in I_r , these test bits must be identical.
- In the verification stage the phases are indexed in reverse order: $r, r - 1, \dots, 1$. In each successive phase j , Alice and Bob test the equality $T^{(A)} = T^{(B)}$ with B_j independent inner product tests. This process stops if it passes a total of E' tests, in which case they report that x and y are *equal* on I_r and *not equal* on $[k] \setminus I_r$, or some verification phase j detects that $T^{(A)} \neq T^{(B)}$. In this case, we know verification phases $r, r - 1, \dots, j + 1$ passed *in error* and that there must also have been an error in refutation phase r . Therefore, Alice and Bob increment E'' by $k_0^{-1/r} B_r/2 + \sum_{j'=j+1}^r B_{j'}$ and halt if $E'' \geq E'$. If not, they rewind the execution of the protocol to phase j of the refutation stage and continue.

Algorithm 1 recapitulates this description in the form of pseudocode from the perspective of Alice. Here $T_A[j, i]$ refers to the sequence of Alice’s test bits in T_A for the i th coordinate produced in the most recent execution of phase j , and $T_A[j_1 \cdots j_2, \cdot]$

refers to the test bits generated from phase j_1 to phase j_2 . Phase $r + 1$ refers to the $E' \times |I_r|$ test bits generated between the refutation and verification stages. $T^{(A)}[j, i]$ refers to the concatenation of $T_A[j, i]$ and $T_B^{(A)}[j, i]$.

4.4.2. Analysis. To prove Theorem 6, it suffices to prove the following Theorem 11.

THEOREM 11. *Fix any $k \geq 1$, $E \geq 1$, and $r \in [1, (\log k_0)/6]$, where $k_0 = \min\{k, E\}$. There exists a randomized protocol for EqualityTesting length- k vectors x, y with Hamming distance $\text{dist}(x, y) \leq k_0$ that uses $O(r)$ rounds, $O(k + Ek_0^{1/r})$ bits of communication, and errs with probability $p_{\text{err}} = 2^{-(E+1)}$.*

The protocol of Lemma 4.1 fails if Bob does not generate the correct $x' = x$, which indicates that the precondition is not met, i.e., $\text{dist}(x, y) > d$. Refutation phase j fails if the condition in line 27 is not satisfied and the else branch at line 37 is executed in order to resume the protocol from phase $j - 1$. Similarly, we say verification phase j fails if the condition in line 45 is not satisfied, which also indicates the else branch at line 49 is executed and the protocol is resumed from refutation phase j .

We begin the proof by showing that the extra communication caused by redoing some of the refutation/verification phases is properly covered by the total error budget. The following two lemmas actually prove that the error budget spent so far is correctly lower bounded in line 38 and line 50, and then Lemma 4.6 upper bounds the total number of extra phases by $O(r)$ and the overall extra communication by $O(k + Ek_0^{1/r})$.

LEMMA 4.4. *Fix any $j \in [2, r]$. If phase j of the refutation stage fails, then the outcome of the most recent execution of phase $j - 1$ happened with probability at most $\exp(-k_0^{-1/r} B_{j-1}/2)$.*

Proof. Recall that there are two types of errors at phase $j - 1$. If the $(j - 1)$ th history check erroneously passed, this occurred with probability $\exp(-B_{j-1})$. The probability that more than k_{j-1} false positives survive in I_{j-1} is less than

Algorithm 1. An EqualityTesting protocol for Theorem 11 (from the perspective of Alice).

```

1: procedure EQUALITYTESTING ▷ main procedure
2:    $I_0 \leftarrow [k]$ 
3:    $k_0 \leftarrow \min\{k, E\}$  ▷ initial bound on Hamming distance
4:    $E' \leftarrow cE$  ▷ error budget
5:    $E'' \leftarrow 0$  ▷ error meter
6:   for  $j \leftarrow 1, \dots, r$  do
7:      $B_j \leftarrow \frac{Ek_0^{1/r}}{\min\{j^2, r\}}$  ▷ phase  $j$  communication budget
8:      $k_j \leftarrow k_0^{1-j/r}$  ▷ ideal upper bound on Hamming distance
9:      $l_j \leftarrow B_j/k_{j-1}$  ▷ tests per coordinate
10:  end for
11:  REFUTATION(1)
12:  VERIFICATION( $r$ )
13:  Output equal on coordinates  $I_r$  and not equal on  $[k] \setminus I_r$ 
14: end procedure

```

Algorithm 1. An EqualityTesting protocol for Theorem 11 (from the perspective of Alice) (cont.).

```

15: procedure INNERPRODUCTTEST( $w, b$ )
16:   perform  $b$  independent inner product tests on  $w$  and return the test bits
17: end procedure

18: procedure REFUTATION( $j$ ) ▷ phase  $j$  of the refutation stage
19:    $T_A[j, \cdot] \leftarrow \perp$  ▷ Clear test bits for phase  $j$ 
20:   for all  $i \in I_{j-1}$  do
21:      $T_A[j, i] \leftarrow \text{INNERPRODUCTTEST}(x_i, l_j)$ 
22:   end for
23:   send  $T_A[j, \cdot]$  to Bob and receive  $T_B^{(A)}[j, \cdot]$  from Bob via Lemma 4.1
24:    $T^{(A)}[j, \cdot] \leftarrow T_A[j, \cdot] \circ T_B^{(A)}[j, \cdot]$ 
25:    $\hat{T}^{(A)} \leftarrow \text{INNERPRODUCTTEST}(T^{(A)}[1 \dots j, \cdot], B_j)$ 
26:   send  $\hat{T}^{(A)}$  to Bob and receive  $\hat{T}^{(B)}$  from Bob directly
27:   if  $\hat{T}^{(A)} = \hat{T}^{(B)}$  then ▷ passed history check
28:      $I_j \leftarrow \{i \in I_{j-1} \mid T_A[j, i] = T_B^{(A)}[j, i]\}$  ▷ all coords. not yet refuted
29:     if  $j < r$  then
30:       REFUTATION( $j + 1$ )
31:     else
32:        $T^{(A)}[r + 1, \cdot] \leftarrow \perp$ 
33:       for all  $i \in I_r$  do
34:          $T^{(A)}[r + 1, i] \leftarrow \text{INNERPRODUCTTEST}(x_i, E')$ 
35:       end for
36:     end if
37:   else
38:      $E'' \leftarrow E'' + k_0^{-1/r} B_{j-1}/2$ , and terminate if  $E'' \geq E'$  ▷ update error meter
39:     REFUTATION( $j - 1$ )
40:   end if
41: end procedure

42: procedure VERIFICATION( $j$ ) ▷ phase  $j$  of the verification stage
43:    $\hat{T}^{(A)} \leftarrow \text{INNERPRODUCTTEST}(T^{(A)}[\cdot, \cdot], B_j)$ 
44:   send  $\hat{T}^{(A)}$  to Bob and receive  $\hat{T}^{(B)}$  from Bob directly
45:   if  $\hat{T}^{(A)} = \hat{T}^{(B)}$  then
46:     if  $\sum_{j'=j}^r B_{j'} < E'$  then ▷ insufficiently confident to halt
47:       VERIFICATION( $j - 1$ )
48:     end if
49:   else ▷ error detected
50:      $E'' \leftarrow E'' + k_0^{-1/r} B_r/2 + \sum_{j'=j+1}^r B_{j'}$ , and terminate if  $E'' \geq E'$  ▷ update
error meter
51:     REFUTATION( $j$ ) ▷ rewind protocol to phase  $j$ 
52:     VERIFICATION( $r$ )
53:   end if
54: end procedure

