Optimal Two-Round Communication Lower Bound for Graph Connectivity via Pointer Chasing

Jaikumar Radhakrishnan
ICTS-TIFR, Bengaluru, India

Chaitanya Reddy
Department of Computer Science and Engineering, IIT Hyderabad, India

Rakesh Venkat
Department of Computer Science and Engineering, IIT Hyderabad, India

Abstract

We consider the communication complexity of the *graph connectivity problem,* where the edges of an n-vertex undirected graph G are distributed between two parties Alice and Bob, who are then required to communicate to determine if G is connected. We show that in any *randomized* protocol with *two-rounds* of communication, Alice and Bob must exchange $\Omega(n \log n)$ bits; such a lower bound for one-round protocols was shown by Sun and Woodruff (APPROX/RANDOM 2015). A one-round deterministic protocol, where Alice sends $O(n \log n)$ bits and Bob determines the answer, was observed by Hajnal, Maass and Turan (STOC 1988); they also showed a matching lower bound of $\Omega(n \log n)$ bits for *deterministic* protocols with unbounded rounds of communication. For randomized protocols, a reduction from the *set disjointness problem* due to Babai, Frankl and Simon (FOCS 1986) implies a randomized lower bound of $\Omega(n)$ even with unbounded rounds of communication. Whether this lower bound can be improved to $\Omega(n \log n)$ has been an outstanding open question, whose algorithmic implications were recently emphasized by Apers, Efron, Gawrychowski, Lee, Mukopadhyay and Nanongkai (FOCS 2022). Our lower bound for randomized two-round protocols is based on a reduction from a restricted version of the two-player *pointer chasing* problem originally studied by Papadimitriou and Sipser (JCSS 1984). Using this reduction, we show an $\omega(n)$ lower bounds on graph connectivity for any constant number of rounds by extending deterministic lower bounds shown by Ponzio, Radhakrishnan and Venkatesh (JCSS 2001) to the randomized setting.

2012 ACM Subject Classification Theory of computation → Communication complexity

Keywords and phrases Communication complexity

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.110

Funding *Jaikumar Radhakrishnan*: Department of Atomic Energy, Government of India, under project number RTI4001.

1 Introduction

Graph problems have been extensively studied in various computational settings – algorithms, decision trees, streaming algorithms, distributed algorithms and communication protocols. In this paper, we focus on the two-party communication setting: the edges of the n-vertex undirected graph $G = (V, E)$ are distributed between two parties Alice and Bob (some of the edges may be given to both), who then need to exchange messages in order to compute some function of G; the goal is to minimize the number of bits exchanged. Early works [4, 7] investigated properties such as connectivity, st-connectivity, bipartiteness, proving both lower and upper bounds on the amount of communication needed in this model. Other properties such as triangle-freeness, existence of Eulerian tours, k-edge connectivity, and cycle freeness among others have been studied and new lower bounds have been obtained in subsequent works [8, 19]. More recently, graph problems have been extensively studied for trade-offs between space requirements and the number of passes for streaming algorithms for estimating

quantities like matchings, cuts, etc., in various settings (see [2] and the references therein). Lower bounds in these settings are also typically shown by considering the communication complexity of some associated graph problem.

A fundamental problem in this domain is the *connectivity problem* Conn_n: Alice and Bob need to determine if G is connected. The randomized complexity of this problem was studied by Babai, Frankl and Simon [4], who showed a reduction from the *set disjointness* problem; when combined with the $\Omega(n)$ lower bound for the set disjointness problem [18, 9], this implies an $\Omega(n)$ lower bound for the randomized communication complexity of Conn_n. There exists an $O(n \log n)$-bit 1-round deterministic communication protocol: Alice sends a spanning forest of her edges to Bob, who then determines the answer. Hajnal, Maas and Turan [7] showed a *deterministic* lower bound of $\Omega(n \log n)$ for connectivity[1]. For randomized protocols with an unbounded number of rounds of communication, closing this gap between the upper and lower bounds has remained a major open question since the work of [4], and has been mentioned [1, 8] as one of the outstanding open questions in communication complexity. An optimal lower bound on Conn_n also implies the optimality of some query algorithms [1]. A randomized lower bound of $\Omega(n \log n)$ was shown for *one-round* protocols by [19]; however, no lower bound better than $\Omega(n)$ was known for randomized protocols with more than one round of communication[2].

We show an optimal randomized lower bound of $\Omega(n \log n)$ on Conn_n against protocols with *two rounds* of communication, by relating it to the *pointer chasing* problem introduced by Papdimitrou and Sipser [14] to study tradeoffs between the number of rounds of communication, and number of bits communicated. In the two-party pointer chasing problem[3], the edges of a *directed* bipartite graph $G = (\mathcal{V}_A \uplus \mathcal{V}_B, E)$ (where $|\mathcal{V}_A|, |\mathcal{V}_B| = n$) are distributed between Alice and Bob in the following way: Alice receives a one-to-one function $f_A : \mathcal{V}_A \to \mathcal{V}_B$ and Bob receives a one-to-one function $f_B : \mathcal{V}_B \to \mathcal{V}_A$, which we view as (directed) matchings E_A and E_B respectively. A fixed starting vertex $v_0 \in \mathcal{V}_A$ and a parameter k are known to both players. We consider two versions of the problem: the *full-version* and the *bit-version*. In the *full-version* of the problem, $\mathsf{PC}(n, k)$, the goal is to find the identity of the unique vertex v_k found by following the k-length path $v_0 \to v_1 := f_A(v_0) \to v_2 := f_B(v_1) \to \ldots \to v_k := f(v_{k-1})$, where $f = f_A$, if k is odd, and $f = f_B$, if k is even. In the *bit-version* of the problem, $\mathsf{bPC}(n, k)$, Alice and Bob need to determine just the *last bit* or the parity of v_k. At the end, *both* players must know the answer.[4]

There is a simple deterministic k-round protocol that exchanges $O(k \log n)$ bits to find v_k, where Alice starts by sending v_1, and in the i-th round the appropriate player sends v_i since they already know v_{i-1}. A major focus of previous work has been to show a separation between the amount of communication needed in the above case versus if only $k-1$ rounds of communication are allowed, or k rounds are allowed but *Bob* (the *wrong* player, since he does not know f_A) starts the communication. With this restriction, for a general k, a sequence of works [13, 20, 11] have resulted in a $\Omega(n/k)$ lower bound against randomized protocols. For

[1] In fact, they show that the bound holds in a stronger *arbitrary* partitioning model of the edges.

[2] Following convention, we only require one of the players to know the answer for Conn_n at the end of the protocol.

[3] Our definition differs from the standard definition, for we insist that the pointers given to Alice and Bob correspond to bijections – our reduction will make crucial use of this property. Note, however, that our lower bounds with this restriction on the inputs imply similar lower bounds for the problem without any restriction.

[4] The extra round that is needed for *both* players to know the answer in pointer chasing is often crucial.

constant k, Ponzio, Radhakrishnan and Venkatesh [15] show a $\Omega(n \log^{(k-1)} n)$ lower bound for the *full version*. As for upper bounds, Nisan and Wigderson [13] gave a randomized protocol with communication $O((k + n/k) \log n)$, and Damm, Jukna and Sgall [5] design a deterministic protocol with communication $O(n \log^{(k-1)} n)$, for constant k.

Our reduction from to Conn_n in fact passes through another related problem. In the Hamiltonian cycle problem, $\mathsf{HamCycle}(n)$, the inputs to Alice and Bob are the same as in $\mathsf{PC}(n, k)$, that is, matchings represented as one-to-one functions f_A and f_B. Alice and Bob need to determine if the union of f_A and f_B is a Hamiltonian Cycle, or it consists of more than one cycle. This particular problem was studied earlier by Raz and Spieker [17], who showed a lower bound of $\Omega(n \log \log n)$ for non-deterministic protocols even with unbounded rounds of communication. It is easy to see that a protocol for Conn_{2n} immediately yields a protocol with the same complexity for $\mathsf{HamCycle}(n)$.

1.1 Our Results

To state our results, we will need the following notation; detailed problem definitions and notation are presented in Section 2. Let $C_\delta^{B,r}(f)$ be the δ-error, r-round randomized communication complexity of the function f when Bob sends the first message; let $C_\delta^{A,r}(f)$ be the corresponding complexity when Alice sends the first message. We write $C_\delta^r(f)$, omitting the superscript, to denote the complexity when either party is allowed to start. $\log^{(i)} n$ denotes the i-th iterated logarithm of n.

Our lower bound for $C_\delta^2(\mathsf{Conn}_n)$ is derived as follows. We show that the pointer chasing problem reduces to the Hamiltonian cycle problem; more precisely, we show (see Theorem 7 in Section 3) that for all $\delta \in [0, 1)$, we have

$$C_\delta^2(\mathsf{HamCycle}(n)) \geq C_\delta^{B,3}(\mathsf{bPC}(16n + 2, 7)).$$

We then show that

$$C_\delta^{B,3}(\mathsf{bPC}(n, 7)) = \Omega(n \log n). \qquad \text{(Theorem 13 in Section 5.1)}$$

Theorem 13 is the main technical contribution of this work. Combining these with our earlier observation that $\mathsf{HamCycle}(n)$ reduces to Conn_{2n}, we obtain the following.

▶ **Theorem 1** (Main Result). *For all $\delta \in [0, 1/2)$, we have*

$$C_\delta^2(\mathsf{Conn}_n), C_\delta^2(\mathsf{HamCycle}(n)) = \Omega(n \log n).$$

Typically, pointer chasing has been considered in the setting the number of rounds of communication is $r = k - 1$; we know of only a few prior works [15, 3] that have considered $r < k - 1$. However, the techniques used in these do not seem to extend to yield an optimal $\Omega(n \log n)$ lower bound.

Furthermore, our reduction can be used to obtain superlinear lower bounds for Conn_n for randomized protocols with more than two rounds of communication. By extending the result [15, Theorem 6] to the randomized setting, we show that for all $\delta \in [0, 1/2)$, constant k and $r \leq k/2 - 1$

$$C_\delta^{B,r}(\mathsf{bPC}(n, k)) = \Omega(n \log^{(k-1)} n);$$

our reduction then implies that for all constant r,

$$C_\delta^r(\mathsf{Conn}_n) = \Omega\left(\frac{n}{r} \log^{(2r+1)} \frac{n}{r}\right).$$

The proof details of the above statement are deferred to the extended version of the paper.

