

Lower Bounds for Number-in-Hand Multiparty Communication Complexity, Made Easy*

Jeff M. Phillips
School of Computing
University of Utah
jeffp@cs.utah.edu

Elad Verbin
Dept. of Computer Science
Aarhus University, Denmark
elad.verbin@gmail.com

Qin Zhang
MADALGO
Dept. of Computer Science
Aarhus University, Denmark
qinzhang@cs.au.dk

Abstract

In this paper we prove lower bounds on randomized multiparty communication complexity, both in the *blackboard model* (where each message is written on a blackboard for all players to see) and (mainly) in the *message-passing model*, where messages are sent player-to-player. We introduce a new technique for proving such bounds, called *symmetrization*, which is natural, intuitive, and often easy to use.

For example, for the problem where each of k players gets a bit-vector of length n , and the goal is to compute the coordinate-wise XOR of these vectors, we prove a tight lower bound of $\Omega(nk)$ in the blackboard model. For the same problem with AND instead of XOR, we prove a lower bound of roughly $\Omega(nk)$ in the message-passing model (assuming $k \leq n/3200$) and $\Omega(n \log k)$ in the blackboard model. We also prove lower bounds for bit-wise majority, for a graph-connectivity problem, and for other problems; the technique seems applicable to a wide range of other problems as well. The obtained communication lower bounds imply new lower bounds in the *functional monitoring model* [11] (also called the *distributed streaming model*). All of our lower bounds allow randomized communication protocols with two-sided error.

We also use the symmetrization technique to prove several direct-sum-like results for multiparty communication.

1 Introduction

In this work we consider multiparty communication complexity in the *number-in-hand* model. In this model, there are k players $\{p_1, \dots, p_k\}$, each with his own n -bit input $x_i \in \{0, 1\}^n$. The players wish to collaborate in order to compute a joint function of their inputs, $f(x_1, \dots, x_k)$. To do so, they are allowed to communicate, until one of them figures out the value of $f(x_1, \dots, x_k)$ and returns it. All players are assumed to have unlimited computational power, so all we care about is the amount of communication used. There are three variants to this model, according to the mode of communication:

1. the *blackboard model*, where any message sent by a player is written on a blackboard visible to all players;
2. the *message-passing model*, where a player p_i sending a message specifies another player p_j that will receive this message;
3. the *coordinator model*, where there is an additional $(k+1)$ -th player called the *coordinator*, who receives no input. Players can only communicate with the coordinator, and not with each other directly.

We will work in all of these, but will mostly concentrate on the message-passing model and the coordinator model. Note that the coordinator model is almost equivalent to the message-passing model, up to a $\log k$ multiplicative factor, since instead of player i sending message x to player j , player i can transmit message (j, x) to the coordinator, and the coordinator forwards it to player j .

Lower bounds in the three models above are useful for proving lower bounds on the space usage of streaming algorithms, and for other models as well, as we explain in Section 1.2. Most previous lower bounds have been proved in the blackboard model, but lower bounds in the message-passing model and the coordinator model suffice for all the applications we have checked (again, see more in Section 1.2).

Note that another, entirely different, model for multiparty communication is the *number-on-forehead* model, where each player can see the inputs of all other players but *not* his own input. This model has important applications for circuit complexity (see e.g. [23]). We do not discuss this model at all in this paper.

We allow all protocols to be randomized, with public coins, i.e. all players have unlimited access to a common infinite string of independent random bits. We allow the protocol to return the wrong answer with

*MADALGO is the Center for Massive Data Algorithmics - a Center of the Danish National Research Foundation. JMP: was supported in part by a subaward to the University of Utah under NSF award 0937060 to the Computing Research Association.

probability ε (which should usually be thought of as a small constant); here, the probability is taken over the sample space of public randomness. Note that the public coin model might seem overly powerful, but in this paper we are mainly interested in proving lower bounds rather than upper bounds, so giving the model such strength only makes our results stronger.

For more on communication complexity, see the book of Kushilevitz and Nisan [23], and the references therein. We give some more definitions in the preliminaries, in Section 2.

1.1 Warm-Up We begin by sketching two lower bounds obtained using our technique, both of them for the coordinate-wise k -XOR problem. These lower bounds can be proved without using symmetrization, but their proofs that use symmetrization are particularly appealing.

First consider the following problem: Each player gets a bitvector $x_i \in \{0, 1\}^n$ and the goal is to compute the coordinate-wise XOR of these vectors. We operate in the *blackboard model*, where messages are posted on a blackboard for all to see.

THEOREM 1.1. *The coordinate-wise k -XOR problem requires communication $\Omega(nk)$ in the blackboard model.*

To see this, first let us specify the *hard distribution*: we prove the lower bound when the input is drawn from this distribution, and by the easy direction of Yao’s Minimax Lemma (see e.g. [23]), it follows that this lower bound applies for the problem as a whole. The hard distribution we choose is just the distribution where the inputs are independently drawn from the uniform distribution.

To prove the lower bound, consider a protocol P for this k -player problem, which works on this distribution, communicates $C(P)$ bits in expectation, and suppose for now that it never makes any errors (it will be easy to remove this assumption). We build from P a new protocol P' for a 2-player problem. In the 2-player problem, suppose that Alice gets input x and Bob gets input y , where $x, y \in \{0, 1\}^n$ are independent random bitvectors. Then P' works as follows: Alice and Bob randomly choose two distinct indices $i, j \in \{1, \dots, k\}$ using the public randomness, and they simulate the protocol P , where Alice plays player i and lets $x_i = x$, Bob plays player j and lets $x_j = y$, and they both play all of the rest of the players; the inputs of the rest of the players is chosen from shared randomness. Alice and Bob begin simulating the running of P . Every time player i should speak, Alice sends to Bob the message that player i was supposed to write on the board, and vice versa. When any

other player p_r ($r \neq i, j$) should speak, both Alice and Bob know his input so they know what he should be writing on the board, thus no communication is actually needed (this is the key point of the symmetrization technique). A key observation is that the inputs of the k players are uniform and independent and thus entirely symmetrical,¹ and since the indices i and j were chosen uniformly at random, then the expected communication performed by the protocol P' is $\mathbf{E}[C(P')] = 2C(P)/k$. Furthermore, when the players have finished simulating P' , Alice now knows the bit-wise XOR of all the vectors x_1, \dots, x_k , from which she can easily reconstruct the vector x_j (since she already knows all the other vectors x_r for $r \neq j$). It follows that using $2C(P)/k$ expected communication, Alice has managed to reconstruct Bob’s entire input; from an easy information-theoretic argument it follows that $2C(P)/k \geq n$, so $C(P) \geq \Omega(nk)$, proving the theorem.

Extending the above argument to cover the case where the protocol P is allowed to return the wrong answer with probability ε is also easy, simply by showing that if Alice managed to learn Bob’s entire bit-vector with probability $1 - \varepsilon$, then $(1 - \varepsilon) \cdot n$ bits must have been communicated: this also follows easily from information-theoretic arguments.

Note that crucial fact that the hard distribution we chose is symmetric: if the distribution of the inputs to the k players was not symmetric, then the protocol P could try to deduce the indices i and j from some statistical properties of the observed inputs, and act according to that. Then the best upper bound we could get on the communication complexity of P' would be $C(P') \leq C(P)$, which is much too weak.

We have just described the version of the symmetrization method for the blackboard model. A similar (slightly more complicated) line of thinking leads to a lower bound of $\Omega(n \log k)$ on the complexity of the coordinate-wise AND problem in the blackboard model; see Section 3.1.

Let us now sketch the symmetrization method as it is used in the *coordinator model*, where players can only send and receive messages to and from the coordinator. We prove a lower bound on the same problem as above, the coordinate-wise k -XOR problem.

THEOREM 1.2. *The coordinate-wise k -XOR problem requires communication $\Omega(nk)$ in the coordinator model.*

¹Here and throughout the paper, *symmetric* means that the inputs are drawn from a distribution where renaming the players does not change the distribution. Namely, a distribution D over X^n is called *symmetric* if exchanging any two coordinates in D keeps the distribution the same.

Note that this theorem actually follows from Theorem 1.1, since the blackboard model is stronger than the coordinator model. However, the proof we sketch here is useful as a warmup exercise, since it shows an easy example of the symmetrization technique as applied to the coordinator model. In the paper we prove multiple lower bounds using this technique, most of which do not follow from corresponding bounds in the blackboard model.

To prove this theorem, we use the same hard distribution as above: all inputs are uniform and independent. Let P be a protocol in the coordinator model, which computes the coordinate-wise XOR, uses communication $C(P)$ and for now assume that P never makes any errors. As before, we build a new protocol P' for a 2-player problem. In the 2-player problem, suppose that Alice gets input x and Bob gets input y , where $x, y \in \{0, 1\}^n$ are independent random bitvectors. Then P' works as follows: Alice and Bob choose a single index $i \in \{1, \dots, k\}$ from public randomness. Alice and Bob simulate the protocol P , where Alice simulates player i and lets $x_i = x$, and Bob plays *all the rest of the players*, including the coordinator, and chooses their inputs uniformly, conditioned on their XOR being equal to y .

To simulate the protocol P , whenever player i needs to send a message to the coordinator, then Alice sends a message to Bob, and whenever the coordinator needs to send a message to player i , Bob sends a message to Alice. Whenever any player $j \neq i$ needs to speak to the coordinator, no communication is needed, since both are played by Bob. Note again that the distribution of the inputs of the k players is independent uniform (for this it is crucial to remember that x and y were uniform and independent in the first place). Once again, from reasons of symmetry, since the index i was chosen uniformly and the inputs are symmetric, the expected communication performed by the protocol P' is $\mathbf{E}[C(P')] \leq 2C(P)/k$. Furthermore, at the end of the running of P' , Alice knows the value of $x \oplus y$ so she can reconstruct the value of y . As before, this implies the theorem. The assumption that we never make errors can once again be easily removed.