```

$$\begin{aligned} \binom{k_0}{k_{j-1}} 2^{-k_{j-1}l_{j-1}} &\leq \left(\frac{ek_0}{k_{j-1}}\right)^{k_{j-1}} 2^{-k_{j-1}l_{j-1}} && \binom{n}{k} \leq \left(\frac{en}{k}\right)^k \\ &\leq 2^{2k_{j-1} \log(k_0/k_{j-1}) - k_{j-1}l_{j-1}} && e \leq k_0^{1/r} \leq \frac{k_0}{k_{j-1}} \text{ due to } r \leq \frac{\log k_0}{6} \\ &\leq 2^{-3l_{j-1}k_{j-1}/4}, \end{aligned}$$

where the last step follows from the inequality

$$\begin{aligned} 2 \log \frac{k_0}{k_{j-1}} &= 2(j-1) \log k_0^{1/r} \\ &\leq \frac{8^{j-1} \log k_0^{1/r}}{4(j-1)^2} && \text{because } 8x^3 \leq 8^x \text{ for } x \in \mathbb{N} \\ &\leq \frac{k_0^{(j-1)/r}}{4(j-1)^2} && \log k_0^{1/r} \leq \frac{k_0^{1/r}}{8} \text{ due to } k_0^{1/r} \geq 2^6 \\ &\leq \frac{B_{j-1}}{4k_{j-2}} && \text{definition of } B_{j-1} \\ &= \frac{l_{j-1}}{4} && \text{definition of } l_{j-1}. \end{aligned}$$

Combining the above two cases, by a union bound, the outcome of the most recent execution of phase $j - 1$ of the refutation stage happens with probability at most $\exp(-B_{j-1}) + \exp(-3l_{j-1}k_{j-1}/4) = \exp(-B_{j-1}) + \exp(-3k_0^{-1/r}B_{j-1}/4) \leq \exp(-k_0^{-1/r}B_{j-1}/2)$, as claimed. \square

LEMMA 4.5. *Fix any $j \in [1, r]$. If phase j of the verification stage fails, then the outcomes of the most recent execution of phases $r, r - 1, \dots, j + 1$ of the verification stage and phase r of the refutation stage happened with overall probability at most $\exp(-k_0^{-1/r}B_r/2 - \sum_{j'=j+1}^r B_{j'})$.*

Proof. Notice that the failure of verification phase j means all previous verification phases $r, r - 1, \dots, j + 1$ failed to detect an inconsistency in the history, which occurs with probability $\exp(-\sum_{j'=j+1}^r B_{j'})$. Meanwhile, the inconsistency is caused by an error of some type in refutation phase r , which, according to Lemma 4.4, occurs with probability at most $\exp(-k_0^{-1/r}B_r/2)$. Therefore, the outcomes of the most recent execution of verification phases $r, r - 1, \dots, j + 1$ and refutation phase r happened with overall probability at most $\exp(-k_0^{-1/r}B_r/2 - \sum_{j'=j+1}^r B_{j'})$. \square

LEMMA 4.6. *Algorithm 1 executes $O(r)$ extra refutation/verification phases and uses $O(k + Ek_0^{1/r})$ extra bits of communication.*

Proof. We first consider the total number of *extra* phases. Each failure of refutation phase j uses at least $k_0^{-1/r}B_{j-1}/2 \geq E/(2r)$ of the error budget and causes the re-execution of two phases, namely $j - 1$ and j . Similarly, each failure of verification phase j uses $k_0^{-1/r}B_r/2 + \sum_{j'=j+1}^r B_{j'} \geq (r - j + 1)E/(2r)$ of the error budget and causes the re-execution of $2(r - j + 1)$ phases. Thus, the total number of extra phases is at most $4cr = O(r)$, where the error budget $E' = cE$.

Turning to the overall *extra* communication, notice that phase j of the refutation stage has communication volume $O(B_j + k_{j-1} \log(k/k_{j-1}))$ and phase j of the verification stage has communication volume $O(B_j)$. For any $j \in [2, r]$, also notice that $B_{j-1}/B_j \leq j^2/(j - 1)^2 \leq 4 \leq k_0^{1/r}$. Thus, the communication caused by each failure

is at most $O(k_0^{1/r})$ times the error budget spent by that failure, if we temporarily ignore the $k_{j-1} \log(k/k_{j-1})$ term.

In order to upper bound the communication contributed by the $k_{j-1} \log(k/k_{j-1})$ term, observe that refutation phase j can only be repeated $O(j^2)$ times before the error budget is exhausted. Thus, the overall extra communication is upper bounded by $O(k + Ek_0^{1/r})$ since

$$\begin{aligned} & O(k_0^{1/r}) \cdot E' + \sum_{j=1}^r O(j^2) \cdot k_{j-1} \log \frac{k}{k_{j-1}} \\ &= O(k_0^{1/r}) \cdot E' + k_0 \sum_{j=1}^r \frac{O(j^2)}{k_0^{(j-1)/r}} \left(\log \frac{k}{k_0} + \log k_0^{(j-1)/r} \right) \\ &= O(k_0^{1/r}) \cdot E' + k_0 \log \frac{k}{k_0} \sum_{j=1}^r \frac{O(j^2)}{k_0^{(j-1)/r}} + k_0 \sum_{j=1}^r \frac{O(j^2) \cdot \log k_0^{(j-1)/r}}{k_0^{(j-1)/r}} \\ &= O(k + Ek_0^{1/r}). \end{aligned}$$

The last step follows from the fact that $k_0^{1/r} \geq 2^6$ and $k_0 \leq k$. \square

Now we are ready to prove Theorem 11.

Proof of Theorem 11. If there are no errors, Algorithm 1 has at most $2r$ phases and uses $O(\sum_{j=1}^r (B_j + k_{j-1} \log(k/k_{j-1}))) = O(k + Ek_0^{1/r})$ communication, where each phase can be implemented in $O(1)$ rounds. Together with Lemma 4.6, we have shown that it is an $O(r)$ -round randomized EqualityTesting protocol using $O(k + Ek_0^{1/r})$ bits of communication. Thus, it suffices to calculate the error probability of the protocol.

Consider a possible execution of the protocol, i.e., the sequence of the refutation/verification phases that are performed. It can be represented by a *unique* 0-1 string of length at most $4cr + 2r$ (by the proof of Lemma 4.6) such that each “1” corresponds to a failed phase. In particular, each execution of the protocol that terminates prematurely because $E'' \geq E'$ is represented as a 0-1 string, which occurs with probability at most $2^{-E'}$, by Lemmas 4.4 and 4.5. Hence the overall probability of terminating prematurely is $2^{4cr+2r} \cdot 2^{-E'}$.