For our main result we used the lower bound for the *bit version* of the pointer chasing problem, where Alice and Bob need to determine the last bit for v_7 with just three rounds of communication. We show that for the *full version* determining v_5 is already equally hard.

▶ **Theorem 2.** *For all constants* $\delta \in [0, 1/2)$, $C_\delta^{B,3}(\mathit{PC}(n,5)) = \Omega(n \log n)$.

Complementing the above, we show that there are *deterministic* protocols that can chase a larger pointer depth with few rounds, by extending the protocols in [5, 15].

▶ **Theorem 3.** *Let* $k' \geq \lceil \frac{k+5}{3} \rceil$. *Then* $C^{B,k'}(\mathit{PC}(n,k)) = O(nk \log\log n)$. *For the bit-version:* $k' \geq \lceil \frac{k+6}{4} \rceil$. *Then* $C^{B,k'}(\mathit{bPC}(n,k)) = O(nk \log\log n)$.

1.2 Organization of the paper

In Section 3, we describe the reduction from $\mathsf{PC}(n,k)$ to Conn_{2n} (and in particular, the $\mathsf{HamCycle}(n)$ problem).

Section 4 proves Theorem 2, a strong lower bound for the *full-version* of the pointer chasing problem. Section 5 builds upon the result in Section 4 to prove the main technical result: a strong lower bound for the *bit-version* Theorem 13. Proofs are detailed in Appendices D and E. Appendix F proves the upper bounds in Theorem 3.

Each technical result is preceded by a proof outline that emphasizes key steps and conceptual insights.

2 Preliminaries

2.1 Problem Definitions

▶ **Definition 4** (Pointer Chasing, full-version). *$\mathit{PC}(n,k)$ is the 2-player pointer chasing problem defined as follows: The input instance is a $(k+1)$-layered, directed graph G. The vertex set is $\mathcal{V}_0, \mathcal{V}_1, \ldots, \mathcal{V}_k$, with $|\mathcal{V}_i| = n$ for all $i \in \{0,1,2,\ldots,k\}$. The edges of G are bijections from $\mathcal{V}_i$ to $\mathcal{V}_{i+1}$ for each $i \in \{0,1,2,\ldots,k-1\}$. The edges are distributed between Alice and Bob as follows:*

Alice gets bijections $F_{A,i} : \mathcal{V}_{2i-2} \to \mathcal{V}_{2i-1}$ Bob gets bijections $F_{B,i} : \mathcal{V}_{2i-1} \to \mathcal{V}_{2i}$

There is a fixed vertex $v_0 \in \mathcal{V}_0$ known to both, and they need to find the unique vertex $v_k \in \mathcal{V}_k$ at distance k from v_0.

▶ **Definition 5** (Pointer Chasing, bit-version). *$\mathit{bPC}(n,k)$ is the problem of finding the* least significant bit *of v_k (as opposed to v_k in $\mathit{PC}(n,k)$). We call the least significant bit of v_k the* parity *of v_k (having identified the vertex set with $[n]$). This is also called the* bit-*version of the pointer-chasing problem.*

For pointer chasing problems, we require that any valid protocol has both players decide on the same answer and output it after the protocol concludes.

We note that Definition 4 differs slightly from the description given in the introduction: in our formal definition above, we allow a different function (bijection) at every step. It is not difficult to see that the two formulations are equivalent up to a factor-k blowup in the vertex-set size. In particular, any instance satisfying Definition 4 is a special case of the introduction's definition with $|\mathcal{V}_A| = |\mathcal{V}_B| = n\lceil (k+1)/2 \rceil$, obtained by taking $\mathcal{V}_A = \bigcup_{i:i \bmod 2=0} \mathcal{V}_i$ and $\mathcal{V}_B = \bigcup_{i:i \bmod 2=1} \mathcal{V}_i$.

▶ **Definition 6.** *The* HamCycle(n) *problem is defined as follows. Let* V_A *and* V_B *be two disjoint sets of* n *vertices each. Alice receives a bijection map* $\sigma_A : V_A \to V_B$*, and Bob receives a bijection* $\sigma_B : V_B \to V_A$*. The goal is to decide whether the composition* $\sigma_A \circ \sigma_B$ *(equivalently, union of the two perfect matchings) forms a Hamiltonian cycle covering all* $2n$ *vertices.*

For the problems Conn_n, HamCycle(n) under consideration, following convention, we require only that one of the players knows the answer after r rounds.

2.2 Notation

We use $[n]$ to denote the set $\{1, 2, \ldots, n\}$. Random variables will typically be denoted by capital letters: e.g. V_5 denotes a random variable, and $V_5 = v_5$ will be its instantiation to the value v_5. Most random variables we deal with will be over a finite, discrete domain. The support of a random variable X is denoted by $\mathrm{supp}(X)$.
All logarithms are to the base 2, and $\log^{(i)} n$ denotes the i-th iterated logarithm of n.
Our proofs make significant use of information theoretic concepts. We describe some preliminaries in Section A.
We assume throughout the proofs that when used, ε, δ are sufficiently small constants.

2.3 Communication Complexity preliminaries

We assume some familiarity with basic notions of two-party communication complexity, and refer the reader to standard books [10, 16] for the basics. Let $f : \mathcal{X} \times \mathcal{Y} \to Z$ be a function.

Rectangles, Subrectangles. A combinatorial rectangle is a subset of inputs of the form $A \times B$, where $A \subseteq \mathcal{X}$ and $B \subseteq \mathcal{Y}$. A deterministic protocol $\Pi = (M_1, M_2, \ldots, M_k)$ splits the input space into a disjoint partition of rectangles, each rectangle corresponding to a certain transcript $(M_1 = m_1, M_2 = m_2, \ldots, M_k = m_k)$ of the protocol.

We call a rectangle R' a subrectangle of a rectangle R if $R' \subseteq R$. R' is a subrectangle of the protocol Π, if it is a subrectangle of some rectangle in the partition induced by the protocol Π. If $R = A \times B$ is a rectangle, then a subrectangle of the form $A' \times B$ for $A' \subseteq A$ is called a horizontal slice of R. Similarly, $A \times B'$, for $B' \subseteq B$ is called a vertical slice.

Let $R = A \times B$ be a rectangle. Let Z be a random variable that depends on the inputs. We denote $Z|R$ to be the random variable having distribution of Z conditioned on inputs restricted to the rectangle R.

3 Lower bounds on HamCycle(n)

▶ **Theorem 7.** *Let* n *be a multiple of* 4*. The* bPC(n, k) *problem reduces to the* HamCycle(N) *problem with* $N = 2 + 2n(k + 1)$ *in the two-party communication setting.*

Proof. For notational simplicity, we describe the reduction with only two functions $\sigma_A : [n] \to [n]$, and $\sigma_B : [n] \to [n]$ in the pointer chasing instance. It will be clear that the reduction can be modified to accommodate the case of different bijections at every level. We specify the reduction for even k, a similar one works for odd k with minor changes.

Let G be the graph for the pointer chasing instance, with permutations $\sigma_A : V_A \to V_B$ and $\sigma_B : V_B \to V_A$ as edges, and a fixed starting vertex $v_0 \in V_A$. We identify V_A and V_B with $[n]$. For notational reasons, we assume that v_0 is the vertex 2 in V_A. Let v_k be the pointer at depth k. We show that Alice and Bob can produce a graph G' without communication, such that G' is a single cycle if and only if $v_k \bmod 2 = 1$.

Though the graph G' described will be a directed graph, the underlying undirected graph of G' (undirected version of G') will be a Hamiltonian cycle iff $v_k \mod 2 = 1$

We first describe a reduction that produces a graph $\tilde{G}'$ satisfying the following: if $v_k \bmod 2 = 1$, then $\tilde{G}'$ is the union of *two* cycles. If $v_k \bmod 2 = 0$, then $\tilde{G}'$ is the union of *four* cycles. A slight modification to the edges in this instance will produce an appropriate YES/NO instance of $\mathsf{HamCycle}(N)$, concluding the reduction.

Vertex set for $\tilde{G}'$. We define symmetric vertex sets V'_A and V'_B forming the bipartite graph $\tilde{G}' = (V'_A \cup V'_B, \tilde{E}')$:

$$V'_A = \{A^\alpha_0, A^\beta_0\} \cup \{A^\alpha_{t,\ell}, A^\beta_{t,\ell} : t \in [k+1], \ell \in [n]\},$$
$$V'_B = \{B^\alpha_0, B^\beta_0\} \cup \{B^\alpha_{t,\ell}, B^\beta_{t,\ell} : t \in [k+1], \ell \in [n]\}$$

We have $(k+1)$ *blocks* of vertices. Each vertex v in the original G has *two* copies in each block: an α-copy and a β-copy (denoted by superscripts). There are 4 additional special vertices labeled with block number 0 for technical reasons. $A^\alpha_{t,\ell}$ denotes the α-copy of the vertex $\ell \in V_A$ in the t-th block.

We note again that v_0 was the vertex labeled 2 in V_A in G.

Hence, $|V'_A| = |V'_B| = 2 + 2n(k+1)$, and the total number of vertices in $\tilde{G}'$ (and G') is $|V'_A| + |V'_B| = 4 + 4n(k+1)$.

Alice's Edges (based on σ_A).

