

An $\Omega((\log n / \log \log n)^2)$ Cell-Probe Lower Bound for Dynamic Boolean Data Structures*

Young Kun Ko[†]

Abstract

We resolve the long-standing open problem of Boolean dynamic data structure hardness, proving an unconditional lower bound of $\Omega((\log n / \log \log n)^2)$ for the Multiphase Problem of Pătraşcu [STOC 2010] (instantiated with Inner Product over $\mathbb{F}_2$). This matches the $\tilde{\Omega}(\log^2 n)$ bound that Larsen [STOC 2012] established for *weighted* problems, where each query returns $\Omega(\log n)$ bits, and closes the gap left by the $\tilde{\Omega}(\log^{1.5} n)$ Boolean bound of Larsen, Weinstein, and Yu [STOC 2018]. It also answers, up to a $(\log \log n)^2$ factor, Open Question 15.24 in Viola’s textbook *Mathematics of the Impossible* [Cambridge 2026].

The previous barrier was methodological: all prior works relied on “one-way” communication games, where the inability to verify query simulations necessitated complex machinery (such as the Peak-to-Average Lemma) that hit a hard ceiling at $\log^{1.5} n$. Our key contribution is conceptual: we introduce a 2.5-round Multiphase Communication Game that augments the standard one-way model with a verification round, where Bob confirms the consistency of Alice’s simulation against the actual memory. This simple, qualitative change removes the obstruction and yields the full bound directly. As a consequence, our analysis naturally extends to other hard Boolean functions — those of small discrepancy under product distributions with large min-entropy — offering a general recipe for translating discrepancy lower bounds into $\Omega((\log n / \log \log n)^2)$ dynamic Boolean data structure lower bounds.

We also argue that this result likely represents the structural ceiling of the Chronogram framework initiated by Fredman and Saks [STOC 1989]: any $\omega(\log^2 n)$ lower bound would require either fundamentally new techniques or major circuit complexity breakthroughs.

*A preliminary version of this paper appears in the Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS 2026).

[†]Department of Computer Science and Engineering, Pennsylvania State University, University Park, PA, USA.
ykkko@psu.edu

1 Introduction

Proving unconditional lower bounds for dynamic data structures in the cell-probe model [Yao81] stands as one of the grand challenges in theoretical computer science. In this model, memory is organized into fixed-size cells of $w = \Theta(\log n)$ bits, and we charge unit cost for accessing a cell while permitting arbitrary computation on the probed data for free. Because of this unbounded computational power, the cell-probe model strictly subsumes all realistic architectures; consequently, any lower bound proved here applies universally. However, this generality comes at a steep price: the model’s ability to arbitrarily compress and encode information makes proving lower bounds notoriously difficult.

Despite more than 35 years of sustained effort since Fredman and Saks’ seminal Chronogram method [FS89], progress has been slow (see Figure 1). The field advanced from $\Omega(\log n / \log \log n)$ to $\Omega(\log n)$ over 15 years [PD06], and then, over another 6 years, to $\Omega((\log n / \log \log n)^2)$ [Lar12] — but only for *weighted* problems, where each query returns $\Omega(\log n)$ bits.

For Boolean (decision) problems, where each query returns a single bit, the $\Omega((\log n / \log \log n)^2)$ barrier seemed insurmountable, and the question was posed from two directions. Pătraşcu believed that dynamic 2-dimensional range parity — report whether the number of points inside a query rectangle is odd — requires $\tilde{\Omega}(\log^2 n)$, having proved an $\tilde{\Omega}(\log n)$ query lower bound for the static problem under near-linear space [Pat07]; Thorup lists this among the five open problems Pătraşcu regarded as most important [Tho13].¹ Viola’s textbook, which spans computational complexity as a whole, poses only twelve open questions; one of them [Vio26, Open Question 15.24] asks for exactly this threshold: “Are there explicit (say in FP) boolean functions $f : [2]^n \rightarrow [2]^m$ with m a power of n that require time $\geq \log^2 n$ for dynamic data structures?”

The obstruction was that no technique reached this threshold at all. Larsen’s $\Omega((\log n / \log \log n)^2)$ bound [Lar12] applies only to the weighted variant, where each point carries an $O(\log n)$ -bit weight and the query asks for the range sum. Such weights can encode more information than a memory address, and Larsen’s technique exploits this heavily; the argument does not survive the passage to a single output bit. Larsen, Weinstein and Yu [LWY20] extended the method to the Boolean setting, but the Peak-to-Average Lemma they introduced loses a $\sqrt{\log n}$ factor and stops at $\tilde{\Omega}(\log^{1.5} n)$ — a loss that is inherent to the lemma [LWY20, Appendix B].




Figure 1: Historical progression of dynamic cell-probe lower bounds. Throughout, $\tilde{\Omega}$ hides a $\text{poly}(\log \log n)$ factor in the denominator.

In this paper, we attain the threshold. We prove an unconditional lower bound of $\Omega((\log n / \log \log n)^2)$ for the Multiphase Problem with Inner Product over $\mathbb{F}_2$, the first such bound for a Boolean problem, matching the highest known bound for weighted problems and answering, up to a $(\log \log n)^2$ factor, Viola’s Open Question 15.24 [Vio26]. Our result directly implies lower bounds for fundamental problems including dynamic matrix-vector multiplication over $\mathbb{F}_2$, dynamic path parity, and the counting version of Erickson’s Problem (Problem 4).

¹Our hard problem is the Multiphase Problem rather than range parity; an $\Omega((\log n / \log \log n)^2)$ bound for range parity itself remains open.

1.1 Our Contribution

As our hard problem, we consider the Multiphase Problem of Pătraşcu [Pat10]. For a Boolean function $f : \{0, 1\}^{2n} \rightarrow \{\pm 1\}$, the problem is:

- $\vec{S} = S_1, \dots, S_m \in \{0, 1\}^n$ are pre-processed.
- The bits of $X \in \{0, 1\}^n$ arrive one at a time, as a sequence of n single-bit updates, each costing t_u probes.
- Per any given $q \in [m]$, output $f(S_q, X)$ using t_{tot} total probes.

Problem 1: Multiphase Problem for f

We stress that a single update sets one bit of X , so that the vector X specifies n updates rather than a single one, at a total update cost of $t_u \cdot n$ probes. It is this sequence of n single-bit updates that the Chronogram method partitions into epochs; see the *Epochs* paragraph of Section 3.

When f is Disjointness or Inner Product over $\mathbb{F}_2$, Pătraşcu conjectured that $\max\{t_u, t_{tot}\} \geq n^{\Omega(1)}$ for $m = n^{1+\Omega(1)}$ [Pat10] — the notorious Multiphase Conjecture [Pat10, Tho13, Bra22]. Pătraşcu showed that this conjecture implies polynomial lower bounds for many fundamental dynamic problems including dynamic s - t reachability.

We show the following main theorem for the Multiphase Problem.