An error can also be caused by at least one false positive surviving all E' independent inner product tests generated after refutation phase r . The probability of this happening is at most $k_0 2^{-E'}$. The last possible source of error is that all verification phases fail to detect the inequality $T^{(A)} \neq T^{(B)}$. According to line 46, the probability of this happening is at most $2^{-E'}$. Hence, the overall probability of error is upper bounded by

$$2^{4cr+2r} \cdot 2^{-E'} + k_0 2^{-E'} + 2^{-E'} = \text{poly}(k_0) 2^{-E'},$$

which is at most 2^{-E} for, say, $E' = 2E$. This concludes the proof. \square

Proof of Theorem 6. Theorem 6 subsequently follows by applying Theorems 8 and 11 in sequence. \square

4.5. A more efficient EqualityTesting protocol. Theorem 10 demonstrates that the $\Omega(Ek^{1/r})$ lower bound can be attained for ExistsEqual. Let us highlight a key property of the protocol that arises naturally in ExistsEqual but is difficult to efficiently recreate in EqualityTesting. In the first round Alice generates about $l_1 = E/k^{1-1/r}$ test bits per coordinate and sends them to Bob. Since there is no possibility of reporting *yes* ($\exists i. x_i = y_i$) in error, Bob can operate under the assumption that for all

$i.x_i \neq y_i$. Therefore, if he finds that $|I_1| = \beta_1 k^{1-1/r}$, he can infer that the adversary has expended a β_1 fraction of his error budget and *adaptively* choose the length of his message to be $\beta_1 E k^{1/r}$, i.e., we are effectively charging $k^{1/r}$ bits of communication to each of the $\beta_1 E$ units of error just spent by the adversary. In Theorem 10 this adaptivity happens transparently: the length of the j th message depends directly on the fraction β_{j-1} of the error budget expended by the adversary in round $j - 1$, even though β_{j-1} is not ever named as a parameter of the algorithm.

A key difference between `ExistsEqual` and `EqualityTesting` is that in the latter, the adversary can effectively *hide* how much of its error budget it has expended. Consider the state of Bob after receiving the first message from Alice. If he finds that $|I_1| = k/2$, there is no way to tell how many false positives are contained in I_1 and how many are true positives. In the *worst case* the number of false positives could be as high as $k^{1-1/r}$. We cannot optimistically assume the false positive number is lower,⁷ and continually using the pessimistic bound leads to $O(rEk^{1/r})$ communication. It seems that any optimal algorithm must *detect* and *adapt* to the fraction of the error budget spent by the adversary.

THEOREM 12. *Fix any $k \geq 1$, $E = \Omega(1)$, and $r \in [1, (\log k_0)/2]$, where $k_0 \leq \min\{k, E\}$. There exists a randomized protocol for `EqualityTesting` length- k vectors x, y with Hamming distance $\text{dist}(x, y) \leq k_0$ that uses r rounds, $O(k + Ek_0^{1/r} \log r + Er \log r)$ bits of communication, and errs with probability $p_{\text{err}} = 2^{-(E+1)}$.*

The remainder of this section constitutes a proof of Theorem 12.

Define $E' = 7E$ to be the *error budget* of the adversary, i.e., it is allowed to make up to E' inner product tests pass on unequal coordinates.

Round 1. Initially $I_0 = [k]$ is guaranteed to contain at most k_0 unequal coordinates. Alice generates $l_1 = E'k_0^{1/r-1}$ test bits for each coordinate in I_0 and transmits them to Bob using Lemma 4.1 with a Hamming distance of $d = k_0$. (We show later that the $k_0 \log(k/k_0)$ terms in this protocol contribute negligibly to the overall communication; thus, for the time being we measure the cost as $k_0 l_1 = E'k_0^{1/r}$.) Bob sets I_1 to be the subset of I_0 that pass all inner product tests.

Round 2. Due to the adversary's error budget, the number of false positives in I_1 is at most

$$k_1 = E'/l_1 = k_0^{1-1/r}.$$

Suppose the *true* number of false positives in I_1 is

$$k_1^* = \beta'_1 k_1 = \beta'_1 k_0^{1-1/r},$$

meaning the adversary just spent a β'_1 fraction of his total error budget. Bob cannot measure β'_1 , but he can send a message to Alice that allows *her* to estimate β'_1 . Bob invokes Lemma 4.1 $\log r$ times. For $i \in [1, \log r]$, Bob generates the next $l_2^{(i)}$ test bits for coordinates in I_1 so that Alice can recover them up to a Hamming distance of $k_1^{(i)}$, where

$$\begin{aligned} l_2 k_1 &= 2E'k_0^{1/r}, \\ l_2^{(i)} &= l_2 \cdot 2^{i-1}, \\ k_1^{(i)} &= k_1/2^{i-1}. \end{aligned}$$

⁷Invoking Lemma 4.1 with a Hamming distance d that is too small can result in an undetected failure of the protocol.

Clearly invocation i will succeed if $k_1^{(i)} \geq k_1^*$ and may fail if $k_1^{(i)} < k_1^*$. In order to detect which invocations of Lemma 4.1 succeed, Bob supplements each with a $\Theta(E)$ -bit hash of the test bits generated. Thus, with probability $1 - 2^{-\Theta(E)}$, Lemma 4.1 has no *silent* failures.

Round 3 onward. Suppose that Alice detects that the invocations of Lemma 4.1 with Hamming distances $k_1^{(1)}, \dots, k_1^{(i^*)}$ succeed but the one with $k_1^{(i^*+1)}$ fails (or that $i^* = \log r$). Alice estimates β'_1 by

$$\beta_1 = 2^{-i^*}.$$

Observe that if $i^* < \log r$ (the $(i^* + 1)$ th invocation of Lemma 4.1 fails), then

$$k_1^* \geq k_1^{(i^*+1)} = 2^{-i^*} \cdot k_0^{1-1/r}$$

and consequently, $\beta_1 \leq \beta'_1$. On the other hand, if $i^* = \log r$, then $\beta_1 = 1/r$, whereas β'_1 may be close to zero. Either way, we have

$$\beta_1 \leq \beta'_1 + 1/r.$$

Since every false positive in I_2 has successfully passed $l_2^{(i^*)}$ inner product tests, we can conclude that the maximum number of false positives remaining in I_2 is

$$k_2 = \frac{E'}{l_2^{(i^*)}} = \frac{E'}{2E'k_0^{1/r} \cdot 2^{i^*-1}/k_1} = \frac{k_1}{k_0^{1/r}2^{i^*}} = \beta_1 k_0^{1-2/r}.$$

As before, $k_2^* = \beta'_2 k_2$ is the true number of false positives in I_2 , where $\beta'_2 \in [0, 1]$ is currently unknown. In the third round Alice selects l_3 (see below) and invokes the test bit exchange protocol (Lemma 4.1) with $l_3^{(i)} = l_3 2^{i-1}$ test bits per coordinate and Hamming distance $k_2^{(i)} = k_2/2^{i-1}$, in parallel for all $i \in [1, \log r]$. The overall communication volume is $k_2 l_3 \log r$, and we select l_3 such that this is linear in the (estimated) error budget spent by the adversary in round one, i.e.,

$$k_2 l_3 \log r = \beta_1 E' k_0^{1/r} \log r$$

and therefore

$$l_3 = \beta_1 2E' k_0^{1/r} / k_2 = 2E' / k_0^{1-3/r}$$

is independent of β_1 .

All the bounds above were specialized to Round 3 but apply to round j by reindexing appropriately. In particular, the receiver of the $(j - 1)$ th message estimates $\beta'_{j-2} = k_{j-2}^* / k_{j-2}$ by

$$\beta_{j-2} = 2^{-i^*} \leq \beta'_{j-2} + 1/r$$

and sets

$$\begin{aligned} k_{j-1} &= \beta_{j-2} k_0^{1-(j-1)/r}, \\ l_j &= 2E' / k_0^{1-j/r}, \end{aligned}$$

then invokes Lemma 4.1 $\log r$ times in parallel with parameters $l_j^{(i)} = l_j 2^{i-1}$ and $k_{j-1}^{(i)} = k_{j-1} / 2^{i-1}$ to send the j th message. It remains to bound the total communication and the probability of error.