- Fixed edges: $A^\alpha_0 \to B^\alpha_0, \quad A^\beta_0 \to B^\beta_0$
- Permutation ($t \in [1, k+1], \ell \in [n]$):

$$A^\alpha_{t,\ell} \to \begin{cases} B^\alpha_{t,\sigma_A(\ell)}, & \text{if } t \le \frac{k}{2}, \\ B^\alpha_{t,\sigma_A^{-1}(\ell)}, & \text{if } t > \frac{k}{2}+1 \end{cases}, \quad A^\beta_{t,\ell} \to \begin{cases} B^\beta_{t,\sigma_A(\ell)}, & \text{if } t \le \frac{k}{2}, \\ B^\beta_{t,\sigma_A^{-1}(\ell)}, & \text{if } t > \frac{k}{2}+1 \end{cases}$$

- Critical ($t = \frac{k}{2}+1$): $\quad A^\alpha_{t,\ell} \to \begin{cases} B^\alpha_{t,\ell}, & \text{if } \ell \bmod 2 = 0, \\ B^\beta_{t,\ell}, & \text{if } \ell \bmod 2 = 1 \end{cases}, \quad A^\beta_{t,\ell} \to \begin{cases} B^\beta_{t,\ell}, & \text{if } \ell \bmod 2 = 0, \\ B^\alpha_{t,\ell}, & \text{if } \ell \bmod 2 = 1 \end{cases}$

Bob's Edges (based on σ_B).

- Fixed: $B^\alpha_0 \to A^\alpha_{1,v_0}, \quad B^\beta_0 \to A^\beta_{1,v_0}$
- Permutation ($t \in [1, k], \ell \in [n]$):

$$B^\alpha_{t,\ell} \to \begin{cases} A^\alpha_{t+1,\sigma_B(\ell)}, & \text{if } t \le \frac{k}{2}, \\ A^\alpha_{t+1,\sigma_B^{-1}(\ell)}, & \text{if } t > \frac{k}{2} \end{cases}, \quad B^\beta_{t,\ell} \to \begin{cases} A^\beta_{t+1,\sigma_B(\ell)}, & \text{if } t \le \frac{k}{2}, \\ A^\beta_{t+1,\sigma_B^{-1}(\ell)}, & \text{if } t > \frac{k}{2} \end{cases}$$

- Block k+1 to Alice's first block: $\quad (v_0 \equiv 2)$

$$B^\alpha_{k+1,\ell} \to \begin{cases} A^\alpha_{1,\ell+1}, & \text{if } \ell \notin \{v_0{-}1, v_0, n\}, \\ A^\alpha_{1,v_0+1}, & \text{if } \ell = v_0{-}1, \\ A^\alpha_0, & \text{if } \ell = v_0, \\ A^\alpha_{1,1}, & \text{if } \ell = n \end{cases}, \quad B^\beta_{k+1,\ell} \to \begin{cases} A^\beta_{1,\ell+1}, & \text{if } \ell \notin \{v_0{-}1, v_0, n\}, \\ A^\beta_{1,v_0+1}, & \text{if } \ell = v_0{-}1, \\ A^\beta_0, & \text{if } \ell = v_0, \\ A^\beta_{1,1}, & \text{if } \ell = n \end{cases}$$

▷ Claim 8. If $v_k \bmod 2 = 0$, $\tilde{G}'$ is the union of 4 cycles.
If $v_k \bmod 2 = 1$, $\tilde{G}'$ is the union of 2 cycles.

Proof of Claim. We defer the analysis of the cases to Appendix C. ◁

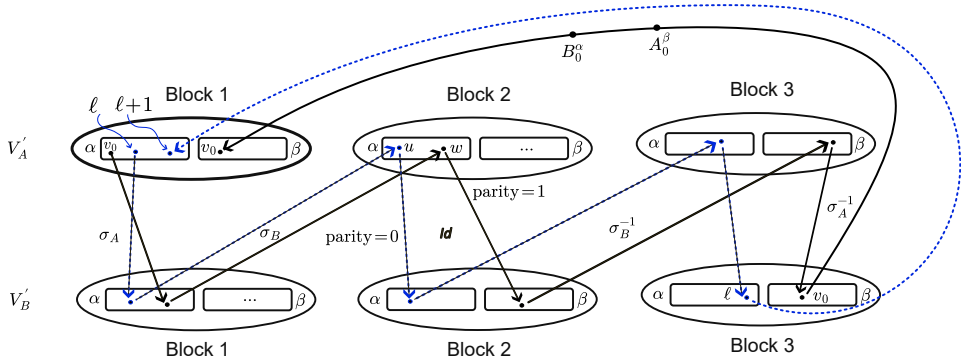

Figure 1 An illustration of typical paths for $k = 2$. The solid path is from v_0, and the dashed path from an arbitrary $\ell \neq v_0$. In the above case, both α and β copies of v_0 will lie in the same cycle.

Modification. The following slight modification that replaces two constant edges in $\tilde{G}'$ gives us the desired G':
(Old edges): $B^{\alpha}_{k+1,n} \to A^{\alpha}_{1,1}$ and $B^{\alpha}_{0} \to A^{\alpha}_{1,v_0}$, we now map
(New edges): $B^{\alpha}_{k+1,n} \to A^{\alpha}_{1,v_0}$ and $B^{\alpha}_{0} \to A^{\alpha}_{1,1}$.

These edges are present in different cycles in the Yes instances of $\tilde{G}'$. Splicing merges the two to a single Hamiltonian cycle (and keeps the NO instances disconnected, since at most two edges are modified). ◀

Assuming our main technical result (Theorem 13), the proof of Theorem 1 follows immediately from the above reduction. Note that since $\mathsf{HamCycle}(N)$ is symmetric between Alice and Bob, we can assume without loss of generality that Bob initiates the communication in any protocol solving it.

4 Randomized lower bound for $\mathsf{PC}(n, 5)$

▶ **Theorem 9** (Theorem 2 restated). *For any constant $\delta \in [0, 1/2)$:*

$$C^{(B,3)}_{\delta}(\mathit{PC}(n, 5)) = \Omega(n \log n)$$

Proof outline. For context, we first describe how pointer chasing lower bounds are typically shown via round elimination [13, 20, 15]. For odd k, consider a k-round protocol $\Pi = (M_1, \ldots, M_k)$ with low distributional error, where Bob sends M_1. Fixing typical values of $M_1 = m_1$ and $V_1 = v_1$, one shows that $F_B(v_1)$ remains highly variable if $|\Pi|$ is small. Iterating this $k - 1$-times, alternating roles between Alice and Bob each time, we can fix $M_1, \ldots, M_{k-1}$ and $V_1, \ldots, V_{k-1}$ while preserving variability in $F_A(v_{k-1})$. Yet Bob must now infer v_k with high confidence, contradicting the protocol's correctness.

The strategy of fixing v_i (i.e., looking at a subrectangle with fixed v_i) after round i that helps isolate entropy in V_{i+1} fails when $\Omega(n \log n)$ bits of communication are allowed. For our lower bound for $\mathsf{PC}(n, 5)$, we avoid fixing v_2, instead allowing it to vary and analyzing a *good assignment* $\sigma : \mathcal{A}_2 \to \mathcal{V}_B$, where $\mathcal{A}_2 \subseteq \mathrm{supp}(V_2)$. Under high-entropy $F_B \mid m_1$, we show that V_4 remains variable over $\Omega(n)$ candidates, each with probability $\Omega(1/n)$. Since $F_A \mid m_2$ also retains entropy (since $|m_2|$ is small), Bob cannot reliably infer the next pointer across these V_4's. Using an unfixed V_2 is what allows us to infer that V_4 is spread out. However, this

requires a careful probabilistic analysis in choosing a good σ while accounting for correlations across layers. The proof relies on the permutation structure of the functions. We now present the proof.

Proof. We show that no 3-round deterministic protocol $\Pi = (M_1, M_2, M_3)$, with Bob starting the communication and using at most $\varepsilon n \log n$ bits exchanged per message for a small constant $\varepsilon > 0$, can correctly solve the problem on more than a (small) constant fraction of inputs drawn from the uniform distribution. The theorem then follows by an application of Yao's lemma [10, Sec 3.4, Theorem 3.20], followed by standard error reduction techniques via repetition.

Recall notation from Section 2: Alice has functions $f_{A,1}$ on $\mathcal{V}_0 \to \mathcal{V}_1$, $f_{A,2}$ on $\mathcal{V}_2 \to \mathcal{V}_3$, and $f_{A,3}$ on $\mathcal{V}_4 \to \mathcal{V}_5$. Bob has $f_{B,1}$ and $f_{B,2}$ for $\mathcal{V}_1 \to \mathcal{V}_2$ and $\mathcal{V}_3 \to \mathcal{V}_4$ respectively. We use $\mathcal{F}_A$ and $\mathcal{F}_B$ to denote the set of all possible inputs to Alice and Bob respectively.

Let Bob send M_1, with $|M_1| = \varepsilon_B n \log n$, and Alice send M_2, with $|M_2| = \varepsilon_A n \log n$. After M_2, Bob should know the answer, which he conveys to Alice, using M_3. We therefore consider just the first two messages, and show that Bob cannot know the correct answer for a constant fraction of the inputs.

Initially, with no messages passed, we have:

- $H[F_{A,1}] = H[F_{A,2}] = H[F_{A,3}] = H[F_{B,1}] = H[F_{B,2}] = \log n! = n \log n - O(n)$
- $H[V_i] = \log n$, for all $1 \le i \le 5$

We prove the theorem in a sequence of three claims, whose proofs we defer to Appendix D.

▷ Claim 10 (Round 1). With probability $\ge 1/2$ over the choice of inputs, (M_1, V_1) satisfy the following:

1. $H[F_{B,1}(v_1) | M_1 = m_1, V_1 = v_1] \ge (1 - 12\varepsilon_B) \log n$
2. $H[F_{A,2}] \ge n \log n - O(n)$
3. $H[F_{B,2} \mid M_1 = m_1] \ge (1 - 4\varepsilon_B) n \log n$

Proof. Deferred to Appendix D. ◁

We note that fixing such a pair of (m_1, v_1) gives us a subrectangle of the protocol after m_1. For any such subrectangle, we show the following:

▷ Claim 11 (Main Claim). For a subrectangle $R \subseteq \mathcal{F}_A \times \mathcal{F}_B$ formed by fixing (m_1, v_1) satisfying the conditions in Claim 10, the following is true: There exists a set $\mathcal{A}_4 := \{b_1, b_2, \ldots, b_r\} \subseteq \mathcal{V}_4$, for $r \ge n/6$, and a constant c_1 such that with probability $\ge 1 - n \exp(-c_1 n^{1-O(\varepsilon_B)}) = 1 - o(1)$ over choice of $F_{A,2} \sim R$:

$$\forall i \in [r] : \qquad \Pr_{f_B \sim R} [V_4 = b_i \mid f_{A,2}] \ge \frac{1}{20n} \tag{1}$$

Proof. Deferred to Appendix D. ◁

Let $\sigma := F_{A,2}$ restricted to the domain $\mathrm{supp}(V_2)$. As mentioned in the outline, the proof of Claim 11 goes via analyzing a random assignment $\sigma : \mathrm{supp}(V_2) \to \mathcal{V}_3$ under $F_{A,2}$. We proceed to analyze what happens in the subrectangle of a fixed σ that satisfies (1) from Claim 11, after Round 2 (i.e. after Alice sends M_2). We note that once we fix a σ, it corresponds to a class of F_A's, and hence a subrectangle within the rectangle R in Claim 11.