1.1.1 Discussion We see that the crux of the symmetrization technique in the coordinator model is to consider the k -player problem that we wish to lower-bound, to find a symmetric distribution which seems hard for it, to give Alice the input of one player (chosen at random) and Bob the input of all other players, and to prove a lower bound for this two-player problem. If the lower bound for the two player problem is L , the lower bound for the k -player problem will be kL . For the blackboard model, the proofs have the same outline,

except in the 2-player problem Alice gets the input of one randomly chosen player, Bob gets the input of another, and they both get the inputs of all the rest of the players. There is one important thing to note here:

- This argument only works when the hard distribution is symmetric.

1.2 Motivation, Previous Work and Related Models

Communication complexity is a widely studied topic. In multiplayer communication complexity, the most studied mode of communication is the blackboard model. The message-passing model was already considered in [15]. (This model can also be called the *private-message model*, but note that this name was used in [18, 16] for a different model.) The coordinator model can be thought of as a server-site setting², where there is one server and k sites. Each site has gathered n bits of information and the server wants to evaluate a function on the collection of these $k \cdot n$ bits. Each site can only communicate with the server, and a server can communicate with any site. This server-site model has been widely studied in the databases and distributed computing communities. Work includes computing top- k [7, 26, 29] and heavy hitters [36, 19].

Another closely related model is the *distributed streaming model*, in which we also have one server and k sites. The only difference is that now the computation is dynamic. That is, each site receives a stream of elements over time and the server would like to maintain continuously at all times some function f of all the elements in the k sites. Thus the server-site model can be seen as a one-shot version of the distributed streaming setting. It follows that any communication complexity lower bound in the message-passing model or the coordinator model also hold in the distributed streaming model. A lot of work on distributed streaming has been done recently in the theory community and the database community, including maintaining random samplings [12], frequency moments [9, 11], heavy hitters [4, 22, 24, 34], quantiles [10, 34], entropy [3], various sketches [13, 10] and some non-linear functions [31, 32].

We will come back to the latter two models in Section 6.

It is interesting to note that despite the large number of upper bounds (i.e. algorithms, communication protocols) in the above models, very few lower bounds have been proved in any of those models, likely because there were few known techniques to prove such results.

A further application of the message-passing model could be for Secure Multiparty Computation: in this

²This terminology is similar as the standard “client-server”, and is used extensively in the literature.

model, there are several players who do not trust each other, but want to compute a joint function of their inputs, with each of them learning nothing about the inputs of the others players except what can be learned from the value of the joint function. Obviously, any lower bound in the message passing model immediately implies a lower bound on the amount of communication required for Secure Multiparty Computation. For more on this model, see e.g. [17].

One final application is for the streaming model. In this model, there is a long stream of data that can only be scanned from left to right. The goal is to compute some function of the stream, and minimize the space usage. It is easy to see that if we partition the stream into k parts and give each part to a different player, then a lower bound of L on the communication complexity of the problem in the coordinator model implies a lower bound of L/k on the space usage. When t passes over the model are allowed, a lower bound of L in the coordinator model translates to a lower bound of L/tk in the streaming model.

1.3 Our Results and Paper Outline Our main technical result in this paper are lower bounds of $\Omega(nk)$ randomized communication for the bitwise k -party AND, OR, and MAJ (majority) functions in the coordinator model. These sidestep clever upper bound techniques (e.g. Slepian-Wolf coding) and can be found in Section 3. In the same section we prove some lower bounds for AND and OR in the blackboard model as well. Back to the coordinator model, we show that the connectivity problem (given k players with subgraphs on a common set of nodes, determine if it is connected) requires $\Omega(nk/\log^2 k)$ communication. This is in Section 4.

The coordinate-wise lower bounds imply lower bounds for the well-studied problems of distinct elements, ε -approximate heavy-hitters, and ε -kernels in the server-site model (or the other related models). We show any randomized algorithm requires at least $\Omega(nk)$, $\Omega(n/\varepsilon)$ and $\Omega(k/\varepsilon^{(d-1)/2})$ communication, respectively. The latter is shown to be tight. This is in Section 6.

We give some direct-sum-like results in Section 5.

2 Preliminaries

In this section we review some basic concepts and definitions.

We denote $[n] = \{1, \dots, n\}$. All logarithms are base-2 unless noted otherwise.

Communication complexity. Consider two players Alice and Bob, given bit vectors A and B , respectively. Communication complexity (see for example the book

[23]) bounds the communication between Alice and Bob that is needed to compute some function $f(A, B)$. The *communication complexity* of a particular protocol is the maximal number of bits that it communicated, taken in the worst case over all pairs of inputs. The *communication complexity* of the problem f is the best communication complexity of P , taken over all protocols P that correctly compute f .

Certain functions (such as $f = \text{EQ}$ which determines if A equals B) can be computed with less communication if randomness is permitted. Let $R^\varepsilon(f)$ denotes the communication complexity when the protocol is allowed to make a mistake with probability ε . The error is taken over the randomness used by the protocol.

Sometimes we are interested in the case where the input of Alice and Bob is drawn from some distribution μ over pairs of inputs. We want to allow an error ε , this time taken over the choice of the input. The worst-case communication complexity in this case is denoted by $D_\mu^\varepsilon(f)$. Yao [33] showed that $R^\varepsilon(f) \geq \max_\mu D_\mu^\varepsilon(f)$, thus in order to prove a lower bounds for randomized protocols it suffices to find a hard distribution and prove a distributional lower bound for it. This is called the Yao Minimax Principle.

In this paper we use an uncommon notion of *expected distributional communication complexity*. In this case we consider the distributional setting as in the last paragraph, but this time consider the expected cost of the protocol, rather than the worst-case cost; again, the expectation is taken over the choice of input. In an abuse of notation, we denote this $\mathbf{E}[D_\mu^\varepsilon(f)]$.

2.1 Two-Party Lower Bounds We state a couple of simple two-party lower bounds that will be useful in our reductions.

2-BITS. Let ζ_ρ be a distribution over bit-vectors of length n , where each bit is 1 with probability ρ and 0 with probability $1 - \rho$. In this problem Alice gets a vector drawn from ζ_ρ , Bob gets a subset S of $[n]$ of cardinality $|S| = m$, and Bob wishes to learn the bits of Alice indexed by S .

The proof of the following lemma is in Appendix A.

LEMMA 2.1. $\mathbf{E}[D_{\zeta_\rho}^{1/3}(2\text{-BITS})] = \Omega(n\rho \log(1/\rho))$.

2-DISJ. In this problem Alice and Bob each have an n -bit vector. If we view vectors as sets, then each of them has a subset of $[n]$ corresponding to the 1 bits. Let x be the set of Alice and y be the set of Bob. It is promised that $|x \cap y| = 1$ or 0. The goal is to return 1 if $x \cap y \neq \emptyset$, and 0 otherwise.

We define the input distribution μ as follows. Let $l = (n+1)/4$. With probability $1/t$, x and y are random

subsets of $[n]$ such that $|x| = |y| = l$ and $|x \cap y| = 1$. And with probability $1 - 1/t$, x and y are random subsets of $[n]$ such that $|x| = |y| = l$ and $x \cap y = \emptyset$. Razborov [30] (see also [20]) proved that for $t = 4$, $D_\mu^{1/100t}(2\text{-DISJ}) = \Omega(n)$. In the following theorem we extend this result to general t and also to the expected communication complexity. In Section 3.1 we only need $t = 4$, and in Section 4 we will need general t .

The proof for the following lemma is in Appendix B.

LEMMA 2.2. *When μ has $|x \cap y| = 1$ with probability $1/t$ then $\mathbf{E} \left[D_\mu^{1/100t}(2\text{-DISJ}) \right] = \Omega(n)$.*

3 Bit-wise Problems

3.1 Multiparty AND/OR We now consider multiparty AND/OR (below we use k -AND and k -OR for short). In the k -AND problem, each player i ($1 \leq i \leq k$) has an n -bit vector I_i and we want to establish the bitwise AND of I_i , that is, $f_j(I_1, \dots, I_k) = \bigwedge_i I_{i,j}$ for $j = \{1, \dots, n\}$. k -OR is similarly defined with OR. Observe that the two problems are isomorphic by $f_j(I_1, \dots, I_k) = \neg g_j(\bar{I}_1, \dots, \bar{I}_k)$ for $j = \{1, \dots, n\}$, where \bar{I}_i is obtained by flipping all bits of I_i . Therefore we only need to consider one of them. Here we discuss k -OR.

3.1.1 Idea for the k -OR Lower Bound in the Coordinator Model We now discuss the hard distribution and sketch how to apply the symmetrization technique for the k -OR problem in the coordinator model. The formal proof can be found in the next subsection.

We in fact start by describing two candidate hard distributions that *do not* work. The reasons they do not work are interesting in themselves. Throughout this subsection, assume for simplicity that $k \geq 100 \log n$.

The most natural candidate for a hard distribution is to make each entry equal to 1 with probability $1/k$. This has the effect of having each bit in the output vector be roughly balanced, which seems suitable for being a hard case. This is indeed the hard distribution for the blackboard model, but for the coordinator model (or the message-passing model) it is an easy distribution: Each player can send his entire input to the coordinator, and the coordinator can figure out the answer. The entropy of each player's input is only $\Theta((n \log k)/k)$, so the total communication would be $\Theta(n \log k)$ in expectation using e.g. Shannon's coding theorem;³ this is much smaller than the lower bound we

wish to prove. Clearly, we must choose a distribution where each player's input has entropy $\Omega(n)$. This is the first indication that the k -player problem is significantly different than the 2-player problem, where the above distribution is indeed the hard distribution.