Theorem 1.1. *Let f be Inner Product over $\mathbb{F}_2$ and let $m = n^{1+\Omega(1)}$. Then any data structure for the Multiphase Problem for f whose update time satisfies $t_u = n^{o(1)}$ must have*

$$t_{tot} \geq \Omega \left(\left(\frac{\log n}{\log(t_u w)} \right)^2 \right).$$

In particular,

$$\max\{t_u, t_{tot}\} \geq \Omega \left(\left(\frac{\log n}{\log \log n} \right)^2 \right).$$

We prove Theorem 1.1 in Section 3.

General Lifting Theorem While Theorem 1.1 focuses on the Inner Product over $\mathbb{F}_2$ to resolve the specific open problem posed by Pătraşcu and Larsen, our technique is not limited to the Inner Product over $\mathbb{F}_2$. Our Simulation Theorem (Section 3.1) establishes a general translation from communication complexity to dynamic data structure lower bounds. Our lower bound applies to *any* function f that is “hard.” Roughly, f is hard if no *product distribution* with large min-entropy allows an advantage greater than $n^{-\Omega(1)}$ (i.e., f has low discrepancy under such distributions). This hardness condition is indeed satisfied by the Inner Product over $\mathbb{F}_2$ function, but unfortunately, the Disjointness function does not meet the criterion [BM13]. A naive strategy of sampling and revealing a few coordinates of X already yields a nontrivial advantage for Disjointness, violating the hardness condition. We refer the reader to Appendix B, specifically Definition B.1, for the precise technical definition. We then prove a general “lifting” theorem (Theorem B.3) showing that if f is “hard,” then the Multiphase Problem with f must have, for $m = n^{1+\Omega(1)}$ and $t_u = n^{o(1)}$,

$$t_{tot} \geq \Omega \left(\left(\frac{\log n}{\log(t_u w)} \right)^2 \right).$$

We defer the full proof to Appendix B to preserve the clarity of the main argument for Theorem 1.1.

1.2 Applications

Via the standard reductions of Pătraşcu [Pat10], this directly implies lower bounds for the following problems.

- A matrix $M \in \mathbb{F}_2^{m \times n}$ is given and pre-processed.
- $X \in \{0, 1\}^n$ is updated dynamically. Output $\langle M_q, X \rangle$ for any given $q \in [m]$.

Problem 2: Dynamic Matrix-Vector Multiplication over $\mathbb{F}_2$

- A directed graph $G = (V, E)$ is initially given. Updates add and remove edges in the graph.
- For a given $(u, v) \in V \times V$, output the parity of the number of paths between u and v . For the k -Path Parity Problem, output the parity of the number of paths of length exactly k .

Problem 3: Dynamic (k -) Path Parity

- A matrix of integers is preprocessed. Updates increment all values in a specified row or column.
- For a given query, output the parity of the number of maximum values in the matrix.

Problem 4: Erickson's Problem, Counting Version

Corollary 1.2. *Problem 2, Problem 3 and Problem 4 require $\max\{t_u, t_{tot}\} \geq \Omega\left(\left(\frac{\log n}{\log \log n}\right)^2\right)$.*

1.3 Previous Works

To put our technical contribution into perspective, we briefly summarize the technical developments in dynamic cell-probe lower bounds. Readers familiar with these developments can skip to Section 1.4 for our technical contribution.

1.3.1 Explicit Quests for Poly-logarithmic Lower Bounds

Chronogram Method Fredman and Saks [FS89] introduced the Chronogram method. At a high level, the Chronogram method (Figure 2) divides the sequence of n updates X into ℓ epochs, $\{X^{(i)}\}_{i=1}^{\ell}$, where $X^{(i)}$ consists of n_i bits and $n_i = |X^{(i)}| = \beta^i$, where β is a parameter set per problem in question (roughly $\text{poly}(wt_u)$). We then process the updates in reverse order, that is, the larger epochs are processed first, and we denote the updates in epoch i by U_i .

The main observation by Fredman and Saks is that we can decompose the updated cells depending on the latest epoch that performed the update. That is $C_1, \dots, C_\ell$, where each $C_i \subseteq U_i$ contains the cells last updated by epoch i , and no longer updated in epochs $< i$. The key point is that there must exist $i \in [\ell]$ such that the query algorithm probes at most $O(t_{tot}/\ell)$ cells from C_i on average. That is $\mathbb{E}_{q \sim Q}[t_q[i]] \leq O(t_{tot}/\ell)$. This was elegantly phrased in [LY25] as a “static” data structure with pre-initialized memory and a cache, where pre-initialized memory refers to $C_{>i}$ and cache refers to $C_{\leq i}$. The goal then is to show a lower bound on t_{tot}/ℓ , which is the average number of probes into C_i . Fredman and Saks proved an $\Omega(1)$ lower bound on t_{tot}/ℓ , leading




Figure 2: The Chronogram technique decomposes n updates into $\ell = \Theta(\log_{\beta} n)$ epochs with geometrically growing sizes $n_i = \beta^i$. Crucially, epochs are **processed in reverse order** (largest first). Each epoch i has associated cells C_i that were last updated in that epoch. A query algorithm for $q \in \mathcal{Q}$ probes $t_q[i]$ cells from each C_i .

to $\max\{t_u, t_{tot}\} \geq \Omega(\log n / \log \log n)$. It then took more than 15 years for further quantitative improvement. A more careful analysis by considering the information transfer between the epochs [PD06] led to shaving the $\log \log n$ factor in the denominator.

Chronogram + Cell-Sampling The breakthrough work of Larsen [Lar12] then successfully merged the cell-sampling technique [Sie04, PTW10] with the Chronogram technique [FS89] to give an $\Omega((\log n / \log \log n)^2)$ bound with the following rough idea. Recall that after the decomposition using the Chronogram method, the goal is to show a lower bound on t_{tot}/ℓ , which is the number of probes into C_i . As ℓ would always be roughly $\Theta(\log n / \log(t_u w))$, if we can show, say, $t_{tot}/\ell \geq \tilde{\Omega}(\log n)$, this would give an $\tilde{\Omega}(\log^2 n)$ bound on t_{tot} .

To show such a lower bound on t_{tot}/ℓ , Larsen [Lar12] utilizes the cell-sampling method [Sie04, PTW10]. The main idea behind cell-sampling is the following. If there exists a too-good-to-be-true (i.e., short) t_{tot} query algorithm, then naively sampling some $p := (\text{poly}(t_u w))^{-1}$ fraction of the cells would give a nontrivial advantage (of roughly $p^{t_{tot}/\ell}$) in answering the queries. This is nontrivial in the regime of interest: with $\ell = \log_{\beta} n = \log n / \log \beta$ and $\beta = (t_u w)^2$, both $1/\ell$ and $\log(1/p)$ contribute a factor of $\log(t_u w)$, so that

$$p^{t_{tot}/\ell} = \exp \left(-\Theta \left(\frac{t_{tot} (\log(t_u w))^2}{\log n} \right) \right) = n^{-o(1)}$$

exactly when $t_{tot} = o((\log n / \log(t_u w))^2)$ — the bound we aim to prove, and the source of its quadratic form. An equivalent way to frame this argument is to find a small subset $C_0 \subset C_i$ such that C_0 along with $C_{<i}$ and $C_{>i}$ will answer $n^{-o(1)}$ fraction of possible queries, without having to know the whole C_i . We can generate such C_0 by simply sampling each cell in C_i at random with probability p . By averaging, there must exist a setting of such C_0 which can answer $n^{-o(1)}$ fraction of possible queries, say $\mathcal{Q}^* \subset \mathcal{Q}$.