▷ Claim 12 (Error after Round 2). Consider a subrectangle $R' \subseteq R \subseteq \mathcal{F}_A \times \mathcal{F}_B$ caused by fixing a $\sigma = \sigma^*$ satisfying (1) in Claim 11. In any such rectangle R', Bob makes an error in knowing the answer with probability at least $1/400$, after receiving Alice's message M_2.

Proof. Deferred to Appendix D. ◁

We finally collate the errors across the subrectangles from Claims 10, 11 and 12. The overall error of the protocol is at least:

$$\Pr(\text{Error }) \geq \frac{1}{2} \cdot \left(1 - n \exp\left(-c_1 n^{1-O(\varepsilon_B)}\right)\right) \cdot \frac{1}{400} = \Omega(1)$$ ◀

5 Randomized lower bound for bPC$(n, 7)$

5.1 Proof of main result

▶ **Theorem 13.** *For any constant* $\delta \in [0, 1/2)$,

$$C_\delta^{(B,3)}(\mathit{bPC}(n, 7)) = \Omega(n \log n)$$

Proof Outline. Before giving the proof, we give a rough outline of the steps involved.

We start with a sub-rectangle R' of the protocol as described in Claim 12, where m_1, v_1, σ are fixed. We have $\Omega(n)$ possibilities for V_4 at this point, called $b_1, \ldots, b_r$. From Bob's point of view, within the rectangle R', he exactly knows V_4; suppose that $V_4 = b_1$.

On receiving m_2, Bob will know the *shadow* of b_1: the set $\Delta(b_1) \coloneqq \{f_{A,3}(b_1) : f_{A,3} \in R'\}$, which typically has size $\approx n^{1-\varepsilon_A}$. For the full-version, this was sufficient to prevent Bob from knowing V_5, but it is not clear that the *parity* of V_7 reachable from potential V_4's will vary within the shadow. Indeed, if the goal was to find parity of V_5, Alice could *explicitly send* this information to Bob in $O(n)$ bits with m_2 (by sending parity$(F_{A,3}(i))$ for every $i \in [n]$). The extra layers in the instance come into play here: when looking at the parity of V_7, even if Alice sends the last bit (or more) of $f_{A,\star}(i)$ to Bob, due to the high entropy of intermediate layers, we show that the *composed function* parity$(F_{A,4}(f_{B,3}(F_{A,3}(b_j))))$ is still unpredictable to Bob within the shadow $\Delta(b_j)$.

Following [15], we view the parity of a vertex in the set $\mathcal{V}_7$ as a *two-coloring* $\chi_7 : \mathcal{V}_7 \to \{0, 1\}$. For a given set of inputs, f_A, f_B, this induces two-colorings on higher levels: $\chi_6 \coloneqq \chi_7 \circ f_{A,4}$, $\chi_5 \coloneqq \chi_6(f_{B,3})$. If the intermediate functions are viewed as random, the resulting coloring at each layer is also a random coloring.

In terms of two-colorings, it is sufficient to show that within R', $\Delta(b_1) \subseteq \mathcal{V}_5$ is two-colored by χ_5 in a roughly balanced way while varying over the inputs consistent with the current subrectangle. We formally show this in the proof via a *composition lemma*. This is a major technical contribution of our work, and we elaborate on it in Appendix E.

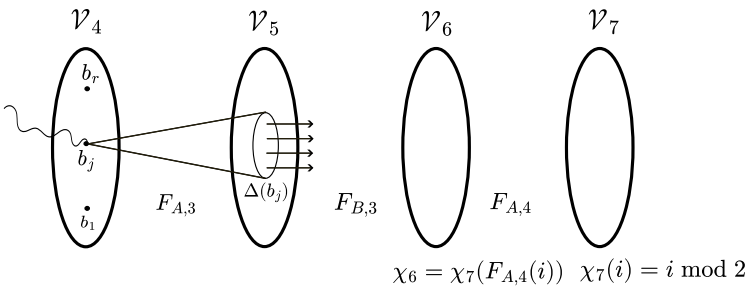

Figure 2 Fixing (m_1, v_1, σ), Bob knows V_4 exactly, say $V_4 = b_j$. $\Delta(b_j)$ will be two-colored by $\chi_6 \circ F_{B,3}|b_j$ by most $F_{B,3}|b_j$ (by composition lemma), preventing Bob from knowing the answer.

Proof. We assume, as before, that all input functions come from the uniform distribution over respective domains and show that any protocol communicating $\varepsilon n \log n$ bits will err on a (small) constant fraction of the inputs (for small enough ε).

We follow the proof of Theorem 9 and arrive at a rectangle R' as in Claim 11. This gives a set $\mathcal{A}_4 = \{b_1, b_2, \ldots, b_r\} \subseteq \mathcal{V}_4$, of size $r \geq n/6$ for some fixed m_1, v_1, σ.

Compared to the full-version, we have two extra layers of functions: $F_{B,3}, F_{A,4}$. The fixing of m_1, v_1, σ does not affect $F_{A,4}$, and the effect on $F_{B,3}$ is described by $H[F_{B,3} \mid R'] \geq (1 - 4\varepsilon_B)n \log n$ from Claim 10.

Due to correlations induced by the message m_1, the possibilities for function $f_{B,3}$ between $\mathcal{V}_5, \mathcal{V}_6$ are influenced by what V_4 is. In particular, let $F_{B,3}^{m_1,b_j} := F_{B,3} | R', V_4 = b_j$. We first show that *many* of the $b_j \in \mathcal{A}_4$ ensure a high entropy of $F_{B,3}^{m_1,b_j}$.

▷ Claim 14. There exists $\mathcal{A}'_4 \subseteq \mathcal{A}_4$ and constants $c_1 > 1, \alpha_1 \in (0,1)$ satisfying:
1. $\forall b_j \in \mathcal{A}'_4 \quad : \quad H[F_{B,3} | R', V_4 = b_j] \geq (1 - c_1\varepsilon_B)n \log n$
2. $|\mathcal{A}'_4| \geq \alpha_1 n$

Proof. Deferred to Appendix E. ◁

We now invoke the following Composition Lemma, whose proof is given in the extended version of this paper.

▶ **Lemma 15** (Composition Lemma)**.** *Let $F : [n] \to [n]$ be a (random) function satisfying $H[F] \geq (1 - \varepsilon)n \log n$. Let $\chi : [n] \to \{0,1\}$ be a random, balanced two-coloring of $[n]$. Then, $\exists J \subseteq [n]$ with $|J| \geq (1 - \sqrt{\varepsilon})n$ such that for any distribution D on J , with $H_\infty[D] \geq k$, with probability $\geq 1 - 1/n^3$ over choice of χ, the composed function $X_i := \chi \circ F(i)$ satisfies:*

$$\Pr_F \left(\left| \sum_i D(i) X_i - \frac{1}{2} \right| \geq 10\varepsilon^{1/32} + t \right) \leq 4\varepsilon^{1/32} + \frac{1}{\varepsilon^{1/32}} \cdot \exp(-\Omega(t^2 \cdot 2^k))$$

Proof. Deferred to the extended version of this paper. ◀

Invoke the lemma for each $b_i \in \mathcal{A}'_4$ in turn, setting $F \equiv F_{B,3}^{m_1,b_i}$ (and $\varepsilon = c_1\varepsilon_B$, where c_1 is the constant from Claim 14), and random choice of $\chi \equiv \chi_6 : \mathcal{V}_6 \to \{0,1\}$. Call a χ_6 *nice*, if it ensures the result of Lemma 15 for *every* $b_i \in \mathcal{A}'_4$. By a union bound, a random χ_6 is nice with probability $\geq 1 - \frac{1}{n}$.

Fix a nice χ_6, this forms a subrectangle R'' of R', being a restriction of $F_{A,4}$. For each such nice χ_6, we get sets $J_1, \ldots, J_{r'} \subseteq \mathcal{V}_5$, with $|J_i| \geq (1 - \sqrt{c_1\varepsilon_B})n$, for every $i \in [r']$. It now remains to show that a typical b_i has a large shadow *within* J_i after m_2. The following claim shows that this happens with good probability. The proof is deferred to Appendix E.

▷ Claim 16. Let $\gamma := \sqrt{c_1\varepsilon_B}$. Let $\mathcal{A}'_4 := \{b_1, \ldots, b_r\} \subseteq \mathcal{V}_4$ with $r \geq \alpha_1 n$, as in Claim 14. For each $i \in [r]$, let $J_i \subseteq \mathcal{V}_5$ be sets with $|J_i| \geq (1 - \gamma)n$. Define the events $\mathcal{E}_i := \{F_{A,3}(b_i) \in J_i\}$, for each $i \in [r]$.

Then, there exist a partitioning of R'' into subrectangles $\{R''_{(m_2,z)}\}_{m_2 \in \mathrm{supp}(M_2), z \in \{0,1\}^r}$ (where m_2 is Alice's message and z is a function of $F_{A,3}$ alone) and constant $\alpha_2 > 0$ such that with probability at least $1/8$ over choice of $M_2 = m_2$ and z, we have: at least $\alpha_2 n$ elements $b_i^* \in \mathcal{A}$ satisfy:

$$H[F_{A,3}(b_i^*) \mid M_2 = m_2, Z = z, \mathcal{E}_i] \quad \geq \quad (1 - 4\sqrt{\varepsilon_A}) \log n,$$

Choose a (m_2, z) to fix a sub-rectangle $\tilde{R} := R''_{(m_2,z)} \subseteq R''$ (a horizontal slice in R'') that satisfies the conclusion of Claim 16. (χ_6, m_2, z) are fixed here. Let $S^* := \{b^*_{i_1}, b^*_{i_2} \ldots, b^*_{i_\ell}\}$ be the indices in $\mathcal{A}'_4$ arising out of Claim 16. We have $\Pr(S^*) \geq \alpha_2 n \times 1/20n = \alpha_2/20$.