The next candidate hard distribution is to randomly partition the n coordinates into two equal-sized sets: The *important set*, where each entry is equal to 1 with probability $1/k$, and the *balancing set*, where all entries are equal to 1. Now the entropy of each player's input is $\Theta(n)$, and the distribution seems like a good candidate, but there is a surprising upper bound for this distribution: the coordinator asks $100 \log n$ players to send him their entire input, and from this can easily figure out which coordinates are in the balancing set and which are in the important set. Henceforth, the coordinator knows this information, and only needs to learn the players' values in the important set, which again have low entropy. We would want the players to send these values, but the players themselves do not know which coordinates are in the important set, and the coordinator would need to send nk bits to tell all of them this information. However, they do not need to know this in order to get all the information across: using a protocol known as *Slepian-Wolf coding* (see e.g. [14]) the players can transmit to the coordinator all of their values in the important coordinates, with only $n \log k$ total communication (and a small probability of error). The idea is roughly as follows: each player p_i chooses $100n \log k/k$ sets $S_{i,j} \subseteq [n]$ independently and uniformly at random from public randomness. For each j , the player XORs the bits of his input in the coordinates of $S_{i,j}$, and sends the value of this XOR to the coordinator. The coordinator already knows the balancing set, so he only has $\Theta(n \log k/k)$ bits of uncertainty about player i 's input, meaning he can reconstruct the player's input with high probability (say by exhaustive search). The upper bound follows.

To get an actual hard distribution, we fix the hard distribution from the last paragraph. We randomly partition the n coordinates into two equal-sized sets: The *important set*, where each entry is equal to 1 with probability $1/n$, and the *noise set*, where each entry is equal to 1 with probability $1/2$.⁴ Clearly, each player's input has entropy $\Theta(n)$. Furthermore, the coordinator can again cheaply figure out which coordinates are in

³Alternatively can skip them entirely, as they are inconsequential for the remainder of the paper and for understanding the symmetrization technique.

⁴We could have chosen each entry in the important set to be equal to 1 with probability $1/k$ as well, but choosing a value of $1/n$ makes the proofs easier. The important thing is to choose each of the noise bits to be equal to 1 with probability $1/2$.

³To show the upper bounds in this subsection we use some notions from information theory without giving complete background for them. The reader can refer to e.g. [14], or alterna-

the important set and which are in the noise set, but the players do not know this information, and nothing like Slepian-Wolf coding exists to help them transmit the information to the coordinator. The distribution that we use in our formal proof is a little different than this, for technical reasons, but this distribution is hard as well.

We now sketch how to apply the symmetrization technique to prove that this distribution is indeed hard. To apply the symmetrization technique, we imagine giving Alice the input of one of the players, and to Bob the input of all of the others. Bob plays the coordinator, so Bob needs to compute the output; the goal is to prove a lower bound of $\Omega(n)$ on the communication complexity between Alice and Bob. What can Bob deduce about the answer? He can immediately take the OR of all the vectors that he receives, getting a vector where with good probability all of the noise bits are equal to 1. (Recall that we assumed that the number of players is $k \geq 100 \log n$.) Among the important bits, roughly one of Alice's bits is equal to 1, and the goal is to discover which bit it is. Alice cannot know which coordinates are important, and in essence we are trying to solve a problem very similar to Set Disjointness. A lower bound of $\Omega(n)$ is easy to get convinced of, since it is similar to the known lower bounds for set disjointness. Proving it is not entirely trivial, and requires making some small modifications to Razborov's classical lower bound on set disjointness [30].

We see that the lower bound for this distribution has to (implicitly) rule out a Slepian-Wolf type upper bound. This provides some evidence that any lower bound for the k -OR problem would have to be non-trivial.

3.1.2 The Proof We prove the lower bound on k -OR by performing a reduction from the promise version of the two-party *set disjointness* problem (2-DISJ) which we lower-bounded in Lemma 2.2. Given an (x, y) for 2-DISJ drawn from the hard distribution μ , we construct an input for k -OR. Note that our mapping is not necessarily one-to-one, that is why we need a lower bound on the *expected* distributional communication complexity of 2-DISJ.

Reduction. We start with Alice's input set x and Bob's input set y from the distribution μ , with $t = 4$. That is, both x and y have $l = n/4$ 1 bits chosen at random under the condition that they intersect at one point with probability $1/4$, otherwise they intersect at no point. For the 2-DISJ problem, we construct k players' input sets I_1, \dots, I_k as follows. Let $z = [n] - y$. Let S_2^l, \dots, S_k^l be random subsets of size l from z , and $S_2^{l-1}, \dots, S_k^{l-1}$ be random subsets of size $l - 1$ from z .

Let S_2^1, \dots, S_k^1 be random elements from y .

$$\begin{cases} I_1 = x \\ I_j (j = 2, \dots, k) = \begin{cases} S_j^l & \text{w.p. } 1 - 1/4 \\ S_j^{l-1} \cup S_j^1 & \text{w.p. } 1/4 \end{cases} \end{cases}$$

Let μ' be this input distribution for k -OR. If I_j ($2 \leq j \leq k$) contains an element S_j^1 , then we call this element a special element.

This reduction can be interpreted as follows: Alice simulates a random player I_1 , and Bob, playing as the coordinator, simulates all the other $k - 1$ players I_2, \dots, I_k . Bob also keeps a set V containing all the special elements that ever appear in I_j ($j = 2, \dots, k$). It is easy to observe the following fact.

LEMMA 3.1. *All I_j ($j = 1, \dots, k$) are chosen from the same distribution.*

Proof. Since by definition I_j ($j = 2, \dots, k$) are chosen from the same distribution, we only need to show that $I_1 = x$ under μ is chosen from the same distribution as any I_j is under μ' . Given y , note that x is a random set of size l in $z = [n] - y$ with probability $1/4$; and with the remaining probability x is the union of a random set of size $l - 1$ in z along with a single random element from y . This is precisely the distribution of each I_j for $j \geq 2$. \square

The following lemma shows the properties of our reduction.

LEMMA 3.2. *If there exists a protocol \mathcal{P}' for k -OR on input distribution μ' with communication complexity C and error bound ε , then there exists a protocol \mathcal{P} for the 2-DISJ on input distribution μ with expected communication complexity $O(C/k)$ and error bound $\varepsilon + 4k/n$.*

Proof. Let us again view I_j ($j = 1, \dots, k$) as n -bit vectors. We show how to construct a protocol \mathcal{P} for 2-DISJ from a protocol \mathcal{P}' for k -OR with the desired communication cost and error bound. \mathcal{P} is constructed as follows: Alice and Bob first run \mathcal{P}' on I_1, \dots, I_k . Let $W \subseteq [n]$ be the set of indices where the results are 1. Bob checks whether there exists some $w \in W \cap y$ such that $w \notin V$. If yes, then \mathcal{P} returns "yes", otherwise \mathcal{P} returns "no".

We start by analyzing the communication cost of \mathcal{P} . Since player I_1 is chosen randomly from the k players, and Lemma 3.1 that all players' inputs are chosen from a same distribution, the expected amount of communication between I_1 (simulated by Alice) and the other $k - 1$ players (simulated by Bob) is at most a $2/k$ fraction of the total communication cost of \mathcal{P} .

Therefore the expected communication cost of \mathcal{P} is at most $O(C/k)$.

For the error bound, we have the following claim: With probability at least $(1 - 4k/n)$, there exists a $w \in W \cap y$ such that $w \notin V$ if and only if $x \cap y \neq \emptyset$. First, if $x \cap y = \emptyset$, then $I_1 = x$ cannot contain any element $w \in V \subseteq y$, thus the resulting bits in $W \cap y$ cannot contain any special element that is not in V . On the other hand, we have $\Pr[((W \cap y) \subseteq V) \wedge (x \cap y \neq \emptyset)] \leq 4k/n$. This is because $((W \cap y) \subseteq V)$ and $(x \cap y \neq \emptyset)$ hold simultaneously if and only if there exist some S_j^1 ($1 \leq j \leq k$) such that $S_j^1 \in x \cap y$. According to our random choices of S_j^1 ($j = 1, \dots, k$), this holds with probability at most $k/l \leq 4k/n$. Therefore if \mathcal{P}' errors at most ε , then \mathcal{P} errors at most $\varepsilon + 4k/n$. \square

Combining Lemma 2.2 and Lemma 3.2, we have the following theorem.

THEOREM 3.1. $D_{\mu'}^{1/800}(k\text{-OR}) = \Omega(nk)$, for $n \geq 3200k$ in the coordinator model.

Proof. If there exists a protocol \mathcal{P}' that computes k -OR on input distribution μ' with communication complexity $o(nk)$ and error bound $1/800$, then by Lemma 3.2 there exists a protocol \mathcal{P} that computes 2-DISJ on input distribution μ with expected communication complexity $o(n)$ and error bound $1/800 + 4k/n \leq 1/400$, contradicting Lemma 2.2 (when $t = 4$). \square

We discuss the applications of this lower bound to the distinct elements problem in Section 6.

3.2 Multiparty AND/OR with a Blackboard

Denote the k -OR problem in the blackboard model by k -OR-board. The general idea to prove a lower bound for k -OR-board is to perform a reduction from a 2-party bit-wise OR problem (2-OR for short) with public randomness. The 2-OR problem is the following. Alice and Bob each have an n -bit vector drawn from the following distribution: Each bit is 1 with probability $1/k$ and 0 with probability $1 - 1/k$. They also use public random bits to generate another $(k - 2)$ n -bit vectors such that each bit of these vectors is 1 with probability $1/k$ and 0 with probability $1 - 1/k$. That is, the bits are drawn from the same distribution as their private inputs. Let ν be this input distribution. Alice and Bob want to compute bit-wise OR of all these k n -bit vectors. For this problem we have the following theorem.

THEOREM 3.2. $\mathbf{E}[D_{\nu}^{1/3}(2\text{-OR})] = \Omega(n/k \cdot \log k)$.