An elegant way to formulate the underlying combinatorial problem is through the following

one-way communication simulation, as the ultimate goal is to show such one-way communication cannot exist.

- Fix $\vec{S}$ known to both Alice and Bob. Bob is given all the updates $\{X^{(i)}\}_{i=1}^\ell$, but not the desired query $Q \in \mathcal{Q}$.
- Alice is given all the updates, except the updates in epoch i (i.e., $X^{(i)}$), and the desired query Q .
- Bob sends a one-way message M of length $|M| \leq p \cdot (wt_u) \cdot n_i + o(n_i)$ bits to Alice. Then Alice announces $f(S_Q, X)$.

Protocol 1: One-Way Communication Simulation

As a one-way message, Bob will send the small sub-sampled subset $\tilde{C}_i$ along with $C_{<i}$. Cell-sampling ensures that $\tilde{C}_i$ is small, while the Chronogram ensures that $C_{<i}$ is small. Alice then can simulate the queries for $\mathcal{Q}^*$. So for any $Q \in \mathcal{Q}^*$, Alice will be outputting the correct answer, achieving a nontrivial advantage! We then just need to show that a short one-way message M cannot answer an $n^{-o(1)}$ fraction of $\mathcal{Q}$.

Though tempting, this argument has a crucial flaw. $\tilde{C}_i$ cannot identify $\mathcal{Q}^*$ (i.e., the $\mathcal{Q}^*$ -identification problem). This occurs because Alice cannot distinguish whether the probed cells are inside $C_i \setminus \tilde{C}_i$ or absent from C_i entirely. As she cannot identify the set $\mathcal{Q}^*$, there is no guarantee that Alice's guess provides zero advantage over completely random guessing for queries in $\mathcal{Q} \setminus \mathcal{Q}^*$. This technical subtlety arises in all subsequent works proving an $\omega(\log n)$ lower bound [Pat07, Lar12, LWY20, LY25].

Larsen [Lar12] tackled this technical subtlety by encoding a subset of $\mathcal{Q}^*$ as a part of Bob's message, but this argument crucially requires the output to be $\Omega(\log n)$ bits. To extend the method to the Boolean setting, Larsen, Weinstein, and Yu [LWY20] introduced the so-called Peak-to-Average Lemma, a proof of which requires a technical dive into Chebyshev polynomials. The lemma allows Bob to carefully choose a subset of cells, knowledge of which yields a nontrivial advantage (of $n^{-o(1)}$) but with a $\sqrt{\log n}$ parameter loss. No improvement can be made to the lemma, as the lemma is tight (Appendix B of [LWY20]).

1.3.2 Previous Work on the Multiphase Problem

Pătraşcu envisioned a bold approach [Pat10, Tho13] to break out of the Chronogram method and its extensions (all the methods in Section 1.3.1), providing a new angle on dynamic Boolean cell-probe lower bounds.

His ambitious goal was to directly prove a polynomial lower bound by considering Problem 1. He conjectured that Problem 1 requires $\max\{t_u, t_{tot}\} \geq n^{\Omega(1)}$ for $m = n^{1+\Omega(1)}$, even when f is Disjointness. This is the notorious Multiphase Conjecture. Pătraşcu then showed reductions to various dynamic problems such as dynamic reachability, among others.

It is worth noting that Pătraşcu [Pat10] initiated the whole area of Fine-Grained Complexity of Dynamic Problems using the Multiphase Problem (Problem 1). He showed that the 3SUM Conjecture² implies the Multiphase Conjecture. We refer the reader to the survey [Wil17] for the context of the Multiphase Conjecture in the area of Fine-Grained Complexity of Dynamic Problems.

²The conjecture states that the 3SUM problem on n integers from $\{-n^4, \dots, +n^4\}$ cannot be solved in $O(n^{2-\varepsilon})$ time for any $\varepsilon > 0$.

While the problem is conjectured to be polynomially hard, the best unconditional lower bound stood at merely $\Omega(\log n)$ for over a decade against the weaker bit-probe model (i.e., the word size $w = 1$) [BL15, CGL15, KW20].³ Pătraşcu’s proposed approach was to consider the following communication game, which is now commonly referred to as the Multiphase Communication Game.

- Alice holds $\vec{S} \in (\{0, 1\}^n)^m$, and $Q \in [m]$. Bob holds $X \in \{0, 1\}^n$ and $Q \in [m]$. Merlin holds $\vec{S}$ and X .
- Merlin sends U to Bob of length $s = O(n \cdot t_u \cdot w)$.
- Alice and Bob engage in standard two-party communication of $w \cdot t_{tot}$ bits to output $f(S_Q, X)$.

Protocol 2: Multiphase Communication Game

The communication model can then simulate the update and query algorithms, as Merlin can simulate all the update algorithms, with $s = w \cdot n \cdot t_u$. Then Alice and Bob can simulate the query algorithm using $w \cdot t_{tot} = \tilde{O}(t_{tot})$ bits of communication to output $f(S_Q, X)$.

In fact, as pointed out in [KW20, Ko25a], even a lower bound against 3-round communication between Alice and Bob (Alice, Bob, then Alice speak) would imply a breakthrough in circuit complexity, namely circuit lower bounds for random linear operators [JS10, Juk12, Juk13, Dru12]. Nevertheless, Ko and Weinstein [KW20] developed an information-theoretic framework to handle 2-round communication (Alice speaks, then Bob) in the zero-error or small-error regime. This framework was further extended by Ko [Ko25a] to handle the low-advantage regime, where the performance of Bob’s final output is compared against a random guess.

Recently, Ko [Ko25b] combined the information-theoretic framework [KW20, Ko25a] with the one-way simulation theorem of Larsen, Weinstein, and Yu [LWY20], to give an $\tilde{\Omega}(\log^{3/2}(n))$ lower bound for Problem 1 when f is Inner Product over $\mathbb{F}_2$. Since that work applied the one-way simulation theorem of Larsen, Weinstein, and Yu [LWY20] directly and then used information-theoretic methods to lower bound the one-way communication, it could yield at most $\tilde{\Omega}(\log^{3/2}(n))$.

1.4 Technical Contribution: The 2.5-Round Communication Game

We resolve the Boolean hardness question by introducing a fundamentally different communication model that eliminates the technical barriers faced by all previous approaches.

1.4.1 The Core Challenge

All previous super-logarithmic lower bounds [Lar12, LWY20, LY25] rely on the Chronogram and cell-sampling techniques, which reduce the problem to proving that no short *one-way* message from Bob to Alice can help answer an $n^{-o(1)}$ fraction of queries in Protocol 1. However, as discussed in Section 1.3.1, there is a fatal technical subtlety: Alice does not know *which queries* she can answer correctly (the set $\mathcal{Q}^*$), because she cannot distinguish between:

- probed cells missing from Bob’s randomly sampled subset;
- probed cells that were never written to memory at all.

³While [BL15, CGL15, KW20] do not explicitly mention the lower bound, one can simulate a t -round adaptive algorithm with a 2^t -round nonadaptive algorithm in the bit-probe model.