Let D'_j be the distribution $F_{A,3}(b^*_{i_j})|\tilde{R}$; we have $H[D'_j] \geq (1 - 4\sqrt{\varepsilon_A}) \log n$. Set δ' to be $\varepsilon_A^{1/32}$; by Lemma 17 (Appendix B), D'_j is within δ' in ℓ_1 distance of a distribution D_j such that $H_\infty[D_j] \geq (1 - 4\sqrt{\varepsilon_A}/\delta') \log n$.
Fix a good $b_j \in \mathcal{A}'_4$ (a vertical slice in $\tilde{R}$).
Set $t = 10\varepsilon^{1/32}$, $\delta := 4\sqrt{\varepsilon_A}/\delta'$ in Lemma 15. We conclude that $\chi_6 \circ F^{b_j}_{B,3}|\tilde{R}$ two-colors D_j with bias $\eta \leq 20\varepsilon^{1/32}$ with probability (over $F^{b_j}_{B,3}|\tilde{R}$) of at least $\beta := (1 - 5\varepsilon^{1/32})$. (Recall that $\varepsilon = c_1\varepsilon_B$.) This means it colors D'_j with bias at most $\eta_2 := \eta + \delta' = \eta + \varepsilon_A^{1/32}$

Call an $f_{B,3}$ good if it ensures a bias of at most η_2 in R''. On such $f_{B,3}$, a deterministic 0 or 1 answer from Bob makes an error with probability at least $1/2 - \eta_2$.

The overall probability of error, is, therefore:

$$\Pr[\text{Error}] \geq \left(\frac{1}{2} - \eta_2\right) \times \frac{\alpha_2}{20} \times \frac{1}{8} \times \left(1 - \frac{1}{n}\right) \times \beta = \Omega(1)$$

This concludes the proof of Theorem 13. ◀

References

1. Simon Apers, Yuval Efron, Paweł Gawrychowski, Troy Lee, Sagnik Mukhopadhyay, and Danupon Nanongkai. Cut query algorithms with star contraction. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 507–518. IEEE, 2022.
2. Sepehr Assadi. Recent advances in multi-pass graph streaming lower bounds. *ACM SIGACT News*, 54(3):48–75, 2023. doi:10.1145/3623800.3623808.
3. Sepehr Assadi, Gillat Kol, Raghuvansh R Saxena, and Huacheng Yu. Multi-pass graph streaming lower bounds for cycle counting, max-cut, matching size, and other problems. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 354–364. IEEE, 2020. doi:10.1109/FOCS46700.2020.00041.
4. László Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory. In *27th Annual Symposium on Foundations of Computer Science (sfcs 1986)*, pages 337–347. IEEE, 1986.
5. Carsten Damm, Stasys Jukna, and Jiří Sgall. Some bounds on multiparty communication complexity of pointer jumping. *Computational Complexity*, 7(2):109–127, 1998. doi:10.1007/PL00001595.
6. Devdatt P Dubhashi and Alessandro Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
7. András Hajnal, Wolfgang Maass, and György Turán. On the communication complexity of graph properties. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 186–191, 1988. doi:10.1145/62212.62228.
8. Gábor Ivanyos, Hartmut Klauck, Troy Lee, Miklos Santha, and Ronald de Wolf. New bounds on the classical and quantum communication complexity of some graph properties. In *32nd International Conference on Foundations of Software Technology and Theoretical Computer Science*, page 148, 2012.
9. Bala Kalyanasundaram and Georg Schintger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545–557, 1992. doi:10.1137/0405044.
10. Eyal Kushilevitz. Communication complexity. In *Advances in Computers*, volume 44, pages 331–360. Elsevier, 1997. doi:10.1016/S0065-2458(08)60342-3.

11 Xinyu Mao, Guangxu Yang, and Jiapeng Zhang. Gadgetless lifting beats round elimination: Improved lower bounds for pointer chasing. In Raghu Meka, editor, *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, volume 325 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 75:1–75:14, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2025.75.

12 Colin McDiarmid. On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188, 1989.

13 Noam Nisan and Avi Widgerson. Rounds in communication complexity revisited. In *Proceedings of the twenty-third annual ACM symposium on Theory of computing*, pages 419–429, 1991. doi:10.1145/103418.103463.

14 Christos H Papadimitriou and Michael Sipser. Communication complexity. In *Proceedings of the fourteenth annual ACM symposium on Theory of computing*, pages 196–200, 1982. doi:10.1145/800070.802192.

15 Stephen J Ponzio, Jaikumar Radhakrishnan, and Srinivasan Venkatesh. The communication complexity of pointer chasing. *Journal of Computer and System Sciences*, 62(2):323–355, 2001. doi:10.1006/JCSS.2000.1731.

16 Anup Rao and Amir Yehudayoff. *Communication complexity: and applications*. Cambridge University Press, 2020.

17 Ran Raz and Boris Spieker. On the "log rank"-conjecture in communication complexity. *Combinatorica*, 15(4):567–588, 1995.

18 Alexander A Razborov. On the distributional complexity of disjointness. In *International Colloquium on Automata, Languages, and Programming*, pages 249–253. Springer, 1990. doi:10.1007/BFB0032036.

19 Xiaoming Sun and David P Woodruff. Tight bounds for graph problems in insertion streams. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2015)*, pages 435–448. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.APPROX-RANDOM.2015.435.

20 Amir Yehudayoff. Pointer chasing via triangular discrimination. *Combinatorics, Probability and Computing*, 29(4):485–494, 2020. doi:10.1017/S0963548320000085.

A Information Theory Preliminaries

Let X be a discrete random variable over a domain $\mathcal{X}$. The entropy of X is defined by $H[X] \coloneqq -\sum_{x\in\mathrm{supp}(X)} \Pr(X=x)\log\Pr(X=x)$ It is always true that $H[X] \le \log|\mathrm{supp}(X)|$. If σ is a random permutation, then $H[\sigma] = \log n! = n\log n - O(n)$.

The *conditional entropy* of a random variable Y given X is defined by

$$H[Y|X] \coloneqq \sum_{x\in\mathrm{supp}(X)} \Pr[X=x]H[Y|X=x] = \mathbb{E}_x[H[Y_x]]$$

Here, Y_x denotes the random variable with distribution $\Pr[Y_x=y] = \Pr[Y=y|X=x]$

The min-entropy of a random variable X is defined as: $H_\infty[X] = \max_{x\in\mathrm{supp}(X)} \log\frac{1}{\Pr[X=x]}$.

Properties of Entropy.

1. Non-negativity: $H(X) \ge 0, \quad H_\infty(X) \ge 0$.
2. Bounds: $H_\infty(X) \le H(X) \le \log|\mathcal{X}|$.
3. Chain rule: $H(X,Y) = H(X) + H(Y \mid X)$.
4. Conditioning: $H(Y \mid X) \le H(Y)$.
5. Data processing: If $Y = f(X)$ then $H(Y) \le H(X)$ and $H_\infty(Y) \le H_\infty(X)$.
6. Subadditivity: $H(X,Y) \le H(X) + H(Y)$.

▶ **Fact 16.0.1** (Fano's Inequality). *Let X be a random variable taking values in a finite set $\mathcal{X}$, and let $\hat{X}$ be an estimate of X based on some observation Y. Define the probability of error as*

$$P_e = \Pr[\hat{X} \neq X].$$

Then,

$$H(X|Y) \leq h(P_e) + P_e \log(|\mathcal{X}| - 1),$$

where $H(X|Y)$ is the conditional entropy of X given Y, and $h(p) = -p\log p - (1-p)\log(1-p)$ is the binary entropy function.

B Statements and proofs of lemmas used

▶ **Lemma 17.** *Let μ be a distribution over $[n]$, and $X \sim \mu$. If $H[X] \geq \log n - A$ for some $A \geq 0$, Then for any $\delta \in (0,1)$, there exist distributions μ_1, μ_2 over $[n]$ such that*

$$\mu = \delta \cdot \mu_1 + (1-\delta) \cdot \mu_2,$$

where

$$H_\infty(\mu_2) \geq \log n - \frac{A+1}{\delta}$$

Proof. Set

$$\theta := \frac{2^{(A+1)/\delta}}{n}, \qquad B := \{i \in [n] : \mu(i) > \theta\}.$$

We claim $\mu(B) \leq \delta$. Suppose, for the sake of contradiction, that $\mu(B) > \delta$. Let $\mathbb{1}_{[B]}$ denote the indicator random variable of the event that $X \in B$. We have:

$$\begin{aligned}
H[X] &\leq H[X, \mathbb{1}_{[B]}] = H[\mathbb{1}_{[B]}] + H[X|\mathbb{1}_{[B]}] \\
&\leq 1 + \mu(B)\Big(\sum_{i\in[B]} \frac{\mu(i)}{\mu(B)} \log \frac{\mu(B)}{\mu(i)}\Big) + \big(1-\mu(B)\big)\log n \\
&\leq 1 + \mu(B)\Big(\log n - \tfrac{A+1}{\delta}\Big) + \big(1-\mu(B)\big)\log n + \mu(B)\log\mu(B) \\
&= \log n + 1 - \mu(B)\left(\frac{A+1}{\delta}\right) + \mu(B)\log\mu(B) \\
&< \log n - A, \qquad \ldots \text{if } \mu(B) > \delta
\end{aligned}$$

contradicting the assumption $H[X] \geq \log n - A$.