Proof. W.l.o.g., let us assume that Bob outputs the final result of 2-OR. It is easy to see that if we take the bit-wise OR of the $(k - 2)$ n -bit vectors generated by public

random bits and Bob's input vector, the resulting n -bit vector b will have at least a constant density of 0 bits with probability at least $1 - o(1)$, by a Chernoff bound. Since Bob can see the $k - 2$ public vectors, to compute the final result, all that Bob has to know are bits of Alice's vector on those indices i ($1 \leq i \leq n$) where $b[i] = 0$. In other words, with probability at least $1 - o(1)$, Bob has to learn a specific, at least constant fraction of bits in Alice's vector up to error $1/3$. Plugging in Lemma 2.1 with $\rho = 1/k$, we know that the expected communication complexity is at least $(1 - o(1)) \cdot \Omega(n/k \cdot \log k) = \Omega(n/k \cdot \log k)$. \square

We reduce this problem to k -OR-board as follows: Alice and Bob each simulates a random player, and they use public random bits to simulate the remaining $k - 2$ players: The $(k - 2)$ n -bit vectors generated by their public random bits are used as inputs for the remaining $k - 2$ random players. Observe that the input of all the k players are drawn from the same distribution, thus the k players are symmetric. Consequently, the expected amount of communication between the two random players simulated by Alice and Bob and other players (including the communication between the two random players) is at most $O(1/k)$ fraction of the total communication cost of protocol for k -OR-board. We have the following theorem.

THEOREM 3.3. $D_{\nu}^{1/3}(k\text{-OR-board}) = \Omega(n \cdot \log k)$.

Proof. It is easy to see that if we have a protocol for the k -OR-board on input distribution ν with communication complexity $o(n \cdot \log k)$ and error bound $1/3$, then we have a protocol for 2-OR on input distribution ν with expected communication complexity $o(n/k \cdot \log k)$ and error bound $1/3$, contradicting Theorem 3.2. \square

It is easy to show a tight deterministic upper bound of $O(n \log k)$ for this problem in the blackboard model: each player speaks in turn and writes the coordinates where he has 1 and no player that spoke before him had 1.

3.3 Majority In the k -MAJ problem, we have k players, each having a bit vector of length n , and they want to compute bit-wise majority, i.e. to determine for each coordinate whether the majority of entries in this coordinate are 1 or 0. We prove a lower bound of $\Omega(nk)$ for this problem in the coordinator model by a reduction to 2-BITS via symmetrization.

For the reduction, we consider $k = 2t + 1$ players and describe the input distribution τ as follows. For each coordinate we assign it either t or $(t + 1)$, each with probability $1/2$, independently over all coordinates.

This indicates whether the k players contain t or $(t+1)$ 1 bits among them in this coordinate, and hence whether that index has a majority of 0 or 1, respectively. Then we place the either t or $(t+1)$ 1 bits randomly among the k players inputs in this coordinate. It follows that under τ : (i) each player's bits are drawn from the same distribution, (ii) each bit of each player is 1 with probability $1/2$ and 0 with probability $1/2$, and (iii) each index has probability $1/2$ of having a majority of 1s.

THEOREM 3.4. $D_\tau^{1/6}(k\text{-MAJ}) = \Omega(nk)$ in the coordinator model.

Proof. Now we use symmetry to reduce to the two-player problem 2-BITS. Alice will simulate a random player under τ and Bob simulate the other $k-1$ players. Notice that, by (ii) any subset of n' indices of Alice's are from $\zeta_{1/2}$. We will show that Alice and Bob need to solve 2-BITS on $\Omega(n)$ bits to solve k -MAJ. And, by (i), since all players have the same distribution and Alice is a random player, her expected cost in communicating to Bob is at most $O(C/k)$ if k -MAJ can be solved in C communication.

Now consider the aggregate number of 1 bits Bob has for each index; he has $(t-1)$ 1 bits with probability $1/4$, t 1 bits with probability $1/2$, and $(t+1)$ 1 bits with probability $1/4$. Thus for at most $(3/4)n$ indices (with probability at least $1 - \exp(-2(n/4)^2/n) = 1 - \exp(-n/8) \geq 1 - 1/7$) that have either $(t-1)$ or $(t+1)$ 1 bits Bob knows that these either will or will not have a majority of 1 bits, respectively. But for the other at least $n/4 = \Omega(n)$ remaining indices for which Bob has exactly t 1 bits, whether or not these indices have a majority of 1 bits depends on Alice's bit. And conditioned on this situation, each of Alice's relevant bits are 1 with probability $1/2$ and 0 with probability $1/2$, hence distributed by $\zeta_{1/2}$. Thus conditioned on at least $n/4$ undecided indices, this is precisely the 2-BITS problem between Alice and Bob of size $n/4$.

Thus a protocol for k -MAJ in $o(nk)$ communication and error bound $1/6$ would yield a protocol for 2-BITS in expected $o(n)$ communication and error bound $1/6 + 1/7 < 1/3$, by running the protocol simulated between Alice and Bob. This contradicts Lemma 2.1, and proves that k -MAJ requires $\Omega(kn)$ communication when allowing error on at most $1/6$ fraction of inputs. \square

Extensions.

This lower bound can easily be extended beyond just majority (threshold $1/2$) to any constant threshold ϕ ($0 < \phi < 1$), by assigning to each coordinate either $\lfloor k\phi \rfloor$ or $(\lfloor k\phi \rfloor + 1)$ 1 bits with probability $1/2$ each.

Let τ_ϕ denote this distribution. Then the analysis just uses ζ_ϕ in place of $\zeta_{1/2}$, which also yields an $\Omega(n)$ lower bound for 2-BITS. We call this extended k -MAJ problem (k, ϕ) -MAJ.

COROLLARY 3.1. $D_{\tau_\phi}^{1/6}((k, \phi)\text{-MAJ}) = \Omega(nk)$ for any constant ϕ ($0 < \phi < 1$) in the coordinator model.

We discuss the applications of this lower bound to the heavy-hitter problem in Section 6.

4 Graph Connectivity

In the k -CONN problem, we have k players, each having a set of edges in an n -vertex graph. The goal is to decide whether the graph consisting of the union of all of these edges is connected. In this section we prove an $\Omega(nk/\log^2 k)$ lower bound for k -CONN in the coordinator model, by performing a symmetry-based reduction from 2-DISJ.

4.1 Proof Idea and the Hard Distribution

Let us start by discussing the hard distribution. Assume for simplicity $k \geq 100 \log n$. We describe a hard distribution which is not quite the same as the one in the proof (due to technical reasons), but is conceptually clearer. In this hard distribution, we consider a graph G , which consists of two disjoint cliques of size $n/2$. We also consider one edge between these two cliques, called the *connector*; the connector is not part of G . Each player gets as input $n/10$ edges randomly and uniformly chosen from the graph G ; furthermore, with probability $1/2$ we choose exactly one random edge in one random player's input, and replace it by the connector. It is easy to see that if one of the players got the connector, then with high probability the resulting set of edges span a connected graph, and otherwise the graph is not connected.

To get convinced that the lower bound holds, notice that the coordinator can easily reconstruct the graph G . However, in order to find out if one of the players has received the connector, the coordinator needs to speak with each of the players to find this out. The situation is roughly analogous to the situation in the k -OR problem, since the players themselves did not get enough information to know G , and no Slepian-Wolf type protocol is possible since the edges received by each player are random-looking. The actual distribution that we use is somewhat more structured than this, in order to allow an easier reduction to the 2-player disjointness problem.

4.2 The Proof We first recall 2-DISJ. Similarly to before, in 2-DISJ, Alice and Bob have inputs x ($|x| = \ell$) and y ($|y| = \ell$) chosen uniformly at random from

$[n]$ ($n = 4\ell - 1$), with the promise that with probability $1/10k$, $|x \cap y| = 1$ and with probability $1 - 1/10k$, $|x \cap y| = 0$. Let φ be this input distribution for 2-DISJ. Now given an input (x, y) for 2-DISJ, we construct an input for k -CONN.

Let $K_{2n} = (V, E)$ be the complete graph with $2n$ vertices. Given Alice's input x and Bob's input y , we construct k players' input I_1, \dots, I_k such that $|I_j| = \ell$ and $I_j \subseteq E$ for all $1 \leq j \leq k$. We first pick a random permutation σ of $[2n]$. Alice constructs $I_1 = \{(\sigma(2i - 1), \sigma(2i)) \mid i \in x\}$.

Bob constructs I_2, \dots, I_k . It is convenient to use σ and y to divide V into two subsets L and R . For each i ($1 \leq i \leq n$), if $i \in y$, then with probability $1/2$, we add $\sigma(2i - 1)$ to L and $\sigma(2i)$ to R ; and with the rest of the probability, we add $\sigma(2i - 1)$ to R and $\sigma(2i)$ to L . Otherwise if $i \notin y$, then with probability $1/2$, we add both $\sigma(2i - 1)$ and $\sigma(2i)$ to L ; and with the rest of the probability, we add both $\sigma(2i - 1)$ and $\sigma(2i)$ to R . Let $K_L = (L, E_L)$ and $K_R = (R, E_R)$ be the two complete graphs on sets of vertices L and R , respectively. Now using E_R and E_L , Bob can construct each I_j . With probability $1 - 1/10k$ I_j is a random subset of disjoint edges (i.e. a matching) from $E_L \cup E_R$ of size ℓ ; and with probability $1/10k$, I_j is a random subset of disjoint edges from $E_L \cup E_R$ of size $\ell - 1$ and one random edge from $E \setminus (E_L \cup E_R)$.

Let φ' be the input distribution for k -CONN defined as above for each I_j ($1 \leq j \leq k$). We define the following two events.

ξ_1 : Both edge-induced subgraphs $\cup_{j=2}^k I_j \cap E_L$ and $\cup_{j=2}^k I_j \cap E_R$ are connected, and span L and R respectively.

ξ_2 : $\cup_{j=2}^k I_j \cap E$ is *not* connected.

It is easy to observe the following two facts by our construction.

LEMMA 4.1. *All I_j ($j = 1, \dots, k$) are chosen from the same distribution.*

Proof. Since σ is a random permutation of $[2n]$, according to the distribution φ' , Bob's input can also be seen as follows: With probability $1 - 1/10k$, it is a random matching of size ℓ from $E_L \cup E_R$; and with probability $1/10k$, it is a matching consists of $\ell - 1$ random edges from $E_L \cup E_R$ and one random edge from $E \setminus (E_L \cup E_R)$, which is $(\sigma(2z - 1), \sigma(2z))$ where $z = x \cap y$. \square

LEMMA 4.2. ξ_1 happens with probability at least $1 - 1/2n$ when $k \geq 68 \ln n + 1$.