Figure 3: Our 2.5-round Multiphase Communication Game. The crucial difference from prior work is the **verification step** (Round 2): Bob checks if Alice’s transcript A_Q is consistent with the actual memory U . The 0.5-round message Φ is sent independently of query Q , making this analyzable with information-theoretic techniques from [KW20, Ko25a].

This ambiguity destroys the advantage argument. Larsen, Weinstein, and Yu [LWY20] resolved this using the Peak-to-Average Lemma, a sophisticated application of Chebyshev polynomials, but this incurs a $\sqrt{\log n}$ parameter loss, yielding only $\tilde{\Omega}(\log^{1.5}(n))$. The lemma is provably tight [LWY20, Appendix B].

1.4.2 Our Key Insight

The one-way communication prevents Alice from verifying whether her simulation is correct. But if we add just *an additional round* — letting Bob verify and certify Alice’s guess — the entire technical barrier disappears.

We introduce a *2.5-round Multiphase Communication Game*:

1. **Round 0 (Merlin $\rightarrow$ Bob):** Merlin (holding $\vec{S}, X$) sends update information U to Bob.
2. **Round 0.5 (Bob $\rightarrow$ Alice):** Bob (holding X) forwards a message Φ to Alice (which may depend on U , X and Bob’s randomness, but not on the query).
3. **Round 1 (Alice $\rightarrow$ Bob):** The query Q is revealed to Alice; Alice (holding $\vec{S}$ and Q) sends her simulation transcript A_Q to Bob.

4. **Round 2 (Bob outputs):** Bob verifies A_Q against U and outputs the final answer.

The protocol structure is illustrated in Figure 3.

Verification Step The verification step is crucial:

- Alice simulates the query algorithm and sends the full transcript to Bob.
- Bob checks if Alice’s simulation is *consistent* with the actual updated memory U .
- If the query probes a cell of C_i that Bob did not sample, Alice can only supply its content from before epoch i , reconstructed from $\vec{S}$ and $X^{(>i)}$; if that content is out of date, her transcript does not match U .
- Bob detects this inconsistency and outputs FAIL, which then results in a random guess.
- Only when Alice’s simulation is correct does Bob output a meaningful answer.

This completely eliminates the $\mathcal{Q}^*$ ambiguity: Bob’s verification determines whether to use Alice’s simulation or default to a random guess, so there is no need for Alice to identify $\mathcal{Q}^*$. We delegate verification to Bob. We no longer need the Peak-to-Average Lemma, and can achieve the full $\Omega((\log n / \log \log n)^2)$ bound.

Detectable Failure What the verification round buys, technically, is that the players *know when the error occurs*. Bob either accepts, having checked Alice’s entire transcript against the true memory, or declares FAIL; there is no third case in which the protocol is wrong without knowing it. In the one-way setting of [LWY20] no such indicator is available: Alice’s answers on $\mathcal{Q} \setminus \mathcal{Q}^*$ are uncontrolled, and a failure is indistinguishable from a success. What our proof uses is only the resulting decomposition. Conditioned on $\neg \text{FAIL}$ Bob is correct with certainty, and conditioned on FAIL he outputs an independent random guess and contributes nothing, so the advantage is exactly $\Pr[\neg \text{FAIL}]$. This is why advantage over a random guess, rather than error probability, is the quantity we track throughout (Definition 3.1). That the players *know* when the error occurs is not itself used anywhere in the argument; whether this additional structure can be exploited to strengthen the bound we leave as future work.

Our Second Contribution: Handling 2.5-Round Protocols Conceptually, the above simulation is simple. Yet, for the past decade, we have lacked the necessary technical tools to prove lower bounds for such 2.5-round protocols against small nontrivial advantage.

While 3-round communication lower bounds remain beyond current techniques (and would imply major circuit lower bounds [JS10, Dru12, Juk13, DGW19]), our key observation is that the information-theoretic framework introduced by Ko and Weinstein [KW20], further refined by Ko [Ko25a] for the small-advantage regime, can handle exactly this 2.5-round structure. The key is that Bob’s “0.5-round” message Φ is independent of the query Q . This allows us to apply the combinatorial lemmas from [KW20, Ko25a], and we can extract a random variable Z from the 2.5-round protocol such that (a) the min-entropy of S_Q and X conditioned on Z is large, and (b) $I(S_Q; X|Z)$ is small.

1.4.3 The Complete Argument

Our proof proceeds in two steps:

1. **Simulation (Section 3.1):** We show that any dynamic data structure with $w \cdot t_u = n^{o(1)}$ and $t_{tot} = o((\log n / \log(w \cdot t_u))^2)$ can be simulated by a 2.5-round Multiphase Communication Game with nontrivial advantage $n^{-o(1)}$.
2. **Communication Lower Bound (Section 3.2):** We prove that no such 2.5-round protocol can achieve advantage better than $n^{-\Omega(1)}$, yielding a contradiction.

We note that the two main components of our proof — the simulation theorem (Section 3.1) and the communication lower bound (Section 3.2) — are self-contained and can be verified independently.

1.5 Why Is $\tilde{\Omega}(\log^2(n))$ the Barrier?

A natural question arises:

Can this approach be extended to prove higher lower bounds for dynamic Boolean problems, i.e., $\omega(\log^2 n)$?

We argue that our result likely matches the “structural ceiling” of the current techniques, specifically the *Chronogram* framework.

The core of the Chronogram technique relies on decomposing the timeline of updates into a hierarchy of ℓ distinct epochs, where the duration of epoch i grows geometrically. This construction naturally yields a hierarchy of depth $\ell = \Theta(\log n / \log(t_u w))$, which is $\Theta(\log n / \log \log n)$ for $t_u = \text{poly} \log(n)$. The lower bound is then obtained by constructing a hard distribution of updates and queries that forces the data structure to probe $\tilde{\Omega}(\log n)$ distinct cells at *each* epoch. We stress that the advance of this paper is not an exception to this. Our change is to the communication model that the simulation produces (Section 1.4), not to the epoch decomposition; the hierarchy, and hence the factor ℓ , is inherited unchanged from [FS89].

Another perspective on the Chronogram framework is that it reduces dynamic cell-probe lower bounds to static cell-probe lower bounds (with pre-initialized memory and cache [LY25]), amplifying the static bound by a factor of ℓ . As ℓ cannot exceed $\Theta(\log n / \log \log n)$, to show an $\omega(\log^2 n)$ lower bound we must show an $\omega(\log n)$ lower bound for a single epoch $i \in [\ell]$.