Let $\mu(B) := \delta' < \delta$. Let $\mu_1 := \mu_{|B}$, and $\mu_2 := \mu_{|B^c}$. Then, clearly, $\mu = \delta'\mu_1 + (1-\delta')\mu_2$. Since $\delta > \delta'$, we can write:

$$\begin{aligned}
\mu &= (1-\delta)\mu_2 + (\delta - \delta')\mu_2 + \delta'\mu_1 \\
&= (1-\delta)\mu_2 + \delta\mu_1'
\end{aligned}$$

Where μ_1' is the normalized version of $(\delta - \delta')\mu_2 + (\delta')\mu_1$. ◀

▶ **Lemma 18** (Entropy vs. Min-Entropy [15]). *Let D be a distribution on $[n]$ with entropy at least $\log n - A$. Then, for all $\varepsilon > 0$, D is within ε of a distribution with min-entropy at least*

$$\log n - \frac{A+1}{\varepsilon}.$$

▶ **Theorem 19** ((The method of averaged bounded differences),[12], [6, Corollary 5.1]). *Let $X_1, \ldots, X_n$ be random variables and f a function such that for each $i \in [n]$ there exists $c_i \geq 0$ with*

$$\left|\mathbb{E}[f \mid X_1, \ldots, X_{i-1}, X_i = a_i] - \mathbb{E}[f \mid X_1, \ldots, X_{i-1}, X_i = a_i']\right| \leq c_i$$

for all choices of a_i, a_i'. Then, for any $t \geq 0$,

$$\Pr(|f - \mathbb{E}[f]| \geq t) \leq 2\exp\left(-\frac{2t^2}{\sum_{i=1}^{n} c_i^2}\right).$$

C Missing proofs from Section 3

Proof of Claim 8. We begin with the following observation:

▶ **Observation 20.** *For a vertex $v' \in V'_A$, and $j \in \mathbb{N}$, let $w_j(v')$ denote the unique vertex at distance j from v'. Then, for all $i \in [n]$,*

$$w_{2k+1}(A^{\alpha}_{1,i}) = \begin{cases} B^{\beta}_{k+1,i}, & \text{if } w_k(A^{\alpha}_{1,i}) \bmod 2 = 1, \\ B^{\alpha}_{k+1,i}, & \text{if } w_k(A^{\alpha}_{1,i}) \bmod 2 = 0 \end{cases},$$

$$w_{2k+1}(A^{\beta}_{1,i}) = \begin{cases} B^{\alpha}_{k+1,i}, & \text{if } w_k(A^{\beta}_{1,i}) \bmod 2 = 1, \\ B^{\beta}_{k+1,i}, & \text{if } w_k(A^{\beta}_{1,i}) \bmod 2 = 0 \end{cases}$$

Given Obs 20, we analyze each case:

- **YES Case: (When $v_k \bmod 2 = 1$)**
  Chasing the pointer starting at the α-copy vertex A^{α}_{1,v_0}, if $v_k \bmod 2 = 1$, then after the $k/2$-th layer, it switches to a path on the β block vertices, ending up at B^{β}_{k+1,v_0}. This switch triggers a path back such that A^{α}_{1,v_0} and A^{β}_{1,v_0} end up on the same cycle, call it C_1. Now consider following the outgoing edges sequentially (chasing the pointer) starting any of the remaining vertices $A^{\alpha}_{1,\star}$. Since $n/2$ is even and all functions are permutations, an odd number of them will switch from α to a copy vertex β in the layer $k/2$ when following outgoing edges. The edges are set up for the vertex $A^{\star}_{1,\ell}$ to lead back to $A^{\star}_{1,\ell+1}$, so we eventually visit all the vertices.
  Due to this, the remaining edges will form a single second cycle C_2. In this cycle, the vertices of $A^{\star}_{1,\star}$ appear in the order $A^{\star}_{1,1} \rightsquigarrow A^{\star}_{1,3} \rightsquigarrow A^{\star}_{1,4} \ldots$.
- **NO Case: (When $v_k \bmod 2 = 0$)**
  In this case, A^{α}_{1,v_0} and A^{β}_{1,v_0} end up in 2 separate cycles since they do not switch from α to β or vice versa.
  By a similar argument as above, chasing pointers from the remaining vertices $A^{\alpha}_{1,\star}$, an even number of them switch from α to a β copy at layer $k/2$ when following pointers. This implies that $A^{\alpha}_{1,\ell}$ and $A^{\beta}_{1,\ell}$ cannot fall in the same cycle. ◁

D Missing proofs from Section 4

Proof of Claim 10

Proof of Claim 10. Since $|M_1| \leq \varepsilon_B n \log n$, $H[F_B \mid M_1] \geq (2-\varepsilon_B)\, n \log n$. [5]

By Markov's inequality, we have

$$\Pr_{m_1}\left(H[F_B \mid M_1 = m_1] \geq (2-4\varepsilon_B) n \log n\right) \geq \frac{3}{4},$$

so for at least $\frac{3}{4}$ of all f_B, the first message $M_1 = m_1$ satisfies

$$H[F_B^{m_1}] \coloneqq H[F_B \mid M_1 = m_1] \geq (2-4\varepsilon_B) n \log n.$$

For such any m_1, we therefore have: $H[F_{B,j} \mid M_1 = m_1] \geq (1-4\varepsilon_B) n \log n$, for $j = 1, 2$.

Since $H[F_{B,1}^{m_1}] \geq (1-4\varepsilon_B) n \log n$, subadditivity of entropy implies that $\mathbb{E}_{i\in[n]}[H[F_{B,1}|M_1 = m_1]] \geq (1-4\varepsilon_B)\log n$. Applying Markov's inequality yields that at least $2n/3$ of $i \in [n]$ satisfy $H[F_{B,1}^{m_1}(i)] \geq (1-12\varepsilon_B)\log n$.

Call a $v_1 \in \mathcal{V}_1$ *good* for m_1, if $H[F_{B,1}^{m_1}(i)] \geq (1-12\varepsilon_B)\log n$. Since the distribution of V_1 is uniform, rectangles corresponding to pairs (m_1, v_1) with $M_1 = m_1$ and good v_1 cover at least a fraction $\frac{3}{4} \times \frac{2}{3} = \frac{1}{2}$ of the total inputs. ◁

Proof of Claim 11

Proof of Claim 11. Define $\mathcal{A}_2 \coloneqq \mathrm{supp}(V_2) = \{a_1, \ldots, a_\ell\}$. Let $F_{B,2}^{m_1,a_j}$ denote a random variable following the distribution of $F_{B,2}$ conditioned on $M_1 = m_1, V_2 = a_j$.

Let μ_B be the distribution induced over $\mathcal{A}_2$ (equivalently, V_2) in R. We know from Claim 10 that $H[\mu_B] = H[V_2 \mid M_1 = m_1, V_1 = v_1] \geq (1-12\varepsilon_B)\log n$.

Since $H[F_{B,2}|M_1 = m_1] \geq (1-4\varepsilon_B) n \log n$, we have that $H[F_{B,2}|M_1 = m_1, V_2] \geq (1-4\varepsilon_B) n \log n - \log n \geq (1-4.5\varepsilon_B) n \log n$.

By Markov's inequality, with probability $\geq 1/2$ over choices for V_2, it holds that $H[F_{B,2}^{m_1,a_j}] \geq (1-9\varepsilon_B) n \log n$.

Call $\mathcal{A}_2' \coloneqq \{a_j \in \mathcal{A}_2 : H[F_{B,2}^{m_1,a_j}] \geq (1-9\varepsilon_B) n \log n\}$; the above implies $\mu_B(\mathcal{A}_2') \geq 1/2$.

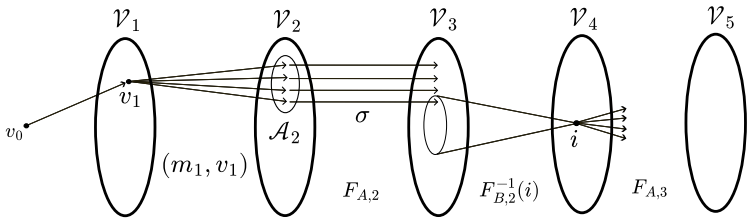

Figure 3 Most σ's confer good probability mass on i, if $F_{B,2}^{-1}(i)$ between $\mathcal{V}_4$, $\mathcal{V}_3$ has high entropy.

▶ **Definition 21.**

- *Call a vertex $i \in V_4$ is* locally good *with respect to $a_j \in \mathcal{A}_2'$ if*

$$H[F_{B,2}^{-1}(i) \mid M_1 = m_1, V_2 = a_j] \geq (1-36\varepsilon_B)\log n$$

- *$i \in V_4$ is* globally good *if it is locally good for at least $\frac{1}{4}$ of $\mathcal{A}_2$ under μ_B.*

[5] We take the lower order $O(n)$ terms due to the functions being permutations as subsumed in the larger $\varepsilon_B n \log n$ term due to M_1 for clarity.

▷ Claim 22. Each $a_j \in \mathcal{A}'_2$ has at least $3n/4$ locally good vertices $i \in \mathcal{V}_4$ with respect to it.

Proof. Since $H[F^{-1}_{B,2} \mid m_1, a_j] \geq (1 - 9\varepsilon_B) \log n$, by sub-additivity we have:

$$\mathbb{E}_{i\in[n]}[H[F^{-1}_{B,2}(i) \mid m_1, a_j]] \geq (1 - 9\varepsilon_B) \log n$$

The statement follows by Markov's inequality on $\log n - H[F^{-1}_{B,2}(i) \mid m_1, a_j]$. ◁

▷ Claim 23. The number of globally good $i \in \mathcal{V}_4$ is at least $\frac{n}{6}$.