Communication volume. With the extra $\Theta(E)$ -bit hash, the cost of each invocation of Lemma 4.2 with parameters d, L is $O(E + dL + d \log(k/d))$. There are at most $\log r$ invocations per round and r rounds, so the total contributed by the first term is $O(Er \log r)$. The total contributed by the second term is

$$\begin{aligned}
 & k_0 l_1 + \sum_{j=2}^r \sum_{i=1}^{\log r} k_{j-1}^{(i)} l_j^{(i)} \\
 &= E' k_0^{1/r} + \log r \cdot \sum_{j=2}^r k_{j-1} l_j \\
 &= E' k_0^{1/r} + 2E' k_0^{1/r} \log r \cdot \left(1 + \sum_{j=3}^r \beta_{j-2} \right) \\
 &\leq E' k_0^{1/r} + 2E' k_0^{1/r} \log r \cdot \left(1 + \sum_{j=3}^r (\beta'_{j-2} + 1/r) \right) \\
 &\leq E' k_0^{1/r} + 6E' k_0^{1/r} \log r \quad \left(\sum_{j \geq 3} \beta'_{j-2} < 1 \right) \\
 &= E' k_0^{1/r} (6 \log r + 1).
 \end{aligned}$$

Next we bound the third term. For any $j \geq 2$, we have

$$\begin{aligned}
 & \sum_{i=1}^{\log r} k_{j-1}^{(i)} \log(k/k_{j-1}^{(i)}) \\
 &= \sum_{i=1}^{\log r} \frac{k_{j-1}}{2^{i-1}} \log \frac{2^{i-1} k}{k_{j-1}} \\
 &= k_{j-1} \log \frac{k}{k_{j-1}} \sum_{i=1}^{\log r} \frac{1}{2^{i-1}} + k_{j-1} \sum_{i=1}^{\log r} \frac{i-1}{2^{i-1}} \\
 &= O(k_{j-1} \log(k/k_{j-1})).
 \end{aligned}$$

Therefore, it suffices to consider only the first invocation of Lemma 4.1 from each round. Now we bound the total across all rounds. For the first two rounds we have $k_0 \log k/k_0 \leq k$ and $k_1 \log k/k_1 \leq k$, so we start counting from round $j = 3$.

$$\begin{aligned}
 & \sum_{j=3}^r k_{j-1} \log \frac{k}{k_{j-1}} \\
 &= \sum_{j=2}^{r-1} \beta_{j-1} k_0^{1-j/r} \log \frac{k}{\beta_{j-1} k_0^{1-j/r}} \\
 &\leq \sum_{j=2}^{r-1} \beta_{j-1} k \frac{\log \frac{k^{j/r}}{\beta_{j-1}}}{k^{j/r}} \\
 &\leq k \sum_{j=2}^{r-1} \frac{\log k^{j/r}}{k^{j/r}} + k \sum_{j=2}^{r-1} \frac{\beta_{j-1} \log \frac{1}{\beta_{j-1}}}{k^{j/r}} \quad \beta_j \leq 1 \\
 &\leq k \sum_{j=2}^{r-1} \frac{\log k^{j/r}}{k^{j/r}} + \frac{k}{e} \sum_{j=2}^{r-1} \frac{1}{k^{j/r}} \\
 &= O(k) \quad k^{1/r} \geq 2.
 \end{aligned}$$

In conclusion, the total communication cost is $O(k + Ek_0^{1/r} \log r + Er \log r)$.

Error probability. We now show that protocol errs with probability less than $2^{-(E+1)}$. If we use a $2E$ -bit hash of the test bits in each invocation of Lemma 4.1 the probability that any failed invocation goes unnoticed is at most

$$r \log r \cdot 2^{-2E}.$$

The algorithm works correctly so long as $k_j^* \leq k_j$ for every j , which holds whenever the adversary does not exceed his *error budget* E' . The probability that the error budget is exceeded is, by a union bound, at most

$$\binom{E' + k_0 - 1}{k_0 - 1} 2^{-E'} \leq \exp(k_0 \log(e(E' + k_0)/k_0) - E'),$$

which is less than 2^{-3E} when $E' = 7E \geq 7k_0$. Finally, every unequal coordinate is ultimately subject to $l_r = 2E'$ inner product tests, and the probability that any goes undetected is at most $k_0 2^{-2E'}$. The total error probability is therefore at most

$$r \log r \cdot 2^{-2E} + 2^{-3E} + k_0 2^{-2E'} \ll 2^{-(E+1)}.$$

This concludes the proof of Theorem 12.

5. Distributed triangle enumeration. One way to solve local triangle enumeration in the CONGEST model is to execute, in parallel, a `SetIntersection` protocol across every edge of the graph, where the set associated with a vertex is a list of its neighbors. Since there are at most $\Delta n/2$ edges, we need the `SetIntersection` error probability to be 2^{-E} , $E = \Theta(\log n)$, in order to guarantee a global success probability of $1 - 1/\text{poly}(n)$. Our lower bound says any algorithm taking this approach must take $\Omega((\Delta + E\Delta^{1/r})/\log n + r)$ rounds since each round of CONGEST allows for one $O(\log n)$ -bit message. The hardest situation seems to be when $\Delta = E = \Theta(\log n)$, in which case the optimum choice is to set $r = \log \Delta$, making the triangle enumeration algorithm run in $O(\log \Delta) = O(\log \log n)$ time. In Theorem 13 we show that it is possible to handle this situation exponentially faster, in $O(\log \log \Delta) = O(\log \log \log n)$ time, and, in general, to solve local triangle enumeration [29] in optimal $O(\Delta/\log n)$ time so long as $\Delta > \log n \log \log \log n$.

THEOREM 13. *Local triangle enumeration can be solved in a CONGEST network $G = (V, E)$ with maximum degree Δ in $O(\Delta/\log n + \log \log \Delta)$ rounds with probability $1 - 1/\text{poly}(n)$. This is optimal for all $\Delta = \Omega(\log n \log \log \log n)$.*

Proof. The algorithm consists of $\min\{\log \log \Delta, \log \log \log n\}$ phases. The goal of the first phase is to transform the original triangle enumeration problem into one with maximum degree $\Delta_1 < (\log n)^{o(1)}$, in $O(\log^* n)$ rounds of communication. The goal of every subsequent phase is to reduce the maximum degree from $\Delta' \leq \sqrt{\log n}$ to $\sqrt{\Delta'}$, in $O(1)$ rounds of communication. Thus, the total number of rounds is $O(\log \log \Delta)$ rounds if the first round is skipped and $O(\log^* n + \log \log(\Delta_1)) = O(\log \log \log n)$ otherwise.