Unfortunately, the barriers to static cell-probe lower bounds are well-known. The best known *explicit* static cell-probe lower bound is subject to the well-known logarithmic barrier [MNSW98, Sie04, PD06, PTW08, PTW10, Pă11, KPI25]

$$t \geq \Omega\left(\frac{\log \frac{|\mathcal{Q}|}{n}}{\log \frac{s}{n}}\right) \tag{1}$$

where s is the number of cells used by the static cell-probe data structure, and $|\mathcal{Q}|$ is the number of possible queries. Even for *semi-explicit* static cell-probe lower bounds, where queries may be nonexplicit linear functions, only a lower bound of

$$t \geq \Omega(\min\{\log(|\mathcal{Q}|/s), n/\log s\}) \tag{2}$$

is known, and only against the weaker bit-probe model [Ko25a]. Both (1) and (2) are merely logarithmic in our regime of interest, since both s and $|\mathcal{Q}|$ are polynomial in n . And due to a

well-known connection to circuit lower bounds [JS10, Juk12, Juk13, Dru12, DGW19, KPI25, Ko25a], showing $t \geq \omega(\log n)$ when $|\mathcal{Q}|$ is $\text{poly}(n)$ would be a major circuit complexity breakthrough.

To conclude, any $\omega(\log^2 n)$ Boolean lower bound would require either completely breaking away from the *Chronogram* framework, which has been the main recipe for more than 35 years, or a major circuit complexity breakthrough.

Remark 1.3 (Natural Proofs). *In recent work, Koucký, Loff, Molli and Saks [KLMS26] propose a natural proofs barrier for data structure lower bounds, and argue — conditional on the existence of suitable pseudorandom constructions — that it takes effect at $\tilde{\Omega}(\log^2 n)$ in the dynamic setting. Whether the argument given here is natural in their sense is unclear to us, and we make no claim either way. We observe only that the threshold they identify is the same $\tilde{\Omega}(\log^2 n)$ that the Chronogram framework has long been conjectured to impose, and that the discussion above reaches from a purely structural direction.*

2 Preliminaries

2.1 Information Theory

In this section, we provide the necessary background on information theory that is used in this paper. For comprehensive background, we direct the reader to [CT06]. Unless noted otherwise, all logarithms are base 2.

Definition 2.1 (Entropy). *The entropy of a random variable X is defined as*

$$H(X) := \sum_x \Pr[X = x] \log \frac{1}{\Pr[X = x]}.$$

Similarly, writing $p(x|y) := \Pr[X = x|Y = y]$, the conditional entropy is defined as

$$H(X|Y) := \mathbb{E}_Y \left[\sum_x p(x|y) \log \frac{1}{p(x|y)} \right].$$

We will also apply $H(\cdot)$ to events: for an event E , we write $H(E)$ for the entropy of its indicator random variable $\mathbb{1}[E]$, and likewise for conditional entropies. In particular, $H(E|Y) = 0$ asserts that E is determined by Y .

Fact 2.2 (Conditioning Decreases Entropy). *For any random variables X and Y*

$$H(X) \geq H(X|Y)$$

With entropy defined, we can also quantify the correlation between two random variables, or how much information one random variable conveys about the other.

Definition 2.3 (Mutual Information). *Mutual information between X and Y (conditioned on Z) is defined as*

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

Similarly, we can quantify how far one distribution is from another.

Definition 2.4 (KL-Divergence). *The KL-divergence between two distributions μ and ν is defined as*

$$D(\mu||\nu) := \sum_x \mu(x) \log \frac{\mu(x)}{\nu(x)}.$$

To bound mutual information, it suffices to bound KL-divergence, due to the following fact.

Fact 2.5 (KL-Divergence and Mutual Information). *The following equality between mutual information and KL-divergence holds:*

$$I(A; B|C) = \mathbb{E}_{B,C} [D(A|_{B=b, C=c} || A|_{C=c})].$$

Fact 2.6 (Pinsker's Inequality). *For any two distributions P and Q ,*

$$\|P - Q\|_{TV} = \frac{1}{2} \|P - Q\|_1 \leq \sqrt{\frac{1}{2 \log e} D(P||Q)}$$

We also make use of the following facts on mutual information throughout the paper.

Fact 2.7 (Chain Rule). *For any random variables A, B, C and D*

$$I(AD; B|C) = I(D; B|C) + I(A; B|CD).$$

Fact 2.8. *For any random variables A, B, C and D , if $I(B; D|C) = 0$*

$$I(A; B|C) \leq I(A; B|CD).$$

Proof. By the chain rule and nonnegativity of mutual information,

$$\begin{aligned} I(A; B|C) &\leq I(AD; B|C) = I(B; D|C) + I(A; B|CD) \\ &= I(A; B|CD). \end{aligned}$$

□

Fact 2.9. *For any random variables A, B, C and D , if $I(B; D|AC) = 0$*

$$I(A; B|C) \geq I(A; B|CD).$$

Proof. By the chain rule and nonnegativity of mutual information,

$$\begin{aligned} I(A; B|CD) &\leq I(AD; B|C) = I(A; B|C) + I(B; D|AC) \\ &= I(A; B|C). \end{aligned}$$

□

We will also use the following version of the Chernoff bound.

Fact 2.10 (Chernoff Bound). *Let $X_1, \dots, X_n$ be n independent random variables with $X_i \in \{0, 1\}$ and $\Pr[X_i = 1] = p$ for all $i \in [n]$. Then for any $\varepsilon > 0$ with $\varepsilon < 1 - p$, we have*

$$\Pr \left[\sum_{i \in [n]} X_i \geq (p + \varepsilon)n \right] \leq \exp \left(-\frac{n}{\log_2 e} \cdot D(\text{Ber}(p + \varepsilon) || \text{Ber}(p)) \right)$$

where $\text{Ber}(p)$ refers to the Bernoulli distribution with success probability p .

2.2 Min-Entropy

Definition 2.11. We define the Rényi entropy $H_2(A)$ and min-entropy $H_\infty(A)$ as

$$H_2(A) := -\log \left(\sum_a \Pr[A = a]^2 \right)$$

$$H_\infty(A) := -\log \left(\max_a \Pr[A = a] \right)$$

Fact 2.12 (Rényi Entropy). *Let A be a random variable. Then*

$$H(A) \geq H_2(A) \geq H_\infty(A)$$

In particular, for any fixed b we have

$$H_2(A|B = b) \geq H_\infty(A|B = b)$$

We use the following notion of “average” min-entropy.

Definition 2.13 (Average Min-Entropy).

$$\begin{aligned} \tilde{H}_\infty(A|B) &:= -\log \left(\mathbb{E}_{b \sim B} \left[\max_a \Pr[A = a|B = b] \right] \right) \\ &= -\log \left(\mathbb{E}_{b \sim B} \left[2^{-H_\infty(A|B=b)} \right] \right) \end{aligned}$$

Lemma 2.14 (Lemma 2.2 of [DORS08]). *Let A, B, C be random variables. If B has at most 2^λ possible values, then*

$$\tilde{H}_\infty(A|B, C) \geq \tilde{H}_\infty(A, B|C) - \lambda \geq \tilde{H}_\infty(A|C) - \lambda.$$

In particular,

$$\tilde{H}_\infty(A|B) \geq \tilde{H}_\infty(A, B) - \lambda \geq H_\infty(A) - \lambda.$$

Claim 2.15.

$$\tilde{H}_\infty(A|B, C) \leq \tilde{H}_\infty(A|B)$$

Claim 2.16 ([JRS09]). *Let $\mathcal{D}$ be a distribution and $\mathcal{U}$ the uniform distribution over some $\mathcal{S} \subseteq \{0, 1\}^n$. If $D(\mathcal{D}||\mathcal{U}) \leq t$ with $t > 1$, then for every $\alpha \geq 4$ there exists an event E such that*

$$\begin{aligned} \mathcal{D}(E) &\geq 1 - \frac{1}{\alpha} \\ H_\infty(\mathcal{D}|_E) &\geq \log |\mathcal{S}| - 2\alpha t \end{aligned}$$

We include proofs of Claim 2.15 and Claim 2.16 in Appendix A for completeness.