Proof. (This is a proof sketch; full proof in Appendix C.) First note that by our construction, both $|L|$ and

$|R|$ are $\Omega(n)$ with high probability. To locally simplify notation, we consider a graph (V, E) of n nodes where edges are drawn in $(k - 1) \geq 68 \ln n$ rounds, and each round $n/4$ disjoint edges are added to the graph. If (V, E) is connected with probability $(1 - 1/4n)$, then by union bound over $\cup_{j=2}^k I_j \cap E_L$ and $\cup_{j=2}^k I_j \cap E_R$, ξ_1 is true with probability $(1 - 1/2n)$. The proof follows four steps.

- (S1) Using the first $28 \ln n$ rounds, we can show that all points have degree at least $8 \ln n$ with probability at least $1 - 1/12n$.
- (S2) Conditioned on (S1), any subset $S \subset V$ of $h < n/10$ points is connected to at least $\min\{h \ln n, n/10\}$ distinct points in $V \setminus S$, with probability at least $1 - 1/n^2$.
- (S3) Iterate (S2) $\ln n$ times to show that there must be a single connected component S_G of size at least $n/10$, with probability at least $1 - 1/12n$.
- (S4) Conditioned on (S3), using the last $40 \ln n$ rounds we can show that all points are connected to S_G with probability at least $1 - 1/12n$. \square

The following lemma shows the properties of our reduction.

LEMMA 4.3. *Assume $k \geq 100 \log n$. If there exists a protocol \mathcal{P}' for k -CONN on input distribution φ' with communication complexity C and error bound ε , then there exists a protocol \mathcal{P} for the 2-DISJ on input distribution φ with expected communication complexity $O(C/k \cdot \log k)$ and error bound $2\varepsilon + 1/2000k$.*

Proof. In \mathcal{P} , Alice and Bob first construct $\{I_1, \dots, I_k\}$ according to our reduction, and then run the protocol \mathcal{P}' on it. By Lemma 4.2 we have that ξ_1 holds with probability at least $1 - 1/2n$. And by our construction, conditioned on that ξ_1 holds, ξ_2 holds with probability at least $(1 - 1/10k)^{k-1} \geq 1 - 1/10$. Thus the input generated by a random reduction encodes the 2-DISJ problem with probability at least $(1 - 1/2n - 1/10) > 8/9$. We call such an input a *good* input. We repeat the random reduction $c \log k$ times (for some large enough constant c) and run \mathcal{P}' on each of the resulting inputs for k -CONN. The probability that we have at least one good input is at least $1 - 1/2000k$ (by picking c large enough). If this happens, Bob randomly pick an input among those good ones and the result given by running protocol \mathcal{P}' on that input gives the correct answer for 2-DISJ with probability at least $1 - 2\varepsilon$, by a Markov inequality, since by our reduction at least a $8/9$ fraction of inputs under φ' are good and \mathcal{P}' has error bound ε under φ' . Therefore we obtained a protocol \mathcal{P} for 2-DISJ

with expected communication complexity $O(C/k \cdot \log k)$ and error bound $2\varepsilon + 1/2000k$. \square

Combining Lemma 2.2 and Lemma 4.3, we have the following theorem.

THEOREM 4.1. $D_{\varphi'}^{1/4000k}(k\text{-CONN}) = \Omega(nk/\log k)$, for $k \geq 100 \log n$ in the coordinator model.

Proof. If there exists a protocol \mathcal{P}' that computes $k\text{-CONN}$ with communication complexity $o(nk/\log k)$ and error $1/4000k$, then by Lemma 3.2 there exists a protocol \mathcal{P} that computes 2-DISJ with expected communication complexity $o(n)$ and error at most $(2 \cdot 1/4000k + 1/2000k) = 1/1000k$, contradicting Lemma 2.2 (when $t = 10k$). \square

Finally, we have $R^{1/3}(k\text{-CONN}) \geq \Omega(R^{1/4000k}(k\text{-CONN})/\log k) \geq \Omega(D_{\varphi'}^{1/4000k}(k\text{-CONN})/\log k) \geq \Omega(nk/\log^2 k)$ as an immediate consequence.

The hardness of the graph connectivity problem may indicate that most graph problems in the coordinator model or message-passing model do not have efficient solutions other than trivial ones.

5 Some Direct-Sum-Like Results

Let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$ be an arbitrary function. Let μ be a probability distribution over $\mathcal{X} \times \mathcal{Y}$. Consider a setting where we have $k + 1$ players: Carol, and P_1, P_2, \dots, P_k . Carol receives an input from $x \in \mathcal{X}$ and each P_i receives an input $y_i \in \mathcal{Y}$. Let $R^\varepsilon(f^k)$ denote the randomized communication complexity of computing f on Carol's input and each of the k other players' inputs respectively; i.e., computing $f(x, y_1), f(x, y_2), \dots, f(x, y_k)$. Our direct-sum-like theorem in the message-passing model states the following.

THEOREM 5.1. *In the message-passing model, for any function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$ and any distribution μ on $\mathcal{X} \times \mathcal{Y}$, we have $R^\varepsilon(f^k) \geq \Omega(k \cdot \mathbf{E}[D_\mu^\varepsilon(f)])$.*

Note that this is not a direct-sum theorem in the strictest sense, since it relates randomized communication complexity to expected distributional complexity. However, it should probably be good enough for most applications. The proof of this theorem is dramatically simpler than most direct-sum proofs known in the literature (e.g. [6]). This is not entirely surprising, as it is weaker than those theorems: it deals with the case where the inputs are spread out over many players, while in the classical direct-sum setting, the inputs are only spread out over two players (this would be analogous to allowing the players P_i to communicate with

each other for free, and only charging them for speaking to Carol). However, perhaps more surprisingly, *optimal* direct-sum results are *not known* for most models, and are considered to be central open questions, while the result above is essentially optimal. Optimal direct-sum results in 2-player models would have dramatic consequences in complexity theory (see e.g. [21] as well as [28]), so it seems interesting to check whether direct-sum results in multiparty communication could suffice for achieving those complexity-theoretic implications.

Proof. Given the distribution μ on $\mathcal{X} \times \mathcal{Y}$, we construct a distribution ν on $\mathcal{X} \times \mathcal{Y}^k$. Let ρ_x be the marginal distribution on \mathcal{Y} induced by μ condition on $X = x$. We first pick $(x, y_1) \in \mathcal{X} \times \mathcal{Y}$ according to μ , and then pick y_2, \dots, y_k independently from \mathcal{Y} according to ρ_x . We show that $D_\nu^\varepsilon(f^k) \geq \Omega(k \cdot \mathbf{E}[D_\mu^\varepsilon(f)])$. The theorem follows by Yao's min-max principle.

Suppose that Alice and Bob get inputs (u, w) from $\mathcal{X} \times \mathcal{Y}$ according to μ . We can use a protocol for f^k to compute $f(u, w)$ as follows: Bob simulates a random player in $\{P_1, \dots, P_k\}$. W.l.o.g, say it is P_1 . Alice simulates Carol and the remaining $k - 1$ players. The inputs for Carol and P_1, \dots, P_k are constructed as follows: $x = u$, $y_1 = w$ and y_2, \dots, y_k are picked from \mathcal{Y} according to ρ_u (Alice knows u and μ so she can compute ρ_u). Let ν be the distribution of (x, y_1, \dots, y_k) in this construction. We now run the protocol for f^k on x, y_1, \dots, y_k . The result also gives $f(u, w)$.

Since y_1, \dots, y_k are chosen from the same distribution and P_1 is picked uniformly at random from the k players other than Carol, we have that in expectation, the expected amount of communication between P_1 and $\{\text{Carol}, P_2, \dots, P_k\}$, or equivalently, the communication between Alice and Bob according to our construction, is at most a $2/k$ fraction of the total communication of the $(k + 1)$ -player game. Thus $D_\nu^\varepsilon(f^k) \geq \Omega(k \cdot \mathbf{E}[D_\mu^\varepsilon(f)])$. \square

5.1 With Combining Functions In this section we extend Theorem 5.1 to the complexity of the AND/OR of k copies of 0/1 function f . Similar to before, AND and OR are essentially the same so we only talk about OR here. Let $R^\varepsilon(f_{\text{OR}}^k)$ denote the randomized communication complexity of computing $f(x, y_1) \vee f(x, y_2) \vee \dots \vee f(x, y_k)$.

THEOREM 5.2. *In the message-passing model, for any function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ and every distribution μ on $\mathcal{X} \times \mathcal{Y}$ such that $\mu(f^{-1}(1)) \leq 1/10k$, we have $R^{\varepsilon/4}(f_{\text{OR}}^k) \geq \Omega(k/\log(1/\varepsilon)) \cdot \mathbf{E}[D_\mu^\varepsilon(f)]$.*

Proof. The proof is similar as that for Theorem 5.1. The reduction is the same as that in the proof of Theorem 5.1. Let $D_\nu^{\varepsilon/4}(f_{\text{OR}}^k) = C$. Note that if $\mu(f^{-1}(1)) \leq$

$1/10k$, then with probability $(1 - 1/10k)^{k-1} \geq 0.9$, we have $f_{\text{OR}}^k(x, y_1, \dots, y_k) = f(u, w)$. We can boost this probability to $(1 - \varepsilon/2)$ by repeating the reduction for $O(\log 1/\varepsilon)$ times and then focus on a random good input (if any), as we did in the proof of Lemma 4.3. The k -player protocol on a random good input succeeds with probability at least $1 - 2 \cdot \varepsilon/4$, by a Markov inequality. Therefore we have a protocol for f with expected communication complexity $O(\log 1/\varepsilon \cdot C/k)$ and error bound $(2 \cdot \varepsilon/4 + \varepsilon/2) = \varepsilon$. \square

It would be interesting to generalize this to other combining functions such as majority, XOR and others.