Phase 1. Suppose $\Delta \geq \sqrt{\log n}$. Each vertex u is identified with the set $A_u = \{\text{ID}(v) \mid \{v, u\} \in E\}$ having size Δ . For each $\{u, v\} \in E$ we reduce `SetIntersection` to `EqualityTesting` by applying Theorem 1, then run the two-party `EqualityTesting` protocol of Theorem 4, with $k = \max\{\Delta, \log n\}$, $r = \log^* n$, and $E = r^{-1}k^{1-1/r}$. (That is, if $\Delta < \log n$ we imagine padding each set to size $\log n$ with dummy elements.) One undesirable property of this protocol is that it can fail “silently” if the preconditions

of Lemma 4.1 are not met. When the Hamming distance between two strings exceeds the threshold d , Bob generates a garbage string $x' \neq x$ but fails to detect this. To rectify this problem, we change the Lemma 4.1 protocol slightly: Alice sends the color $\phi(x)$ of her string, as well as an $O(\log n)$ -bit hash $h(x)$. Bob reconstructs x' as usual and terminates the protocol if $h(x) \neq h(x')$. Clearly the probability of an undetected failure (i.e., $x \neq x'$ but $h(x) = h(x')$) is $1/\text{poly}(n)$. Define $G_1 = (V, E_1)$ such that $\{u, v\} \in E_1$ iff the `SetIntersection` protocol over $\{u, v\}$ detected a failure. In other words, with high probability, all triangles in G have been discovered, except for those contained *entirely* inside G_1 . The probability that any particular edge appears in E_1 is $2^{-E} = 2^{-k^{1-1/\log^* n}/\log^* n}$ and independent of all other edges. In particular, if $\Delta \gg (\log n)^{1+1/\log^* n}$, then no errors occur, with probability $1 - 1/\text{poly}(n)$. Define Δ_1 to be the maximum degree in G_1 . Thus,

$$\begin{aligned} & \Pr [\Delta_1 \geq (\log n)^{2\epsilon}] \\ & \leq n \cdot \binom{\Delta}{(\log n)^{2\epsilon}} \cdot (2^{-E})^{(\log n)^{2\epsilon}} \qquad \epsilon = 1/r = 1/\log^* n \\ & \leq n \cdot \exp(O((\log n)^{2\epsilon} \log \log n) - \epsilon(\log n)^{1-\epsilon} \cdot (\log n)^{2\epsilon}) \\ & \leq 1/\text{poly}(n). \end{aligned}$$

Phases 2 and above. Suppose that at some round, we have detected all triangles except for those contained in some subgraph $G' = (V, E')$ having maximum degree $\Delta' < \sqrt{\log n}$. Express Δ' as $(\log n)^\gamma$, where $\gamma < 1/2$. We execute the `EqualityTesting` protocol of Theorem 9 with $k = \Delta'$, $r = 2$, and $E = C(\log n)^{1-\gamma/2}$ for a sufficiently large constant C . Note that $1 - \gamma/2 > \gamma$, so $E > k$, as required by Theorem 9. The protocol takes $O(Ek^{1/2}/\log n + r) = O(1)$ rounds since the communication volume is $O(Ek^{1/2}) = O(\log n)$ and $r = 2$. Let G'' be the subgraph of G' consisting of edges whose protocols detected a failure and Δ'' be the maximum degree in G'' . Once again,

$$\begin{aligned} & \Pr [\Delta'' \geq (\log n)^{\gamma/2}] \\ & \leq n \cdot \binom{\Delta'}{(\log n)^{\gamma/2}} \cdot (2^{-E})^{(\log n)^{\gamma/2}} \\ & \leq n \cdot \exp(O((\log n)^{\gamma/2} \log \log n) - C(\log n)^{1-\gamma/2} \cdot (\log n)^{\gamma/2}) \\ & \leq 1/\text{poly}(n). \end{aligned}$$

Thus, once $\Delta \leq \sqrt{\log n}$, $\log \log \Delta \leq \log \log \log n - 1$ of these two-round phases suffice to find all remaining triangles in G . \square

Theorem 13 depends critically on the duality between edges and `SetIntersection` instances and between edge endpoints and elements of sets. In particular, when an execution of a `SetIntersection` over $\{u, v\}$ is successful, this effectively *removes* $\{u, v\}$ from the graph, thereby removing many occurrences of `ID(u)` and `ID(v)` from adjacent sets.

Consider a slightly more general situation where we have a graph of *arboricity* λ (but unbounded Δ), witnessed by a given acyclic orientation having out-degree at most λ . Redefine the set A_u to be the set of out-neighbors of u .

$$A_u = \{\text{ID}(v) \mid \{u, v\} \in E \text{ with orientation } u \rightarrow v\}.$$

By definition $|A_u| \leq \lambda$. Because the orientation is acyclic, every triangle on $\{x, y, z\}$ is (up to renaming) oriented as $x \rightarrow y$, $x \rightarrow z$, $y \rightarrow z$. Thus, it will *only* be detectable by the `SetIntersection` instance associated with $\{x, y\}$.

THEOREM 14. *Let $G = (V, E)$ be a CONGEST network equipped with an acyclic orientation with out-degree at most λ . We can solve local triangle enumeration on G in $O(\lambda/\log n + \log \lambda)$ time.*

Proof. We apply Theorem 1 to reduce each SetIntersection instance to an EqualityTesting instance, then apply Theorem 6 with $E = \Theta(\log n)$ and $r = \log \lambda$ to solve each with $O(\lambda + E\lambda^{1/r}) = O(\lambda + E)$ communication in $O((\lambda + E)/\log n + r) = O(\lambda/\log n + \log \lambda)$ time. Note that the dependence on λ here is exponentially worse than the dependence on Δ in Theorem 13. \square

It may be that G is known to have arboricity λ , but an acyclic orientation is unavailable. The well-known “peeling algorithm” (see [16] or [4]) computes a $C\lambda$ orientation in $O(\log_C n)$ time for C sufficiently large, say, $C \geq 3$. Using this algorithm as a preprocessing step, we can solve local triangle enumeration optimally when $\lambda = \Omega(\log^2 n)$.

THEOREM 15. *Let $G = (V, E)$ be a CONGEST network having arboricity λ (with no upper bound on Δ). Local triangle enumeration can be solved in optimal $O(\lambda/\log n)$ time when $\lambda = \Omega(\log^2 n)$ and sublogarithmic time $O(\log n/\log(\log^2 n/\lambda))$ otherwise.*

Proof. The algorithm computes a $\gamma \cdot \lambda$ orientation in $O(\log_\gamma n)$ time and then applies Theorem 14 to solve local triangle enumeration in $O(\gamma\lambda/\log n + \log(\gamma\lambda))$ time. The only question is how to set γ . If $\lambda = \Omega(\log^2 n)$ we set $\gamma = 3$, making the total time $O(\lambda/\log n)$, which is optimal [29]. Otherwise we choose γ to balance the $\log_\gamma n$ and $\gamma\lambda/\log n$ terms, so that

$$\gamma \log \gamma = \log^2 n/\lambda.$$

Thus, the total running time is slightly sublogarithmic $O(\log n/\log(\log^2 n/\lambda))$. Specifically, it is $O(\log n/\log \log n)$ whenever $\lambda < \log^{2-\epsilon} n$. \square

6. Conclusions and open problems. We have established a new three-way trade-off between rounds, communication, and error probability for many fundamental problems in communication complexity such as SetDisjointness and EqualityTesting. Our lower bound is largely incomparable to the round-communication lower bounds of [41, 11] and stylistically very different from both [41] and [11]. We believe that our method *can* be extended to recover Sağlam and Tardos’s [41] trade-off (in the constant error probability regime), but with a more “direct” proof that avoids some technical difficulties arising from their round-elimination technique. It is still open whether EqualityTesting can be solved in r rounds with precisely $O(Ek^{1/r})$ communication and error probability $2^{-E} < 2^{-k}$. Our algorithms match this lower bound only when $r = O(1)$ or $r = \Omega(\log k)$, or for any r when solving the easier ExistsEqual problem.