3 Main Proof

In this section, we prove our main theorem.

Theorem 1.1. *Let f be Inner Product over $\mathbb{F}_2$ and let $m = n^{1+\Omega(1)}$. Then any data structure for the Multiphase Problem for f whose update time satisfies $t_u = n^{o(1)}$ must have*

$$t_{tot} \geq \Omega \left(\left(\frac{\log n}{\log(t_u w)} \right)^2 \right).$$

In particular,

$$\max\{t_u, t_{tot}\} \geq \Omega \left(\left(\frac{\log n}{\log \log n} \right)^2 \right).$$

Before we delve into the proof of the main theorem, we introduce formal notation and the model necessary for the proof.

Epochs To give a super-logarithmic lower bound, the first main technical component is the Chronogram technique [FS89], where we divide the sequence of n updates into epochs, epoch i containing $n_i := \beta^i$ updates. We take ℓ to be the largest integer with $\sum_{i=1}^{\ell} n_i \leq n$, appending any remaining updates to the last epoch, so that all n updates are covered. By maximality of ℓ , $n < \sum_{i=1}^{\ell+1} \beta^i \leq \frac{\beta}{\beta-1} \beta^{\ell+1} \leq 2\beta^{\ell+1}$ for $\beta \geq 2$, which gives

$$\beta^{\ell+1} > \frac{n}{2}, \quad \ell = \Theta(\log_{\beta} n) = \Theta\left(\frac{\log n}{\log \beta}\right). \quad (3)$$

For each epoch of n_i updates, denoted as $X^{(i)}$, the update algorithms will use at most $t_u n_i$ cell-probe operations (write operations) to overwrite the memory. We denote the epochs of updates as $\{U_i\}_{i=1}^{\ell}$, each of size $|U_i| \leq 2w \cdot t_u n_i$ (an address and a content of w bits for each written cell), and we process the updates in reverse order, that is, the larger epochs get processed first.

The essence of the Chronogram technique is that every query algorithm can be decomposed in the following manner. Let C_i denote the cells last updated in epoch i . The query algorithm for the query q will then probe into $C_1, \dots, C_{\ell}$. We denote by $t_q[i]$ the number of probes the query for $q \in \mathcal{Q}$ makes into C_i . Then it must be the case that

$$\sum_{i=1}^{\ell} t_q[i] \leq t_{tot}$$

Communication Model We define the following 2.5-round Multiphase Communication Game, as an extension of the 1.5-round Multiphase Communication Game introduced in [CEEP16]. This replaces the one-way communication model used in [LWY20, LY25] for the lower bound. One difference from Protocol 2 is worth noting: there the query is known to both players, whereas here only Alice receives it. Our model sits between two extremes. Pătraşcu’s model [Pat10] has so far resisted every attempt at a lower bound, whereas the restricted models in which lower bounds have been obtained [CEEP16, KW20] do not imply lower bounds for general dynamic data structures. The 2.5-round model is weaker than Pătraşcu’s, yet strong enough to yield such an implication, and it is what the simulation of Section 3.1 actually produces.

- Alice holds $\vec{S}$. Bob holds X . Merlin holds $\vec{S}$ and X .
- Merlin sends U to Bob.
- As the 0.5-round message, Bob simply forwards to Alice a sub-sampled partial view Φ of the updated memory. It may depend on U , on Bob's own input X and on his private randomness, but never on the query.
- The query $q \in \mathcal{Q}$ is then revealed to Alice. Alice sends a message A_q to Bob, and Bob outputs an estimate $B_q \in \{\pm 1\}$ of $f(S_q, X)$ as a function of U , A_q and his own randomness.

Protocol 3: 2.5-Round Multiphase Communication Game

Advantage of the Protocol We will be considering a fixed $\vec{S}$ for the simulation theorem, while for the lower bound analysis we will assume that these $\vec{S}$ are distributed uniformly at random. Regardless of the chosen $\vec{S}$, we will be using a uniform distribution over X as the update, and a uniform distribution over $\mathcal{Q}$ as the possible query. We do not require f to be balanced: the only property of f used in the lower bound is that it has small bias under product distributions of large min-entropy (Section 3.2; Definition B.1 for general f). This already accounts for rare biased inputs, such as $S_Q = 0$ for Inner Product over $\mathbb{F}_2$. Under these distributions, we measure the advantage of the 2.5-round protocol Π . Throughout the rest of the paper we abbreviate the messages exchanged on query q as

$$\mathcal{M}_q := (\Phi, A_q, B_q),$$

that is the whole 2.5-round transcript conversed between Alice and Bob.

Definition 3.1. Fix $\vec{S}$. The advantage of a 2.5-round protocol $\Pi = (\Phi, A_Q, B_Q)$ under the distribution Q over the queries is defined as

$$\text{adv}_{\vec{S}}(G_f, \Pi) := \mathbb{E}_{Q, \mathcal{M}_Q} \left[\mathbb{E}_{(S_Q, X) | Q, \mathcal{M}_Q} [B_Q \cdot f(S_Q, X)] \right]$$

where $B_Q \in \{\pm 1\}$ is Bob's final estimate.

For the rest of the section, we prove the simulation theorem (Section 3.1) showing that any efficient data structure yields an efficient 2.5-round protocol achieving good advantage. This results in a contradiction via information-theoretic arguments (Section 3.2).

3.1 Simulation Theorem

First, we show the following simulation theorem, which argues that an efficient dynamic data structure implies an efficient 2.5-round Multiphase Communication Game protocol Π with large advantage. Throughout this section we write

$$\xi := \frac{100 t_{\text{tot}}}{\ell} \ln(1/p) \quad (4)$$

for the cost, in the exponent, of the event on which the simulation succeeds: Bob's sample $\tilde{C}_i$ containing every cell of C_i that the query probes.

Theorem 3.2 (New Simulation Theorem). Fix $\vec{S}$, put $\beta := (t_u w)^2$, and let $\ell = \Theta(\log n / \log \beta)$ be the resulting number of epochs as in (3). If there exists a data structure for the Multiphase Problem

for f with query time t_{tot} and update time t_u , then for every $p \in (0, 1/3)$ with $p \leq 1/(8t_u w)$ there exists a 2.5-round Multiphase Communication Game protocol Π with

$$\begin{aligned} |U| &\leq 2wt_u n + \log \ell, \\ |\Phi| &\leq n - \frac{1}{2} (1 - 4pwt_u - o(1)) n^{2/3}, \\ |A_q| &\leq O(w t_{tot}), \end{aligned}$$

achieving advantage

$$\text{adv}_{\vec{S}}(G_f, \Pi) \geq \Omega \left(e^{-\xi} - e^{-\Omega(p n^{2/3})} \right).$$

Proof. Suppose there exists a dynamic data structure for the Multiphase Problem for f with t_{tot} query time and t_u update time. We simulate the data structure with a 2.5-round Multiphase Communication Game protocol as follows.