One can also use symmetrization to prove similar direct-sum results for other settings. One such setting is when there are $2k$ players: k players that get x_1, \dots, x_k and k players that get y_1, \dots, y_k , and the goal is to compute $f(x_1, y_1), f(x_2, y_2), \dots, f(x_k, y_k)$. Another setting is the same except that there are only $k + 1$ players, and Carol receives all of the inputs x_1, \dots, x_k . We omit the details here. Versions of direct-sum in the blackboard model are also possible for some of these settings.

6 Applications

We now consider applications where multiparty communication complexity lower bounds such as ours are needed. As mentioned in the introduction, our multiparty communication problems are strongly motivated by research on the server-site model and the more general distributed streaming model. We discuss three problems here: the heavy hitters problem, which asks to find the approximately most frequently occurring elements in a set which is distributed among many clients; the distinct elements problems, which lists the distinct elements from a fixed domain where the elements are scattered across distributed databases possibly with multiplicity; and the ε -kernel problem, which asks to approximate the convex hull of a set which is distributed over many clients.

Distinct elements. Consider of domain of n possible elements and k distributed databases each which contains a subset of these elements. The exact distinct elements problem is to list the set of all distinct elements from the union of all elements across all distributed databases. This is precisely the k -OR problem and follows from Theorem 3.1, since the existence of each element in each distributed data point can be signified by a bit, and the bit-wise OR represents the set of distinct elements.

THEOREM 6.1. *For a set of k distributed databases, each containing a subset of n elements, it requires $\Omega(nk)$*

communication total between the databases to list the set of distinct elements with probability at least $2/3$.

ε -Kernels. Given a set of n points $P \subset \mathbb{R}^d$, the width in direction u is denoted by

$$\text{wid}(P, u) = \left(\max_{p \in P} \langle p, u \rangle \right) - \left(\min_{p \in P} \langle p, u \rangle \right),$$

where $\langle \cdot, \cdot \rangle$ is the standard inner product operation. Then an ε -kernel [2, 1] K is a subset of P so that for any direction u we have

$$\text{wid}(P, u) - \text{wid}(K, u) \leq \varepsilon \cdot \text{wid}(P, u).$$

An ε -kernel K approximates the convex hull of a point set P , such that if the convex hull of K is expanded in any direction by an ε -factor it contains P . As such, this coresot has proven useful in many applications in computational geometry such as approximating the diameter and smallest enclosing annulus of point sets [2, 1]. It has been shown that ε -kernels may require $\Omega(1/\varepsilon^{(d-1)/2})$ points (on a $(d-1)$ -sphere in \mathbb{R}^d) and can always be constructed of size $O(1/\varepsilon^{(d-1)/2})$ in time $O(n + 1/\varepsilon^{d-3/2})$ [35, 8].

We can note a couple of other properties about ε -kernels. Composibility: If K_1, \dots, K_k are ε -kernels of P_1, \dots, P_k , respectively, then $K = \bigcup_{i=1}^k K_i$ is an ε -kernel of $P = \bigcup_{i=1}^k P_i$. Transitivity: If K_1 is an ε_1 -kernel of P and K_2 is an ε_2 -kernel of K_1 , then K_2 is an $(\varepsilon_1 + \varepsilon_2)$ -kernel of P . Thus it is easy to see that each site i can simply send an $(\varepsilon/2)$ -kernel K_i of its data of size $n_\varepsilon = O(1/\varepsilon^{(d-1)/2})$ to the server, and the server can then create an $(\varepsilon/2)$ -kernel of $\bigcup_{i=1}^k K_i$ of size $O(1/\varepsilon^{(d-1)/2})$. This is asymptotically optimal size for an ε -kernel of the full distributed data set. We next show that this procedure is also asymptotically optimal in regards to communication.

THEOREM 6.2. *For a distributed set of k sites, it requires $\Omega(k/\varepsilon^{(d-1)/2})$ communication total between the sites and the server for the server to create an ε -kernel of the distributed data with probability at least $2/3$.*

Proof. We describe a construction which reduces k -OR to this problem, where each of k players has $n_\varepsilon = \Theta(1/\varepsilon^{(d-1)/2})$ bits of information. Theorem 3.1 shows that this requires $\Omega(n_\varepsilon k)$ communication.

We let each player have very similar data $P_i = \{p_{i,1}, \dots, p_{i,n_\varepsilon}\}$, each player's data points lie in a unit ball $B = \{q \in \mathbb{R}^d \mid \|q\| \leq 1\}$. For each player, their n_ε points are in similar position. Each player's j th point $p_{i,j}$ is along the same direction u_j , and its magnitude is either $\|p_{i,j}\| = 1$ or $\|p_{i,j}\| = 1 - 2\varepsilon$. Furthermore,

the set of directions $U = \{u_i\}$ are well-distributed such that for any player i , and any point $p_{i,j}$ that $P_i \setminus p_{i,j}$ is not an ε -kernel of P_i ; that is, the only ε -kernel is the full set. The existences of such a set follows from the known lower bound construction for size of an ε -kernel.

We now claim that the k -OR problem where each player has n_ε bits can be solved by solving the distributed ε -kernel problem under this construction. Consider any instance of k -OR, and translate to the ε -kernel problem as follows. Let the j th point $p_{i,j}$ of the i th player have norm $\|p_{i,j}\| = 1$ when j th bit of the player is 1, and have norm $\|p_{i,j}\| = 1 - 2\varepsilon$ if the j th bit is 0. By construction, an ε -kernel of the full set must acknowledge (and contain) the j th point from some player that has such a point with norm 1, if one exists. Thus the full ε -kernel encodes the solution to the k -OR problem: it must have n_ε points and, independently, the j th point has norm 1 if the j th OR bit is 1, and has norm $1 - 2\varepsilon$ if the j th OR bit is 0. \square

Heavy Hitters. Given a multi-set S that consists of n elements, a threshold parameter ϕ , and an error parameter ε , the *approximate heavy hitters* problem asks for a set of elements which contains all elements that occur at least ϕ times in S and contains no elements that occur fewer than $\phi(1 - \varepsilon)$ times in S . On a static non-distributed data set this can easily be done with sorting. This problem has been famously studied in the streaming literature where the Misra-Gries [27] and SpaceSaving [25] summaries can solve the problem in $O(1/\varepsilon)$ space. In the distributed setting the best known algorithms use random sampling of the indices and require either $O((1/\varepsilon^2)n \log n)$ or $O(k + \sqrt{kn}/\varepsilon \cdot \log n)$ communication to guarantee a correct set with constant probability [19]. We will prove a lower bound of $\Omega(n/\varepsilon)$ and leave tightening this bound as an open problem.

We now present a specific formulation of the approximate heavy hitters problem as (k, ϕ, ε) -HH as follows. Consider k players, each with a bit sequence (either 0 or 1) of length n where each coordinate represents an element. The goal is to answer YES for each index with at least ϕk elements, NO for each index with no more than $\phi k(1 - \varepsilon)$ elements, and either YES or NO for any count in between. The reduction is based on a distribution $\tau_{\phi, \varepsilon}$ where independently each index has either ϕk or $\phi k(1 - \varepsilon)$ 1 bits, each with probability $1/2$. In the reduction the players are grouped into sets of $k\varepsilon$ players each, and all grouped players for each index are either given a 1 bit or all players are given a 0 bit. These 1 bits are distributed randomly among the $1/\varepsilon$ groups. The proof then uses Corollary 3.1.

THEOREM 6.3. $D_{\tau_{\phi, \varepsilon}}^{1/6}((k, \phi, \varepsilon)\text{-HH}) = \Omega(n/\varepsilon)$.

Proof. To lowerbound the communication for (k, ϕ, ε) -HH we first show another problem is hard: $(1/\varepsilon, \phi)$ -HH (assume $1/\varepsilon$ is an integer). Here there are only $1/\varepsilon$ players, and each player at each index has a count of 0 or $k\varepsilon$, and we again want to distinguish between total counts of at least $k\phi$ (YES) and at most $k\phi(1 - \varepsilon)$ (NO). By distribution $\tau_{\phi, \varepsilon}$ each index has a total of either $k\phi$ or $k\phi(1 - \varepsilon)$ exactly. And then we distribute these bits to players so each player has precisely either 0 or $k\varepsilon$ at each index. When k is odd, this is precisely the $(1/\varepsilon, \phi)$ -MAJ problem, which by Corollary 3.1 takes $\Omega(n/\varepsilon)$ communication.

Now it is easy to see that $D_{\tau_{\phi, \varepsilon}}^{1/6}((1/\varepsilon, \phi)\text{-HH}) \leq D_{\tau_{\phi, \varepsilon}}^{1/6}((k, \phi, \varepsilon)\text{-HH})$, since the former on the same input allows $(1/\varepsilon)$ sets of $k\varepsilon$ players to talk to each other at no cost. \square

7 Conclusions and Open Problems

In this paper we have introduced the symmetrization technique, and have shown how to use it to prove lower bounds for k -player communication games. This technique seems widely applicable, and we expect future work to find further uses.

7.1 A Brief Comparison to the icost Method

In this section we make a brief comparison between our symmetrization method and the celebrated *icost* method by [5]. Readers who are familiar with the *icost* method may notice that the k -XOR, k -MAJ and blackboard k -OR/AND problems discussed in this paper can also be handled by the *icost* method. However, for problems whose complexities are different in the blackboard model and the message-passing model, e.g., k -OR and k -CONN, the *icost* method cannot be used to obtain tight lower bounds in the message-passing model, while the symmetrization method still applies.

If we view the input to k players as a matrix with players as rows each having an n -bit input, the *icost* method first “divides” the whole problem to n copies of primitive problems column-wise, and then analyzes a single primitive problem. While the symmetrization method first reduces the size of a problem in the row space, that is, it first reduce a k -player problem to a 2-player problem, and then analyze the 2-player problem. We can certainly use the *icost* method again when analyzing the resulting 2-player problem, which gives us an elegant way to combine these two techniques.