We developed some CONGEST algorithms for triangle enumeration that employ two-party SetIntersection protocols. It is known that this strategy is suboptimal when $\Delta \gg n^{1/3}$ [13, 14]. However, for the *local* triangle enumeration problem,⁸ our $O(\Delta/\log n + \log \log \Delta)$ algorithm is optimal [29] for every $\Delta = \Omega(\log n \log \log \log n)$. Whether there are faster algorithms for triangle *detection*⁹ is an intriguing open problem. It is known that one-round LOCAL algorithms must send messages of $\Omega(\Delta \log n)$ bits deterministically [1] or $\Omega(\Delta)$ bits randomized [23]. Even for two-round triangle detection algorithms, there are no nontrivial communication lower bounds known.

⁸Every triangle must be reported by one of its three constituent vertices.

⁹At least one vertex must announce there is a triangle; there is no obligation to list them all.

Appendix A. Reductions and near equivalences. Brody et al. [11] proved that `SetIntersection` on sets of size k is reducible to `EqualityTesting` on vectors of length $O(k)$, at the cost of one round and $O(k)$ bits of communication. However, the reduction is *randomized* and fails with probability at least $\exp(-\tilde{O}(\sqrt{k}))$. This is the probability that when k balls are thrown uniformly at random into k bins, some bin contains $\omega(\sqrt{k})$ balls.

Recall the statement of Theorem 1:

$$\begin{aligned} \text{Eq}(k, r, p_{\text{err}}) &\leq \text{SetInt}(k, r, p_{\text{err}}), & \text{SetInt}(k, r + 1, p_{\text{err}}) &\leq \text{Eq}(k, r, p_{\text{err}}) + \zeta, \\ \exists \text{Eq}(k, r, p_{\text{err}}) &\leq \text{SetDisj}(k, r, p_{\text{err}}), & \text{SetDisj}(k, r + 1, p_{\text{err}}) &\leq \exists \text{Eq}(k, r, p_{\text{err}}) + \zeta, \end{aligned}$$

where $\zeta = O(k + \log \log p_{\text{err}}^{-1})$. In other words, under any error regime p_{err} , the communication complexity of `SetIntersection` and `EqualityTesting` is the same, up to one round and $O(k + \log \log p_{\text{err}}^{-1})$ bits of communication, and the same relationship holds between `SetDisjointness` and `ExistsEqual`. The proof is inspired by the probabilistic reduction of Brody et al. [11] but uses succinct encodings of perfect hash functions rather than random hash functions.

Proof of Theorem 1. The leftmost inequalities have been observed before [41, 11]. Given inputs x, y to `ExistsEqual` or `EqualityTesting`, Alice and Bob generate sets $A = \{(1, x_1), \dots, (k, x_k)\}$ and $B = \{(1, y_1), \dots, (k, y_k)\}$ before the first round of communication and then proceed to solve `SetIntersection` or `SetDisjointness` on (A, B) . Knowing $A \cap B$ or whether $A \cap B = \emptyset$ clearly allows them to determine the correct output of `EqualityTesting` or `ExistsEqual` on (x, y) .

The reverse direction is slightly more complicated. Let (A, B) be the instance of `SetIntersection` or `SetDisjointness` over a universe U with size at most $|U| = O(k^2/p_{\text{err}})$. Alice examines her set A , and picks a *perfect* hash function $h : U \mapsto [k]$ for A , i.e., h is injective on A . (This can be done in $O(k)$ time, in expectation, using only *private* randomness. In principle Alice could do this step deterministically, given sufficient time.) Most importantly, h can be described using $O(k + \log \log |U|) = O(k + \log \log p_{\text{err}}^{-1})$ bits [42], using a variant of the Fredman–Komlós–Szemerédi [24] two-level perfect hashing scheme.¹⁰ Alice sends the $O(k + \log \log p_{\text{err}}^{-1})$ -bit description of h to Bob. Bob calculates $B_j = B \cap h^{-1}(j)$ and responds to Alice with the distribution $|B_0|, |B_1|, \dots, |B_{k-1}|$, which takes at most $2k$ bits. They can now generate an instance of equality testing where the k equality tests are the pairs $A_0 \times B_0, A_1 \times B_1, \dots, A_{k-1} \times B_{k-1}$. By construction, $A_j = A \cap h^{-1}(j)$ is a one-element set. There is clearly a 1-1

¹⁰We sketch how the encoding of h works, for completeness. First, pick a function $h' : U \mapsto [O(k^2)]$ that is collision-free on A . Fredman, Komlós, and Szemerédi [24] proved that a function of the form $h'(x) = (ax \bmod p) \bmod O(k^2)$ works with constant probability, where $p = \Omega(k^2 \log |U|)$ is prime and $a \in [0, p)$ is random. Pick another function $h_* : [O(k^2)] \mapsto [k]$ that has at most twice the expected number of collisions on A , namely $2 \cdot \binom{k}{2}/k < k$, and partition A into k buckets $A_j = A \cap h_*^{-1}(j)$. The sizes $|A_0|, |A_1|, \dots, |A_{k-1}|$ can be encoded with $2k$ bits. We now pick $O(\log k)$ pairwise independent hash functions $h_1, h_2, \dots, h_{O(\log k)} : [O(k^2)] \mapsto [O(k^2)]$. For each bucket A_j , we define $h_{(j)}$ to be the function with the minimum i for which $h_{(j)}(x) = h_i(x) \bmod |A_j|^2$ is injective on A_j . In order to encode which function $h_{(j)}$ is (given that $h_1, \dots, h_{O(\log k)}$ are fixed and that $|A_j|$ is known), we simply need to write i in unary, i.e., using the bit-string $0^{i-1}1$. This takes less than 2 bits per j in expectation since each h_i is collision-free on A_j with probability at least $1/2$. Combining $h', h_*, |A_0|, \dots, |A_{k-1}|$ and $h_{(0)}, \dots, h_{(k-1)}$ into a single injective function from $U \mapsto [O(k)]$ is straightforward and done exactly as in [24]. By marking which elements in this range are actually used ($O(k)$ more bits), we can generate the perfect $h : U \mapsto [k]$ whose range has size precisely k . Encoding h' takes $O(\log k + \log \log |U|)$ bits and encoding h_* takes $O(\log k)$ bits. The distribution $|A_0|, \dots, |A_{k-1}|$ can be encoded with $2k$ bits. The functions $h_1, \dots, h_{O(\log k)}$ can be encoded in $O(\log^2 k)$ bits and the functions $h_{(0)}, \dots, h_{(k-1)}$ with less than $2k$ bits in expectation.

correspondence between equal pairs and elements in $A \cap B$. We have Bob speak first in the EqualityTesting/ExistsEqual protocol; thus, the overhead for this reduction is just one round of communication and $O(k + \log \log p_{\text{err}}^{-1})$ bits. \square

Appendix B. Section 3 proofs.

B.1. Proof of Lemma 3.1.