The cost guarantee of the protocol is straightforward. Merlin's message is simply the update together with an index $i \in [2(\ell + 1)/3, \ell]$. Therefore $|U| \leq 2t_u w n + \log \ell = O(t_u w n)$. Each cell in Φ is sent as an address and a content, $2w$ bits. Next, we bound the length of Bob's 0.5-round message:

$$\begin{aligned} |\Phi| &= \left| X^{(>i)} \right| + |C_{<i}| + \left| \tilde{C}_i \right| \\ &\leq \sum_{k>i} n_k + \sum_{k<i} (2wt_u) n_k + 4pwt_u n_i \\ &\leq (n - n_i) + 4pwt_u n_i + o(n_i) \end{aligned} \tag{5}$$

where the upper bound follows from the guarantee on the epoch sizes, $n_i := \beta^i = (t_u w)^{2i}$. Here $\sum_{k<i} (2wt_u) \cdot n_k \leq 4wt_u n_{i-1} \leq 4wt_u n_i / \beta = o(n_i)$ as it is a geometric series and $\beta = (t_u w)^2$. As Merlin chooses $i \in [2(\ell + 1)/3, \ell]$ and $\beta^{\ell+1} > n/2$ by (3), we get $n_i \geq \beta^i \geq \beta^{2(\ell+1)/3} > (n/2)^{2/3} \geq n^{2/3}/2$ (with $\geq$ rather than $=$, as the leftover updates enlarge the last epoch). Therefore, assuming $1 - 4p \cdot wt_u - o(1) > 0$, as we will choose the parameters accordingly, (5) becomes

$$\begin{aligned} (5) &\leq n - (1 - 4pwt_u - o(1)) n_i \\ &\leq n - \frac{1}{2} (1 - 4pwt_u - o(1)) n^{2/3}. \end{aligned} \tag{6}$$

Next, we bound the length of Alice's message,

$$|A_q| = 2w \cdot t_{tot} + 1 \leq 3w \cdot t_{tot}. \tag{7}$$

(6) and (7) give the desired bound on the cost of Π . We now lower bound $\text{adv}_{\vec{S}}(G_f, \Pi)$ by bounding the probability that Alice's simulation is consistent with U , so that Bob's verification accepts.

Claim 3.3.

$$\text{adv}_{\vec{S}}(G_f, \Pi) \geq \Omega \left(e^{-\xi} - e^{-\Omega(p n^{2/3})} \right)$$

Proof. Recall that $\vec{S}$ is fixed and that Merlin chose $i \in [2(\ell + 1)/3, \ell]$ with

$$\mathbb{E}_{\mathcal{Q}|\vec{S}, X} [t_q[i]] \leq 10 \cdot t_{tot} / \ell.$$

We first note the observation that reduces the analysis to a probability estimate. If Bob's verification accepts, then $B_q = f(S_q, X)$; if it rejects, B_q is an independent random guess. The advantage on a

- First Merlin chooses $i \in [2(\ell + 1)/3, \ell]$, which satisfies

$$\mathbb{E}_{q \in \mathcal{Q} | \vec{S}, X} [t_q[i]] \leq 10 \cdot t_{tot}/\ell.$$

Note that by a simple averaging argument, such i is guaranteed to exist, as the interval contains at least $(\ell - 2)/3 \geq \ell/10$ integers for $\ell \geq 3$. Merlin then forwards $U = U_1, \dots, U_\ell$ along with the index i to Bob. From U , Bob can construct $C_1, \dots, C_\ell$.

- As the 0.5-round message, Bob sends the following. For C_i , he samples each cell with probability p at random, and creates $\tilde{C}_i$. If $\tilde{C}_i$ contains more than $2p \cdot t_u n_i$ cells, then Bob simply aborts, and outputs FAIL. Otherwise, he forwards $X^{(>i)}, C_{<i}$ and $\tilde{C}_i$ along with i (which is evident in the message already) to Alice. Observe that the message depends only on U, X, i and Bob's randomness, and not on q .
- Alice does the following.

- Using $\vec{S}$, and $X^{(>i)}$, the first $\ell - i$ epochs of updates, Alice constructs $U_{>i}$ from $X^{(>i)}$.
- Alice answers a probe to an address a as follows. If $a \in \tilde{C}_i \cup C_{<i}$, she uses the content supplied in Φ . Otherwise she uses the content given by her own reconstruction from $\vec{S}$ and $X^{(>i)}$, which is correct whenever $a \notin C_i$. Note that in the latter case she cannot tell whether a in fact lies in $C_i \setminus \tilde{C}_i$, in which case her content is wrong; this is precisely the ambiguity that Bob's verification resolves.
- Alice simulates the query algorithm for $q \in \mathcal{Q}$ under $\vec{S}$, generating the query transcript

$$A_q := ((a_1, c_1), \dots, (a_{t_{tot}}, c_{t_{tot}}), z_{out})$$

where a_τ refers to the address probed at time τ , c_τ refers to its content, and z_{out} refers to the final output of the query algorithm.

- Alice sends the query transcript A_q to Bob.
- Then Bob simply verifies A_q . That is, he checks whether (a_τ, c_τ) matches U for every $\tau \in [t_{tot}]$. (Note that if the cell a_τ is never updated by U , Alice's content c_τ must be correct, as Alice holds $\vec{S}$.) If every probe matches, he declares z_{out} as the answer B_q . Otherwise, he outputs FAIL.
- Whenever Bob aborts or the verification fails, he outputs FAIL and an independent random guess as B_q .

Protocol 4: Simulation Protocol II

query q is therefore exactly $\Pr[\text{verification accepts}]$, and in particular it is never negative. It suffices, then, to exhibit a constant fraction of queries on which acceptance is likely.

Step 1: discarding the heavy queries. By Markov's inequality, at most a $1/10$ fraction of the queries have $t_q[i] > 100 t_{tot}/\ell$. We claim nothing on those; by the observation above they contribute a nonnegative amount. Fix a query q with $t_q[i] \leq 100 t_{tot}/\ell$.

Step 2: when does the verification accept? Consider the execution of the query algorithm for q on the actual memory, and let $P_q \subseteq C_i$ be the set of $t_q[i]$ cells of C_i that it probes. Note that P_q is determined by $\vec{S}, X$ and i alone; in particular it does not depend on Bob's sampling. We claim that the verification accepts whenever $P_q \subseteq \tilde{C}_i$. Indeed, every content Alice then supplies agrees

with U : an address in $\tilde{C}_i \cup C_{<i}$ is answered from Φ , and any other address lies outside C_i , where her reconstruction from $\vec{S}$ and $X^{(>i)}$ is correct. Her simulation therefore reproduces the actual execution probe by probe, so Bob accepts.