Proof. Let s be the number of globally good $i \in \mathcal{V}_4$. Since every $a_j \in \mathcal{A}'_2$ has at least $3n/4$ locally good vertices, and $\mathcal{A}'_2$ has measure at least $1/2$ under μ_B, we have:

$$s \cdot 1 + (n - s) \cdot \frac{1}{4} \geq \frac{3n}{4} \cdot \frac{1}{2}$$

Rearranging gives the stated bound on s. ◁

We now make a probabilistic argument to continue the proof of Claim 11. Let $\sigma : \mathcal{A}_2 \to \mathcal{V}_3$ be a random one-to-one mapping induced by $F_{A,2}$. Define:

$$\nu_B(r \to i \mid j) \coloneqq \Pr_{F_{B,2}}\left[F^{-1}_{B,2}(i) = r \mid m_1, v_1, V_2 = a_j\right]$$

For each globally good $i \in \mathcal{V}_4$, define the random variable:

$$Z(i) := \sum_{j=1}^{\ell} \mu_B(a_j) \sum_{r=1}^{n} \mathbf{1}_{\{\sigma(a_j)=r\}} \nu_B(r \to i \mid j),$$

$Z(i)$ captures the probability mass (under F_B) that fixing a σ leads to $V_4 = i$. We note that $Z(i)$ is random purely due to the random assignment σ, which is completely determined by $F_{A,2}$. Furthermore, $F_{A,2}$ is a uniformly random permutation in the subrectangle R under consideration, so σ is a uniformly random one-to-one mapping.

▷ Claim 24. For any constant $\gamma \in (0, 1)$, and any globally good $i \in \mathcal{V}_4$ as in Definition 21,

$$\Pr_{\sigma}\left[Z(i) \leq \frac{1-\gamma}{16n}\right] \leq n \exp\left(-2\gamma^2 n^{1-O(\varepsilon_B)}\right),$$

Proof. Partition the indices $j \in [\ell]$ into

$$J_{\text{good}} := \{j \in [\ell] : i \text{ is locally good for } a_j\}, \quad J_{\text{bad}} := [\ell] \setminus J_{\text{good}},$$

with $\mu_B(a_{J_{\text{good}}}) \geq 1/4$, since i is globally good.

Let $S = \{j \in [\ell] : \mu_B(a_j) \leq n^{-(1-100\varepsilon_B)}\}$. From a calculation as in Lemma 17 (Appendix B) with $\delta = 1/8$, $A = 12\varepsilon_B \log n$ applied on the distribution μ_B (recall that $H[\mu_B] \geq (1 - 12\varepsilon_B) \log n$) , we infer that $\mu_B(a_j : j \in S^c) \leq 1/8$. This implies that:

$$\sum_{j\in J_{\text{good}}\cap S} \mu_B(a_j) \geq 1/4 - 1/8 = 1/8 \tag{2}$$

We can thus write:

$$Z(i) = \sum_{j\in J_{\text{good}}\cap S} \mu_B(a_j) \sum_{r=1}^{n} \mathbf{1}_{\{\sigma(a_j)=r\}} \nu_B(r \to i \mid j) + \sum_{j\notin J_{\text{good}}\cap S} \mu_B(a_j) \sum_{r=1}^{n} \mathbf{1}_{\{\sigma(a_j)=r\}} \nu_B(r \to i \mid j)$$

We now concentrate on the contribution from the good indices $J_{\text{good}} \cap S$ and apply the decomposition lemma (Lemma 17) to the ν_B distributions therein, since they have high entropy.

For ν_B, from the entropy bound $H[F_B^{-1}(i)|m_1, a_j] \geq (1 - 36\varepsilon_B)\log n$, we apply Lemma 17 with $A := 36\varepsilon_B \log n$ and $\delta := \frac{1}{2}$, giving:

$$\nu_B = \delta\nu_B^{(1)} + (1-\delta)\nu_B^{(2)}, \quad H_\infty[\nu_B^{(2)}] \geq \log n - 72\varepsilon_B \log n - O(1).$$

Using this decomposition, we can write $Z(i)$ as follows:

$$Z(i) = Z_1(i) + Z_2(i),$$

where

$$Z_2(i) := \frac{1}{2} \sum_{j \in J_{\text{good}} \cap S} \mu_B(a_j) \sum_{r=1}^{n} \mathbf{1}_{\{\sigma(a_j)=r\}} \nu_B^{(2)}(r \to i \mid j)$$

and $Z_1(i)$ has the remaining terms of $Z(i)$, and is a non-negative random variable. Intuitively, $Z_2(i)$ captures components of the distribution ν_B having high *min-entropy*.

We now bound the probability of deviation of $Z_2(i)$ from its mean. Clearly, for any β, $\Pr(Z(i) \leq \beta) \leq \Pr(Z_2(i) \leq \beta)$. We compute the expectation:

$$\begin{aligned}\mathbb{E}[Z_2(i)] &= \frac{1}{2} \sum_{j \in J_{\text{good}} \cap S} \mu_B(a_j) \sum_{r=1}^{n} \mathbb{E}[\mathbf{1}_{\{\sigma(a_j)=r\}}] \nu_B^{(2)}(r \to i \mid j) \\ &= \frac{1}{2} \sum_{j \in J_{\text{good}} \cap S} \mu_B(a_j) \sum_{r=1}^{n} \frac{1}{n} \nu_B^{(2)}(r \to i \mid j) \geq \frac{1}{16n} \end{aligned} \tag{3}$$

The last step follows from Eq. 2. We now apply a deviation bound due to McDiarmid Theorem 19:

▶ **Theorem 25** ([12], [6, The method of averaged bounded differences, Corollary 5.1]). *Let $X_1, \ldots, X_n$ be random variables and f a function such that for each $i \in [n]$ there exists $c_i \geq 0$ with*

$$\left|\mathbb{E}[f \mid X_1, \ldots, X_{i-1}, X_i = a_i] - \mathbb{E}[f \mid X_1, \ldots, X_{i-1}, X_i = a_i']\right| \leq c_i$$

for all choices of a_i, a_i'. Then, for any $t \geq 0$,

$$\Pr(|f - \mathbb{E}[f]| \geq t) \leq 2\exp\left(-\frac{2t^2}{\sum_{i=1}^{n} c_i^2}\right).$$

We apply the above, for f being $Z_2(i)$ and the random variable $X_j \equiv \sigma(a_j)$. We bound $c := \sum_{j=1}^{n} c_j^2$ for $Z_2(i)$: each term in $Z_2(i)$ is of the form $\mu_B(a_j) \cdot \nu_B^{(2)}(r \to i \mid j)$, and is non-zero only for one $r = \sigma(a_j)$. Thus, each contributes at most:

$$\max \mu_B(a_j) \cdot \max \nu_B^{(2)}(r \to i \mid j) \leq n^{-1+100\varepsilon_B} \cdot n^{-1+72\varepsilon_B+o(1)} = n^{-2+172\varepsilon_B+o(1)}.$$

Each c_j^2 is at most $n^{-4+O(\varepsilon_B)}$, and we have at most n terms. This gives us:

$$c \leq n \cdot n^{-4+O(\varepsilon_B)} = n^{-3+O(\varepsilon_B)}.$$

Applying the concentration bound with $t = \gamma/16n$, we get:

$$\Pr_\sigma\left[Z(i) \leq \frac{1-\gamma}{16n}\right] \leq \Pr_\sigma\left[Z_2(i) \leq \frac{1-\gamma}{16n}\right] \leq \exp\left(-\frac{2\gamma^2}{256n^2 \cdot c}\right) = \exp\left(-\frac{1}{128}\gamma^2 n^{1-O(\varepsilon_B)}\right) \quad \triangleleft$$

Setting $b_1, \ldots b_r$, to be the set of globally good indices i, $\gamma = 1/5$, and a union bound, we conclude the proof of Claim 11. $\triangleleft$

Proof of Claim 12

Proof of Claim 12. Let Alice's message M_2 have $|M_2| \leq \varepsilon_A n \log n$. We know that before M_2 is passed, for *every* σ, since $F_{A,3}$ is independent of σ (and v_1) $H\left[F_{A,3} \mid \sigma\right] = n \log n - O(n)$.

In particular, for $\sigma = \sigma^*$, with probability $\geq 3/4$ over the choice of M_2, $H\left[F_{A,3} \mid \sigma^*, M_2 = m_2\right] \geq (1 - 4\varepsilon_A) n \log n$. Thus, within R' on fixing such an m_2, we get a sub-rectangle R'' where, by subadditivity of entropy:

$$\begin{aligned} & \mathbb{E}_{i \in [n]}[H[F_{A,3}(i) | R'']] \geq (1 - 4\varepsilon_A) n \log n \\ \implies \exists S \subseteq \mathcal{V}_4 : |S| \geq (1 - 2\sqrt{\varepsilon_A}) n \text{ such that} & \\ \forall i \in S : \quad & H[F_{A,3}(i) | R'']] \geq (1 - 2\sqrt{\varepsilon_A}) \log n \end{aligned}$$

Call S the *uncertain* vertices, since Bob cannot predict the next pointer (V_5) starting at these with reasonable certainty after receiving $M_2 = m_2$. Intuitively, it is sufficient to show that in R'', V_4 occurs as an uncertain vertex with constant probability. We formalize this next.

Within R'', the value of V_4 is known exactly to Bob: it is the same across $F_{A,3}$, for any fixed f_B. This is because $V_1 = v_1, \sigma = \sigma^*$ are fixed. Thus, for a fixed f_B, if $v_4 \in S$, $H[V_5] \geq (1 - 2\sqrt{\varepsilon_A}) \log n$. Applying Fano's inequality (Fact 16.0.1), the probability of Bob making an error in finding the value of V_5 is at least $P_e \geq (1 - 2\sqrt{\varepsilon_A} - \frac{1}{\log n}) \geq 0.9$, for ε_A being small enough.

Since σ^* satisfies Eq. (1), we get a set $\mathcal{B} = \{b_1, \ldots, b_r\}$ of globally good vertices satisfying Claim 11. Then $|\mathcal{B} \cap S| \geq (1/6 - 2\sqrt{\varepsilon_A}) n \geq n/10$, and each has probability at least $\frac{1}{20n}$. Thus, the probability of Bob seeing an uncertain vertex as v_4 in R'' is at least $1/200$.