7.2 Limitations The symmetrization technique also has several limitations, which we wish to discuss here.

Firstly, there are problems that might be impossible to lower bound using symmetrization. Consider for

example the k -player disjointness problem, where each player gets a subset of $\{1, \dots, n\}$, and the goal is to decide whether there is an element that appears in all of the sets. This problem looks to be easier than the coordinate-wise AND problem. But we believe that this problem has a communication lower bound of $\Omega(nk)$ in the coordinator model as well. However, it seems impossible to prove this lower bound using symmetrization, for the following reason. Suppose we give Alice the input of a randomly-chosen player p_i , and give Bob the inputs of all the other players. It seems that for any symmetric input distribution, this problem can be solved using $O(n \log k/k)$ bits in expectation, which is much lower than the $\Omega(n)$ lower bound we are aiming for. It is not out of the question that some variant of the symmetrization technique can be used to get the $\Omega(nk)$ lower bound, but we do not know how to do it, and it might well be impossible. We leave this problem as an intriguing open question.

The second limitation is that symmetrization seems to require proving distributional lower bounds for 2-player problems, over somewhat-convoluted distributions. This presents some difficulty for the researcher, who needs to start proving lower bounds from scratch and cannot use the literature, since lower bounds in the literature are proved for other distributions. Yao’s min-max principle cannot be used, since it only guarantees that there is *some* hard distribution, but it doesn’t guarantee anything for the distribution of interest. This is often only a methodical difficulty, since it is often easy to get convinced that the distribution of interest is indeed hard, but a proof of this must still be found, which is a tedious task. It would be useful if there was some way to circumvent this difficulty, for example by finding a way that standard randomized lower bounds can be used.

The third limitation is that in order to use symmetrization, one needs to find a hard distribution for the k -player problem which is symmetric. This is usually impossible when the problem itself is not symmetric, i.e. when the players have different roles. For example, one could envision a problem where some of the players get as input elements of some group and the rest of the players get as input integers. However, note that for such problems, symmetrization can still be useful in a somewhat-generalized version. For example, suppose there are two sets of players: in set P , the players get group elements, and in set P' , the players get integers. Assume each of the sets contains exactly $k/2$ players. To use symmetrization, we would try to find a hard distribution that is symmetric inside of P and also inside of P' ; namely, a distribution where permuting the players inside P has no effect on the distribution, and simi-

larly for permuting the players inside P' . Then, to use symmetrization we can have Alice simulate two random players, p_i and p_j , where p_i is from P and p_j is from P' ; Bob will simulate all the rest of the players. Now symmetrization can be applied. If, alternatively, the set P contained just 3 players and P' contained $k - 3$ players, we can have Alice simulate one of the players in P' , Bob can simulate the rest of the players in P' , and either Alice or Bob can play the three players from P . As can be seen, with a suitable choice of distribution, it should still be possible to apply symmetrization to problems that exhibit some amount of symmetry.

The main topic for future work seems to be to find more settings and problems where symmetrization can prove useful. Since discovering this approach we have found repeatedly that problems that we encounter in other settings seem susceptible to this approach, and we think it has the potential to be a widely useful tool.

7.3 List of Problems Symmetrization is a technique that can be applied to a wide range of k -player communication problems, as well as having many possible applications in related models. Any problem which has a symmetric hard-looking distribution is a good candidate. Below is a list of candidate problems, some of which we lower-bound in this paper, and some that we leave for future work. This list is far from exhaustive.

- *Coordinate-wise problems:* Each player gets a vector of length n . Some symmetric coordinate-wise function $g : \{0, 1\}^k \rightarrow \{0, 1\}$ is applied, resulting in a length n vector. Then a “combining function” $h : \{0, 1\}^n \rightarrow Z$ is applied to the bits of the result. Z could be any output domain. In the examples we dealt with in this paper, g was AND, OR, XOR or MAJ, and h was typically the identity function. k -player disjointness is the case when g is AND and h is OR. We believe that for most choices of g and h that are non-constant symmetric functions, the complexity of this problem is $\Omega(nk)$. It might be interesting to conjecture this holds for any such h and g .
- *Equality:* Each player gets a vector of length n , and the goal is to decide whether all players have received the same vector. Another variant is to decide whether all players have received distinct vectors. This problem should have a lower bound of $\Omega(nk)$ in the coordinator model, with a proof using symmetrization; we defer the details to a later version.
- *Graph problems:* Each player gets a set of m edges in an n -vertex graph, and the goal is to compute

some graph property of the graph that consists of all of the edges in the input. In this paper we treat this problem when the graph property is connectivity. Obviously, many other properties could be interesting.

- *Pointer Chasing*: Each player has a set of directed edges in an n -vertex graph, where for each vertex v , only one player has an edge whose source is v . The “sink” of the instance is the vertex obtained by starting at vertex number 1 and walking according to the directed edges until reaching a vertex with no outgoing edges. We assume there is always a sink, and the goal is to find it.

Acknowledgements We would like to thank Andrew McGregor for helpful discussions, and for suggesting to study the connectivity problem. And David Woodruff for some insightful comments.

References

- [1] P. K. Agarwal, S. Har-Peled, and K. Varadarajan. Geometric approximations via coresets. *C. Trends Comb. and Comp. Geom. (E. Welzl)*, 2007.
- [2] P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Approximating extent measure of points. *J. ACM*, 51:606–635, 2004.
- [3] C. Arackaparambil, J. Brody, and A. Chakrabarti. Functional monitoring without monotonicity. In *ICALP*, 2009.
- [4] B. Babcock and C. Olston. Distributed top-k monitoring. In *SIGMOD*, 2003.
- [5] Z. Bar-Yossef, T. S. Jayram, R. Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. Comput. Syst. Sci.*, 68:702–732, June 2004.
- [6] B. Barak, M. Braverman, X. Chen, and A. Rao. How to compress interactive communication. In *STOC*, pages 67–76, 2010.
- [7] P. Cao and Z. Wang. Efficient top-k query calculation in distributed networks. In *PODC*, 2004.
- [8] T. Chan. Faster core-set constructions and data-stream algorithms in fixed dimensions. *Computational Geometry: Theory and Applications*, 35:20–35, 2006.
- [9] G. Cormode and M. Garofalakis. Sketching streams through the net: Distributed approximate query tracking. In *VLDB*, 2005.
- [10] G. Cormode, M. Garofalakis, S. Muthukrishnan, and R. Rastogi. Holistic aggregates in a networked world: Distributed tracking of approximate quantiles. In *SIGMOD*, 2005.
- [11] G. Cormode, S. Muthukrishnan, and K. Yi. Algorithms for distributed functional monitoring. In *SODA*, 2008.
- [12] G. Cormode, S. Muthukrishnan, K. Yi, and Q. Zhang. Optimal sampling from distributed streams. In *PODS*, 2010.
- [13] G. Cormode, S. Muthukrishnan, and W. Zhuang. What’s different: Distributed, continuous monitoring of duplicate-resilient aggregates on data streams. In *ICDE*, 2006.
- [14] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 1991.
- [15] P. Duris and J. D. P. Rolim. Lower bounds on the multiparty communication complexity. *J. Comput. Syst. Sci.*, 56(1):90–95, 1998.
- [16] A. Gál and P. Gopalan. Lower bounds on streaming algorithms for approximating the length of the longest increasing subsequence. In *FOCS*, 2007.
- [17] O. Goldreich. Secure multi-party computation. *Working Draft*, 2002. available at <http://www.wisdom.weizmann.ac.il/~oded/pp.html>.
- [18] S. Guha and Z. Huang. Revisiting the direct sum theorem and space lower bounds in random order streams. In *ICALP*, 2009.
- [19] Z. Huang, K. Yi, Y. Liu, and G. Chen. Optimal sampling algorithms for frequency estimation in distributed data. In *INFOCOM*, 2011.
- [20] B. Kalyanasundaram and G. Schintger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5:545–557, 1992.
- [21] M. Karchmer, R. Raz, and A. Wigderson. Super-logarithmic depth lower bounds via the direct sum in communication complexity. *Computational Complexity*, 5(3/4):191–204, 1995.
- [22] R. Keralapura, G. Cormode, and J. Ramamirtham. Communication-efficient distributed monitoring of thresholded counts. In *SIGMOD*, 2006.
- [23] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge, 1997.
- [24] A. Manjhi, V. Shkapenyuk, K. Dhamdhere, and C. Olston. Finding (recently) frequent items in distributed data streams. In *ICDE*, 2005.
- [25] A. Metwally, D. Agrawal, and A. E. Abbadi. An integrated efficient solution for computing frequent and top-k elements in data streams. *ACM Transactions on Database Systems*, 2006.
- [26] S. Michel, P. Triantafillou, and G. Weikum. Klee: a framework for distributed top-k query algorithms. In *VLDB*, 2005.
- [27] J. Misra and D. Gries. Finding repeated elements. *Sc. Comp. Prog.*, 2:143–152, 1982.
- [28] M. Patrascu. Towards polynomial lower bounds for dynamic problems. In *STOC*, pages 603–610, 2010.
- [29] B. Patt-Shamir and A. Shafir. Approximate distributed top-k queries. *Distributed Computing*, 21(1):1–22, 2008.
- [30] A. A. Razbarov. On the distributional complexity of disjointness. In *Theor. Comput. Sci.*, 106(2):385–390, 1990.
- [31] I. Sharfman, A. Schuster, and D. Keren. A geometric

approach to monitoring threshold functions over distributed data streams. In *SIGMOD*, 2006.

- [32] I. Sharfman, A. Schuster, and D. Keren. Shape sensitive geometric monitoring. In *PODS*, 2008.
- [33] A. C.-C. Yao. Probabilistic computations: Toward a unified measure of complexity (extended abstract). In *FOCS*, 1977.
- [34] K. Yi and Q. Zhang. Optimal tracking of distributed heavy hitters and quantiles. In *PODS*, 2009.
- [35] H. Yu, P. K. Agarwal, R. Poredy, and K. R. Varadara-jan. Practical methods for shape fitting and kinetic data structures using coresets. In *SoCG*, 2004.
- [36] Q. Zhao, M. Ogihara, H. Wang, and J. Xu. Finding global icebergs over distributed data sets. In *PODS*, 2006.