Proof of Lemma 3.1. Let \mathcal{X} be the whole universe. From our assumptions, the entropy of q can be lower bounded as

$$H(q) = \sum_{x \in \mathcal{X}} q(x) \log \frac{1}{q(x)}$$

and since $\sum_{x \in \mathcal{X}} q(x) = 1$, this is

$$\begin{aligned} &= \frac{1}{\alpha} \sum_{x \in \mathcal{X}} \alpha q(x) \log \frac{1}{\alpha q(x)} + \log \alpha \\ &\geq \frac{1}{\alpha} \sum_{x \in \mathcal{X}} \left[p(x) \log \frac{1}{p(x)} - (p(x) - \alpha q(x)) \log \frac{1}{p(x) - \alpha q(x)} \right] + \log \alpha. \end{aligned}$$

The previous step follows from Assumption 2 and the fact that $x \log x^{-1} + y \log y^{-1} \geq (x + y) \log(x + y)^{-1}$ for any $x, y \geq 0$. Continuing, by Assumption 1,

$$\geq \frac{1}{\alpha} \left[s - g - \sum_{x \in \mathcal{X}} (p(x) - \alpha q(x)) \log \frac{1}{p(x) - \alpha q(x)} \right] + \log \alpha,$$

and by the concavity of logarithm,

$$\begin{aligned} &\geq \frac{1}{\alpha} \left[s - g - (1 - \alpha) \log \frac{2^s}{1 - \alpha} \right] + \log \alpha \\ &= s - \frac{g}{\alpha} + \frac{1 - \alpha}{\alpha} \log(1 - \alpha) + \log \alpha = s - \frac{g}{\alpha} - \frac{H(\alpha)}{\alpha}. \quad \square \end{aligned}$$

B.2. Proof of Lemma 3.2. Recall the Kullback–Leibler divergence (also known as relative entropy) is defined to be $D_{\text{KL}}(q||p) = \sum_x q(x) \log \frac{q(x)}{p(x)}$, where $\text{supp}(q) \subseteq \text{supp}(p)$. That is, it is the expected value of $\log \frac{q(x)}{p(x)}$ when $x \sim q$. This lemma bounds the probability that $\log \frac{q(x)}{p(x)}$ deviates too far from its expectation. It is syntactically similar to Markov's inequality, but note that Markov's inequality is inapplicable as $\log \frac{q(x)}{p(x)}$ is not nonnegative.

Proof of Lemma 3.2. Define \mathcal{X}_0 to be

$$\mathcal{X}_0 = \{x \in \text{supp}(q) \mid q(x)/p(x) \leq 2^{g_1/\alpha + g_2 - (1-\alpha) \log(1-\alpha)/\alpha}\}$$

and $\mathcal{X}_1 = \text{supp}(q) \setminus \mathcal{X}_0$. Suppose, for the purpose of obtaining a contradiction, that the conclusion of the lemma is false, i.e., $q(\mathcal{X}_1) = \alpha_0$, for some $\alpha_0 > \alpha$. Notice that for each value $x \in \mathcal{X}_1$, Assumption 2 implies that

$$(B.1) \quad q(x) > p(x) \cdot 2^{g_1/\alpha + g_2 - (1-\alpha) \log(1-\alpha)/\alpha} \geq 2^{-s + g_1/\alpha - (1-\alpha) \log(1-\alpha)/\alpha}.$$

Then we can upper bound the entropy of q as follows:

$$\begin{aligned} H(q) &= \sum_{x \in \mathcal{X}_0} q(x) \log \frac{1}{q(x)} + \sum_{x \in \mathcal{X}_1} q(x) \log \frac{1}{q(x)} \\ &< \sum_{x \in \mathcal{X}_0} q(x) \log \frac{1}{q(x)} + \alpha_0 \left[s - \frac{g_1}{\alpha} + \frac{1-\alpha}{\alpha} \log(1-\alpha) \right]; \end{aligned}$$

the previous step follows from (B.1). By the concavity of logarithm, this is

$$\begin{aligned} &\leq (1-\alpha_0) \log \frac{2^s}{1-\alpha_0} + \alpha_0 \left[s - \frac{g_1}{\alpha} + \frac{1-\alpha}{\alpha} \log(1-\alpha) \right] \\ &= s - \frac{\alpha_0}{\alpha} \cdot g_1 + \alpha_0 \left[\frac{1-\alpha}{\alpha} \log(1-\alpha) - \frac{1-\alpha_0}{\alpha_0} \log(1-\alpha_0) \right] \\ &< s - g_1, \end{aligned}$$

where the last step follows from the monotonicity of $(1-\alpha) \log(1-\alpha)/\alpha$. This contradicts Assumption 1. \square

Acknowledgment. We would like to thank two anonymous reviewers for their careful reading of the paper and suggestions for improvement.

REFERENCES

- [1] A. ABBOUD, K. CENSOR-HILLEL, S. KHOURY, AND C. LENZEN, *Fooling Views: A New Lower Bound Technique for Distributed Computations Under Congestion*, CoRR, <https://arxiv.org/abs/1711.01623>, 2017.
- [2] L. BABAI, P. FRANKL, AND J. SIMON, *Complexity classes in communication complexity theory (preliminary version)*, in Proceedings of the 27th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 1986, pp. 337–347, <https://doi.org/10.1109/SFCS.1986.15>.
- [3] B. BARAK, M. BRAVERMAN, X. CHEN, AND A. RAO, *How to compress interactive communication*, SIAM J. Comput., 42 (2013), pp. 1327–1363, <https://doi.org/10.1137/100811969>.
- [4] L. BARENBOIM AND M. ELKIN, *Sublogarithmic distributed MIS algorithm for sparse graphs using Nash-Williams decomposition*, Distrib. Comput., 22 (2010), pp. 363–379.
- [5] M. BRAVERMAN, *Interactive information complexity*, SIAM Rev., 59 (2017), pp. 803–846, <https://doi.org/10.1137/17M1139254>.
- [6] M. BRAVERMAN AND A. MOITRA, *An information complexity approach to extended formulations*, in Proceedings of the 45th ACM Symposium on Theory of Computing Conference (STOC), 2013, pp. 161–170, <https://doi.org/10.1145/2488608.2488629>.
- [7] M. BRAVERMAN AND A. RAO, *Information equals amortized communication*, IEEE Trans. Inform. Theory, 60 (2014), pp. 6058–6069, <https://doi.org/10.1109/TIT.2014.2347282>.
- [8] M. BRAVERMAN, A. RAO, O. WEINSTEIN, AND A. YEHUDAYOFF, *Direct products in communication complexity*, in Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2013, pp. 746–755, <https://doi.org/10.1109/FOCS.2013.85>.
- [9] M. BRAVERMAN AND O. WEINSTEIN, *A discrepancy lower bound for information complexity*, Algorithmica, 76 (2016), pp. 846–864, <https://doi.org/10.1007/s00453-015-0093-8>.
- [10] J. BRODY, A. CHAKRABARTI, R. KONDAPALLY, D. P. WOODRUFF, AND G. YAROSLAVTSEV, *Beyond set disjointness: The communication complexity of finding the intersection*, in Proceedings of the 33rd ACM Symposium on Principles of Distributed Computing (PODC), 2014, pp. 106–113, <https://doi.org/10.1145/2611462.2611501>.
- [11] J. BRODY, A. CHAKRABARTI, R. KONDAPALLY, D. P. WOODRUFF, AND G. YAROSLAVTSEV, *Certifying equality with limited interaction*, Algorithmica, 76 (2016), pp. 796–845, <https://doi.org/10.1007/s00453-016-0163-6>.
- [12] H. BUHRMAN, D. GARCÍA-SORIANO, A. MATSLIAH, AND R. DE WOLF, *The non-adaptive query complexity of testing k -parities*, Chic. J. Theoret. Comput. Sci., 2013 (2013).
- [13] Y. CHANG, S. PETTIE, AND H. ZHANG, *Distributed triangle detection via expander decomposition*, in Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2019, pp. 821–840, <https://doi.org/10.1137/1.9781611975482.51>.