Overall, the probability of error within R' is:

$$\frac{3}{4} \cdot \frac{1}{200} \cdot 0.9 \geq \frac{1}{400}$$

This concludes the proof of Claim 12. ◁

E Missing proofs from Section 5

Proof of Claim 14

Proof of Claim 14. Let $F := F_{B,3}$. Since $|\mathcal{A}_4| \geq n/6$, $H[F \mid R'] \geq (1 - 4\varepsilon_B) n \log n$, and $H(V_4) \leq \log n$, we obtain

$$H[F \mid R', V_4] \geq (1 - 4\varepsilon_B) n \log n - \log n \geq (1 - 5\varepsilon_B) n \log n.$$

There are at least $n/12$ elements in $\mathcal{A}_4$ satisfying $\Pr[V_4 = i] \in (\frac{1}{20n}, \frac{1}{10n})$; call this set $\mathcal{D}$. Let $I[F] := n \log n - H[F]$. From the above, $\mathbb{E}_{v \sim V_4}[I[F|R', V_4 = v]] \leq 5\varepsilon_B n \log n$. Let c_1 be a large constant to be set later, and $\mathcal{D}_{\text{good}} = \{i \in \mathcal{D} \ : \ I[F|R', V_4 = i] \leq c_1 \varepsilon_B n \log n\}$.

Let μ be the induced distribution on V_4 in R'. We have $\mu(\mathcal{D}) \geq 1/240$. By Markov's inequality, $\mu(\mathcal{D}_{\text{good}}) \geq \mu(\mathcal{D}) - \frac{5}{c_1} \geq \frac{1}{240} - \frac{5}{c_1} \geq \frac{1}{480}$, for $c_1 := 5 \times 480$.

But this means that we should have at least $\frac{1}{480 \cdot \frac{1}{10n}} = \frac{n}{48} =: \alpha_1 n$ elements in $\mathcal{D}_{\text{good}}$, since each element has probability at most $1/10n$. ◁

Proof of Claim 16

Proof of Claim 16. Define $\{Z_i := \mathbb{1}_{[\mathcal{E}_i]}\}_{i \in [r]}$ to be indicator variables for the events $\mathcal{E}_i$ under a random choice of $F_{A,3}$. For any i, we have $\mathbb{E}[Z_i] \geq (1 - \gamma)$. Thus, $\mathbb{E}[\sum_{i=1}^{r} Z_i] \geq (1 - \gamma) r$. Let $\mathcal{E} \equiv \sum_i Z_i \geq (1 - 10\gamma) r$.

▷ Claim 26. $\Pr\left(\sum_{i=1}^{r} Z_i \geq (1-10\gamma)r\right) \geq 1 - \exp(-\Omega(\gamma^2 n))$

Proof. We use McDiarmid's inequality (see Theorem 19) on the function $f \coloneqq \sum_{i\in[r]} Z_i$. The average bounded difference at i satisfy $c_i \leq r/(n-r) \leq 2\alpha_1$.

$$\Pr(Z \leq (1-10\gamma)r) \;\leq\; \exp\left(-\Omega(81\gamma^2(n-r)^2)\right) = \exp(-\alpha_2\gamma^2 n)\ldots \text{ for some constant } \alpha_2$$

◁

Let $F \equiv F_{A,3}$ for convenience. Let $Z = (Z_1, \ldots, Z_r)$. Note that $\mathcal{E} \implies |Z| \geq (1-10\gamma)r$. Since $H[F] \geq n\log n - O(n)$, we have:

$$\begin{aligned} & H[F \mid M, Z] \geq (1-\varepsilon_A)n\log n - O(n) \\ \implies\; & H[F \mid M, Z, \mathcal{E}] \geq \frac{(1-2\varepsilon_A)n\log n - \Pr(\mathcal{E}^c)n\log n}{\Pr(\mathcal{E})} \\ \implies\; & H[F \mid M, Z, \mathcal{E}] \geq (1-4\varepsilon_A)n\log n \end{aligned}$$

Hence, conditioned on $\mathcal{E}$, with probability at least $3/4$ over choices of $(M, Z) = (m, z)$, we have that $H[F \mid M = m, Z = z, \mathcal{E}] \geq (1-16\varepsilon_A)n\log n$. Fix such a (m, z) pair, and let $z(i_1) = 1, \ldots, z(i_\ell) = 1$ for some $\ell \geq (1-10\gamma)r$. By subadditivity of entropy:

$$\mathbb{E}_{i\in[n]} H[F(i)|M = m, Z = z, \mathcal{E}] \geq (1-16\varepsilon_A)\log n$$

By Markov's inequality, there are at least $(1-4\sqrt{\varepsilon_A})n$ indices i where $H[F(i)|M = m, Z = z, \mathcal{E}] \geq (1-4\sqrt{\varepsilon_A})\log n$. From these, at least $(\alpha_1 - 4\sqrt{\varepsilon_A})n$ coordinates have to be from $\mathcal{A}'_4$. For each of these, we have that $H[F(b_i) \mid M = m, \mathcal{E}_i] \geq (1-4\sqrt{\varepsilon_A})\log n$. Let $\alpha_2 \coloneqq \alpha_1 - 4\sqrt{\varepsilon_A}$.

Thus, the probability (over M, Z, under the distribution of F_A) that we get at least $\alpha_2 n$ indices from $\mathcal{A}$, such that $H[F_{A,3}(b_i)|M = m, Z = z, \mathcal{E}_i] \geq (1-4\sqrt{\varepsilon_A})\log n$ is at least:

$$\Pr[\mathcal{E}] \cdot \frac{3}{4} \geq 1/8$$

◁

F Upper bounds

We adapt the protocols from [5, 15] to infer what depth of pointer can be chased with $O(n\log\log n)$ bits. For simplicity of notation, while describing protocols, we let $f_{A,i} \equiv f_{A,j}$ and $f_{B,i} \equiv f_{B,j}$ all i, j. We will incur a multiplicative factor of k in the communication cost if the functions are different.

▶ **Theorem 27.** *Let $k' \geq \left\lceil \frac{k+5}{3} \right\rceil$. Then $C^{B,k'}(\mathsf{PC}(n,k)) = O(nk\log\log n)$.*

Proof. Let $n' \coloneqq \frac{n}{\log n}$.

- Round 1 (Bob → Alice): Bob sends the last $\log\log n$ bits of $f_B(v_b)$ for every $v_b \in V_B$.
- Round 2 (Alice → Bob): Alice sends the last $\log\log n$ bits of $f_A(v_a)$ for every $v_a \in V_A$, and also sends $v_1 = f_A(v_0)$ along with $f_A(v)$ for every possible vertex in Bob's list that v_1 could point to.
- Round 3 (Bob → Alice): Bob computes $v_2 = f_B(v_1)$. From Alice's message, he knows $v_3 = f_A(v_2)$. He then computes $v_4 = f_B(v_3)$ and sends v_4 along with $f_B(v)$ for every possible vertex in Bob's list that v_4 could point to.

- Round 4 (Alice → Bob): Alice computes $v_5 = f_A(v_4)$. From Bob's message, she knows $v_6 = f_B(v_5)$. She then computes $v_7 = f_A(v_6)$ and sends v_7 along with $f_A(v)$ for every possible vertex in Bob's list that v_7 could point to.
- ...
- Round i: If i is even, Alice computes $v_{3i-7} = f_A(v_{3i-8})$, knows the possibilities for $v_{3i-6} = f_B(v_{3i-7})$ from Bob's previous message, computes $v_{3i-5} = f_A(v_{3i-6})$, and sends v_{3i-5} along with $f_A(v)$ for all n' possible vertices;
  if i is odd, Bob performs the analogous computation and sends v_{3i-5} along with $f_B(v)$ for all n' possible vertices.
- ...
- **Round k'**: If k' is even, Alice knows $v_{3k'-5}$ and sends it to Bob; if k' is odd, Bob knows $v_{3k'-5}$ and sends it to Alice, where $k' \geq \lceil \frac{k+5}{3} \rceil$.

Intuitively, Alice knowing the pointer values for one layer ahead from Bob allows her to "jump" forward 3 pointers per round: two from her layers and one from Bob.

Rounds 1 and 2 together take $O(n \log\log n)$ bits. All other rounds (except the last round) take $O(\log n + \frac{n}{\log n} \cdot \log n) = O(n)$ bits per round. Since the number of rounds k' is constant, the total communication is $O(n \log\log n)$ bits. ◀

▶ **Theorem 28.** *Let $k' \geq \lceil \frac{k+6}{4} \rceil$. Then $C^{B,k'}(\mathsf{bPC}(n,k)) = O(nk \log\log n)$.*

Proof. The protocol follows the same communication structure as in the $\mathsf{PC}(n,k)$ protocol, with each party in every round (starting from round 2), each party *additionally* sending the following n bits every round:

- Round 1 (Bob → Alice): Bob sends the last $\log\log n$ bits of $f_B(v)$ for every vertex v, which also implicitly communicates the bit vectors $\chi_1(v) = \mathit{parity}(f_B(v))$ across all vertices.
- Round 2 (Alice → Bob): The bit vector $\chi_2(v) := \chi_1(f_A(v))$ for all $v \in V_A$.
- Round 3 (Bob → Alice): The bit vector $\chi_3(v) := \chi_2(f_B(v))$ for all $v \in V_B$.
- Round 4 (Alice → Bob): The bit vector $\chi_4(v) := \chi_3(f_A(v))$ for all $v \in V_A$.
- Round 5 (Bob → Alice): The bit vector $\chi_5(v) := \chi_4(f_B(v))$ for all $v \in V_B$.
- ...
- Round i: If i is even, Alice sends $\chi_i(v)$ for all $v \in V_A$, else Bob sends $\chi_i(v)$ for all $v \in V_B$.
- ...
- Round k': If k' is even, Alice knows $v_{3k'-5}$ from the base $\mathsf{PC}(n,k)$ protocol, and she learns $\chi_{k'}$ from the extra parity information exchanged from Round 2. Thus, she sends $\chi_{k'-\delta}(v_{3k'-5}) = \mathrm{parity}(v_{4k'-6-\delta})$ to Bob, where $\delta \in \{0,1,2,3\}$ is chosen so that the resulting vertex is exactly v_k. If k' is odd, Bob analogously computes and sends $\mathrm{parity}(v_k)$ to Alice.

Rounds 1 and 2 together take $O(n \log\log n)$ bits. All other rounds (except the last round) have n additional bits compared to the $\mathsf{PC}(n,k)$ protocol, but since the number of rounds k' is constant, we still have the same total communication of $O(n \log\log n)$ bits. ◀