A Omitted Proof for 2-BITS

Lemma 2.1 (restated). $\mathbf{E}[D_{\zeta\rho}^{1/3}(2\text{-BITS})] = \Omega(n\rho \log(1/\rho))$.

Proof. Here we will make use of several simple tools from information theory. Given a random variable X drawn from a distribution μ , we can measure the amount of randomness in X by its entropy $H(X) = -\sum_x \mu(x) \log_2 \mu(x)$. The conditional entropy $H(X | Y) = H(XY) - H(Y)$ describes the amount of entropy in X , given that Y exists. The mutual information $I(X : Y) = H(X) + H(Y) - H(XY)$ measures the randomness in both random variables X and Y .

Let \mathcal{P} be any valid communication protocol. Let X be Alice's (random) input vector. Let Y be Bob's output as the Alice's vector he learns after the communication. Let Π be the transcript of \mathcal{P} . Let $\varepsilon = 1/3$ be the error bound allows by Bob. First, since after running \mathcal{P} , with probability at least $1 - \varepsilon$, we have $Y = X$, thus Bob knows X . Therefore

$$I(X : Y | \Pi) \leq \varepsilon H(X).$$

Consequently,

$$\begin{aligned} \varepsilon H(X) &\geq I(X : Y | \Pi) \\ &= H(X | \Pi) + H(Y | \Pi) - H(XY | \Pi) \\ &= H(X\Pi) - H(\Pi) + H(Y\Pi) - H(\Pi) \\ &\quad - (H(XY\Pi) - H(\Pi)) \\ &= H(X\Pi) + H(Y\Pi) - H(XY\Pi) - H(\Pi) \\ &= I(X\Pi : Y\Pi) - H(\Pi) \\ &\geq (1 - \varepsilon)H(X\Pi) - H(\Pi) \\ &\geq (1 - \varepsilon)H(X) - H(\Pi) \end{aligned}$$

Therefore, $\mathbf{E}[|\Pi|] \geq H(\Pi) \geq (1 - 2\varepsilon)H(X) \geq \Omega(nH(\rho)) \geq \Omega(n\rho \log(1/\rho))$. \square

B Omitted Proofs for the Biased 2-party Set Disjointness

Lemma 2.2 (restated). *When μ has $|x \cap y| = 1$ with probability $1/t$ then $\mathbf{E}[D_\mu^{1/100t}(2\text{-DISJ})] = \Omega(n)$.*

The proof is based on [30]. Before giving the proof, we first introduce some notations and a key technical lemma. Define

$$A = \{(x, y) : (\mu(x, y) > 0) \wedge (x \cap y = \emptyset)\}$$

and

$$B = \{(x, y) : (\mu(x, y) > 0) \wedge (x \cap y \neq \emptyset)\}.$$

Thus $\mu(A) = 1 - 1/t$ and $\mu(B) = 1/t$. We need the following key lemma, which is an easy extension of the main lemma in Razbarov [30] by rescaling the measures on the YES and NO instances.

LEMMA B.1. [30] *Let A, B, μ be defined as above. Let $R = C \times D$ be any rectangle in the communication protocol. Then we have $\mu(B \cap R) \geq 1/40t \cdot \mu(A \cap R) - 2^{-0.01n}$.*

Proof. (for Lemma 2.2) Let $\mathcal{R} = \{R_1, \dots, R_t\}$ be the minimal set of disjoint rectangles in which the protocol outputs "1", i.e., $x \cap y = \emptyset$. Imagine that we have a binary decision tree built on top of these rectangles. If we can show that there exists $\mathcal{O} \subseteq \mathcal{R}$ such that $\mu(\bigcup_{R_i \in \mathcal{O}} R_i) \geq 0.5 \cdot \mu(\bigcup_{R_i \in \mathcal{R}} R_i)$ and each of $R_i \in \mathcal{O}$ lies on a depth at least $0.005n$ in the binary decision tree, then we are done. Since $\mu(\bigcup_{R_i \in \mathcal{O}} R_i) \geq 0.5 \cdot \mu(\bigcup_{R_i \in \mathcal{R}} R_i) \geq 0.5 \cdot (\mu(A) - 1/100t) = \Omega(1)$ and querying inputs in each rectangle in \mathcal{O} costs $\Omega(n)$ bits.

We prove this by contradiction. Suppose that there exists $\mathcal{O}' \subseteq \mathcal{R}$ such that $\mu(\bigcup_{R_i \in \mathcal{O}'} R_i) > 0.5 \cdot \mu(\bigcup_{R_i \in \mathcal{R}} R_i)$ and each of $R_i \in \mathcal{O}'$ lies on a depth less than $0.005n$ in the binary decision tree. We have the following two facts.

1. There are at most $2^{0.005n}$ disjoint rectangles that lie on depths less than $0.005n$, i.e., $|\mathcal{O}'| \leq 2^{0.005n}$.
2. $\mu(\bigcup_{R_i \in \mathcal{O}'} (R_i \cap A)) > 0.5 - 1/100t$.

Combining the two facts with Lemma B.1 we reach the following contradiction of our error bound.

$$\begin{aligned} \mu\left(\bigcup_{i=1}^t (R_i \cap B)\right) &\geq \mu\left(\bigcup_{R_i \in \mathcal{O}'} (R_i \cap B)\right) \\ &\geq \sum_{R_i \in \mathcal{O}'} (1/40t \cdot \mu(R_i \cap A) - 2^{-0.01n}) \\ &> 1/40t \cdot (0.5 - 1/100t) - 2^{0.005n} \cdot 2^{-0.01n} \\ &> 1/100t \square \end{aligned}$$

C Omitted Proofs Graph Connectivity

We provide here a full proof for the probability of the event ξ_1 , that both subset of the graph L and R are connected.

Lemma 4.2 (restated). ξ_1 happens with probability at least $1 - 1/2n$ when $k \geq 68 \ln n + 1$.

Proof. First note that by our construction, both $|L|$ and $|R|$ are $\Omega(n)$ with high probability. To locally simplify notation, we consider a graph (V, E) of n nodes where edges are drawn in $(k - 1) \geq 68 \ln n$ rounds, and each round $n/4$ disjoint edges are added to the graph. If (V, E) is connected with probability $(1 - 1/4n)$, then by union bound over $\cup_{j=2}^k I_j \cap E_L$ and $\cup_{j=2}^k I_j \cap E_R$, ξ_1 is true with probability $(1 - 1/2n)$. The proof follows four steps.

(S1): *All points have degree at least $8 \ln n$.* Since for each of k rounds each point's degree increases by 1 with probability $1/2$, then the expected degree of each point is $4 \ln n$ after the first $28 \ln n$ rounds. A Chernoff-Hoeffding bound says that the probability that a point has degree less than $8 \ln n$ is at most $2 \exp(-2(6 \ln n)^2 / (14 \ln n)) \leq 2 \exp(-5 \ln n) \leq 2/n^5 \leq 1/12n^2$. Then by the union bound, this holds for none of the n points with probability at least $1 - 1/12n$.

(S2): *Conditioned on (S1), any subset $S \subset V$ of $h < n/10$ points is connected to at least $\min\{h \ln n, n/10\}$ distinct points in $V \setminus S$.* At least $9n/10$ points are outside of S , so each point in S expects to be connected at least $(9/10)8 \ln n \geq 7 \ln n$ times to a point outside of S . Each of these edges occur in different rounds, so they are independent. Thus we can apply a Chernoff-Hoeffding bound to say the probability that the number of edges outside of S for any point is less than $3 \ln n$ is at most $2 \exp(-2(4 \ln n)^2 / (8 \ln n)) = 2 \exp(-4 \ln n) = 2/n^4$. Thus the probability that no point in S has fewer than $\ln n$ edges outside S is (since $h < n/10$) at most $1/5n^3$.

If the $h \cdot 3 \ln n$ edges outside of S (for all h points) are drawn independently at random, then we need to bound the probability that these go to more than $n/10$ distinct points or $h \ln n$ distinct points. Since the edges are drawn to favor going to distinct points in each round, it is sufficient to analyze the case where all of the edges are independent, which can only increase the chance they collide. In either case $h \ln n < n/10$ or $h \ln n > n/10$ each time an edge is chosen (until $n/10$ vertices have been reached, in which case we can stop), $9/10$ of all

possible vertices are outside the set of edges already connected to. So if we select the $3h \ln n$ edges one at a time, each event connects to a distinct points with probability at least $9/10$, so we expect at least $(9/10)(3h \ln n) > 2h \ln n$ distinct points. Again by a Chernoff-Hoeffding bound, the probability that fewer than $h \ln n$ distinct points have been reached is at most $2 \exp(-2(h \ln n)^2 / (3hn)) \leq 2 \exp(-2/3)h \ln n < 2 \cdot \exp(-5 \ln n) \leq 2/n^5 \leq 1/5n^2$ (for $h \geq 8$). Together the probability of these events not happening is at most $1/2n^2$.

(S3): *There is a single connected component S_G of size at least $n/10$.* Start with any single point, we know from Step 1, its degree is at least $8 \ln n$. Then we can consider the set S formed by these $h_1 = 8 \ln n$ points, and apply Step 2 to find another $h_1 \ln n = 8 \ln^2 n = h_2$ points; add these points to S . The process iterates and at each round $h_i = 8 \ln^i n$, by growing only from h_i the newly added points. So, by round $i = \ln n$ the set $S = S_G$ has grown to at least size $n/10$. Taking the union bound over these $\ln n$ rounds shows that this process fails with probability at most $1/12n$.

(S4): *All points in $V \setminus S_G$ are connected to S_G .* Each round each point p is connected to S_G with probability at least $1/20$. So by coupon collector's bound, using the last $40 \ln n$ rounds all points are connected after $2 \ln n$ sets of 20 rounds with probability at least $1 - 1/12n$.

By union bound, the probability of steps (S1), (S3), and (S4) are successful is at least $1 - 1/4n$, proving our claim. \square