- [14] Y. CHANG AND T. SARANURAK, *Improved distributed expander decomposition and nearly optimal triangle enumeration*, in Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing (PODC), 2019, pp. 66–73, <https://doi.org/10.1145/3293611.3331618>.
- [15] A. CHATTOPADHYAY AND T. PITASSI, *The story of set disjointness*, SIGACT News, 41 (2010), pp. 59–85, <https://doi.org/10.1145/1855118.1855133>.
- [16] N. CHIBA AND T. NISHIZEKI, *Arboricity and subgraph listing algorithms*, SIAM J. Comput., 14 (1985), pp. 210–223.
- [17] A. CZUMAJ AND C. KONRAD, *Detecting cliques in CONGEST networks*, in Proceedings of the 32nd International Symposium on Distributed Computing (DISC), LIPIcs Leibniz Int. Proc. Inform. 121, Schloss Dagstuhl Leibniz-Zentrum für Informatik, Wadern, 2018, pp. 16:1–16:15.
- [18] A. DASGUPTA, R. KUMAR, AND D. SIVAKUMAR, *Sparse and lopsided set disjointness via information theory*, in Proceedings of the 15th International Workshop on Approximation, Randomization, and Combinatorial Optimization (APPROX), 2012, pp. 517–528, https://doi.org/10.1007/978-3-642-32512-0_44.
- [19] A. DRUCKER, F. KUHN, AND R. OSHMAN, *On the power of the congested clique model*, in Proceedings of the 33rd ACM Symposium on Principles of Distributed Computing (PODC), 2014, pp. 367–376.
- [20] D. P. DUBHASHI AND A. PANCONESI, *Concentration of Measure for the Analysis of Randomized Algorithms*, Cambridge University Press, Cambridge, UK, 2009.
- [21] T. FEDER, E. KUSHILEVITZ, M. NAOR, AND N. NISAN, *Amortized communication complexity*, SIAM J. Comput., 24 (1995), pp. 736–750, <https://doi.org/10.1137/S0097539792235864>.
- [22] U. FEIGE, D. PELEG, P. RAGHAVAN, AND E. UPFAL, *Computing with unreliable information*, in Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC), 1990, pp. 128–137, <https://doi.org/10.1145/100216.100230>.
- [23] O. FISCHER, T. GONEN, F. KUHN, AND R. OSHMAN, *Possibilities and impossibilities for distributed subgraph detection*, in Proceedings of the 30th Symposium on Parallelism in Algorithms and Architectures (SPAA), 2018, pp. 153–162, <https://doi.org/10.1145/3210377.3210401>.
- [24] M. L. FREDMAN, J. KOMLÓS, AND E. SZEMERÉDI, *Storing a sparse table with $O(1)$ worst case access time*, J. ACM, 31 (1984), pp. 538–544, <https://doi.org/10.1145/828.1884>.
- [25] T. GONEN AND R. OSHMAN, *Lower bounds for subgraph detection in the CONGEST model*, in Proceedings of the 21st International Conference on Principles of Distributed Systems (OPODIS), LIPIcs Leibniz Int. Proc. Inform. 95, Schloss Dagstuhl Leibniz-Zentrum für wadern, 2018, pp. 6:1–6:16.
- [26] M. GÖÖS, T. PITASSI, AND T. WATSON, *The landscape of communication complexity classes*, Comput. Complex., 27 (2018), pp. 245–304, <https://doi.org/10.1007/s00037-018-0166-6>.
- [27] M. GÖÖS AND T. WATSON, *Communication complexity of set-disjointness for all probabilities*, Theory Comput., 12 (2016), pp. 1–23, <https://doi.org/10.4086/toc.2016.v012a009>.
- [28] J. HÅSTAD AND A. WIGDERSON, *The randomized communication complexity of set disjointness*, Theory Comput., 3 (2007), pp. 211–219, <https://doi.org/10.4086/toc.2007.v003a011>.
- [29] T. IZUMI AND F. L. GALL, *Triangle finding and listing in CONGEST networks*, in Proceedings of the 36th ACM Symposium on Principles of Distributed Computing (PODC), 2017, pp. 381–389, <https://doi.org/10.1145/3087801.3087811>.
- [30] R. JAIN, A. PERESZLÉNYI, AND P. YAO, *A direct product theorem for two-party bounded-round public-coin communication complexity*, Algorithmica, 76 (2016), pp. 720–748, <https://doi.org/10.1007/s00453-015-0100-0>.
- [31] B. KALYANASUNDARAM AND G. SCHNITGER, *The probabilistic communication complexity of set intersection*, SIAM J. Discrete Math., 5 (1992), pp. 545–557, <https://doi.org/10.1137/0405044>.
- [32] H. KLAUCK, *A strong direct product theorem for disjointness*, in Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC), 2010, pp. 77–86, <https://doi.org/10.1145/1806689.1806702>.
- [33] J. H. KORHONEN AND J. RYBICKI, *Deterministic subgraph detection in broadcast CONGEST*, in Proceedings of the 21st International Conference on Principles of Distributed Systems (OPODIS), LIPIcs Leibniz Int. Proc. Inform. 95, Schloss Dagstuhl Leibniz-Zentrum Informatik, Wadern, 2018, pp. 4:1–4:16.
- [34] E. KUSHILEVITZ AND N. NISAN, *Communication Complexity*, Cambridge University Press, Cambridge, UK, 1997.
- [35] L. LOVASZ, *Communication Complexity: A Survey*, Technical Report TR-204-89, Computer Science Department, Princeton University, Princeton, NJ, 1989.

- [36] K. MEHLHORN AND E. M. SCHMIDT, *Las Vegas is better than determinism in VLSI and distributed computing*, in Proceedings of the 14th Annual ACM Symposium on Theory of Computing (STOC), 1982, pp. 330–337, <https://doi.org/10.1145/800070.802208>.
- [37] V. NIKISHKIN, *Amortized communication complexity of an equality predicate*, in Proceedings of the 8th International Computer Science Symposium in Russia (CSR), Lecture Notes in Comput. Sci. 7913, Springer, New York, 2013, pp. 212–223, https://doi.org/10.1007/978-3-642-38536-0_19.
- [38] A. RAO AND A. YEHUDAYOFF, *Communication Complexity*, unpublished manuscript.
- [39] A. A. RAZBOROV, *On the distributional complexity of disjointness*, Theoret. Comput. Sci., 106 (1992), pp. 385–390, [https://doi.org/10.1016/0304-3975\(92\)90260-M](https://doi.org/10.1016/0304-3975(92)90260-M).
- [40] T. ROUGHGARDEN, *Communication complexity (for algorithm designers)*, Found. Trends Theoret. Comput. Sci., 11 (2016), pp. 217–404, <https://doi.org/10.1561/04000000076>.
- [41] M. SAĞLAM AND G. TARDOS, *On the communication complexity of sparse set disjointness and exists-equal problems*, in Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2013, pp. 678–687, <https://doi.org/10.1109/FOCS.2013.78>.
- [42] J. P. SCHMIDT AND A. SIEGEL, *The spatial complexity of oblivious k -probe hash functions*, SIAM J. Comput., 19 (1990), pp. 775–786, <https://doi.org/10.1137/0219054>.
- [43] R. SHALTIEL, *Towards proving strong direct product theorems*, Comput. Complex., 12 (2003), pp. 1–22, <https://doi.org/10.1007/s00037-003-0175-x>.
- [44] A. C. YAO, *Probabilistic computations: Toward a unified measure of complexity (extended abstract)*, in Proceedings of the 18th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 1977, pp. 222–227, <https://doi.org/10.1109/SFCS.1977.24>.
- [45] A. C. YAO, *Some complexity questions related to distributive computing (preliminary report)*, in Proceedings of the 11th Annual ACM Symposium on Theory of Computing (STOC), 1979, pp. 209–213, <https://doi.org/10.1145/800135.804414>.