Step 3: the sampling. Bob samples each cell of C_i independently with probability p , and aborts only when $\tilde{C}_i$ contains more than $2p \cdot t_u n_i$ cells. Since $|P_q| \leq 100 t_{tot}/\ell$,

$$\begin{aligned} \Pr[\text{accepts}] &\geq \Pr[P_q \subseteq \tilde{C}_i] - \Pr[\text{aborts}] \\ &\geq p^{100 t_{tot}/\ell} - e^{-\Omega(n_i D(\text{Ber}(2p) || \text{Ber}(p)))} \\ &\geq e^{-\xi} - e^{-\Omega(p n^{2/3})}, \end{aligned}$$

where the last step uses (4), the standard estimate $D(\text{Ber}(2p) || \text{Ber}(p)) = \Theta(p)$, and $n_i \geq n^{2/3}/2$. The abort bound is a Chernoff bound over the at most $t_u n_i$ cells of C_i .

Averaging over the at least 9/10 fraction of queries covered by Steps 2 and 3 gives

$$\text{adv}_{\vec{S}}(G_f, \Pi) \geq \Omega \left(e^{-\xi} - e^{-\Omega(p n^{2/3})} \right) \quad (8)$$

which then completes the proof of the claim. $\square$

(6), (7) and Claim 3.3 complete the proof of the simulation theorem. $\square$

Remark 3.4 (Detected Error). *Protocol 4 errs only on an event the players detect: whenever Bob does not declare FAIL, he has verified Alice's transcript against U in full, and therefore $B_q = f(S_q, X)$ with certainty. Every loss in $\text{adv}_{\vec{S}}(G_f, \Pi)$ thus comes from $\Pr[\text{FAIL}]$, on which Bob outputs an independent random guess and contributes zero advantage. Consequently $\text{adv}_{\vec{S}}(G_f, \Pi) = \Pr[\neg \text{FAIL}]$, and the lower bound of Section 3.2 need only bound this single quantity, in contrast to the one-way analyses of [Lar12, LWY20].*

Now we would like to select a range of parameters to derive a contradiction while simplifying the expressions.

Corollary 3.5. *Assume $t_u = n^{o(1)}$ and set $t_{tot} = o \left(\left(\frac{\log n}{\log(t_u w)} \right)^2 \right)$, $\beta = (w \cdot t_u)^2$, $p := \beta^{-1}/2$. Then*

$$|U| \leq n^{1+o(1)}, \quad |\Phi| \leq n - \frac{n^{2/3}}{4}, \quad |A_q| \leq o(\log^3(n))$$

while having

$$\text{adv}_{\vec{S}}(G_f, \Pi) \geq n^{-o(1)}.$$

Proof. Let ℓ be as in (3) with $\beta = (t_u w)^2$, and let the sampling probability be $p := \beta^{-1}/2 = 1/(2(t_u w)^2)$, so that $p \leq 1/(8t_u w)$ and $4p w t_u = 2/(t_u w) = o(1)$. As $w = \Theta(\log n)$, the hypothesis $t_u = n^{o(1)}$ gives $t_u w = n^{o(1)}$ and hence $\log \beta = o(\log n)$, so that $\ell = \omega(1)$ and $\ell = \Theta(\log n / \log(t_u w))$. The cost bounds are then

$$\begin{aligned} |U| &\leq O(w n t_u) \leq n^{1+o(1)} \\ |\Phi| &\leq n - \frac{1}{2}(1 - o(1)) n^{2/3} \leq n - \frac{n^{2/3}}{4} \\ |A_q| &\leq O(t_{tot} \cdot w) \leq o(\log^3 n), \end{aligned}$$

where the last bound uses $\log(t_u w) \geq \log w = \Omega(\log \log n)$, whence $t_{tot} = o((\log n / \log \log n)^2)$, together with $w = \Theta(\log n)$. For the advantage, $\ln(1/p) = \ln(2\beta) = \Theta(\log(t_u w))$, so that

$$\xi = \Theta \left(\frac{t_{tot} (\log(t_u w))^2}{\log n} \right) = o(\log n)$$

while $p n^{2/3} = n^{2/3-o(1)}$, so that (8) is at least $\Omega(n^{-o(1)}) - e^{-n^{2/3-o(1)}} \geq n^{-o(1)}$. $\square$

3.2 2.5-Round Multiphase Communication Game Lower Bound

The goal of this section is to show that such an advantage is impossible, building on the framework of [KW20, Ko25a, Ko25b]. The key idea, developed across [KW20, Ko25a, Ko25b], is to extract a random variable Z containing the output Z_{out} from the 2.5-round protocol such that (a) the min-entropy of S_Q and X conditioned on Z is large, and (b) $I(S_Q; X|Z)$ is small. Then it must be the case that $\mathbb{E}_{(S_Q, X)|Z} [Z_{out} \cdot f(S_Q, X)]$ is small. Throughout this section we abbreviate the pair of protocol messages — Bob's 0.5-round message and Alice's messages **for all queries** — as the *transcript*

$$\mathcal{T} := (\Phi, A_{\mathcal{Q}_{all}}),$$

and we follow the analogous argument, selecting Z as

$$Z := (\mathcal{T}, S_{<Q}, B_Q, Q)$$

where Q is chosen uniformly at random from $\mathcal{Q} = [m]$ with $m = n^{1+\Omega(1)}$. $A_{\mathcal{Q}_{all}}$ refers to A_q for all $q \in \mathcal{Q}$, and $S_{<Q}$ refers to $S_1, \dots, S_{Q-1}$. Note that B_Q plays the role of the output Z_{out} . We need to include A_q for all $q \in \mathcal{Q}$ to apply direct-sum techniques.

We first show that Corollary 3.5 implies the following claim on the advantage conditioned on Z .

Claim 3.6.

$$\mathbb{E}_Z \left[\left| \mathbb{E}_{(S_Q, X)|Z} [\chi_{S_Q}(X)] \right| \right] \geq n^{-o(1)}$$

where $\chi_{S_Q}(X) := (-1)^{\langle S_Q, X \rangle}$.

Proof. As Q is chosen uniformly at random, Corollary 3.5 guarantees that for any fixed $\vec{S}$,

$$\mathbb{E}_{\mathcal{M}_Q, Q|\vec{S}} [B_Q \cdot \chi_{S_Q}(X)] \geq n^{-o(1)}$$

This then implies

$$\begin{aligned} n^{-o(1)} &\leq \mathbb{E}_{\mathcal{M}_Q, Q, \vec{S}} \left[\mathbb{E}_{X|\mathcal{M}_Q, Q, \vec{S}} [B_Q \cdot \chi_{S_Q}(X)] \right] \\ &= \mathbb{E}_{\mathcal{M}_Q, Q, S_{<Q}, S_Q} \left[\mathbb{E}_{X|\mathcal{M}_Q, Q, S_{<Q}, S_Q} [B_Q \cdot \chi_{S_Q}(X)] \right] \\ &= \mathbb{E}_{\mathcal{M}_Q, Q, S_{<Q}} \left[\mathbb{E}_{(S_Q, X)|\mathcal{M}_Q, Q, S_{<Q}} [B_Q \cdot \chi_{S_Q}(X)] \right] \\ &\leq \mathbb{E}_{\mathcal{M}_Q, Q, S_{<Q}} \left[\left| \mathbb{E}_{(S_Q, X)|\mathcal{M}_Q, Q, S_{<Q}} [\chi_{S_Q}(X)] \right| \right] \end{aligned}$$

But then

$$\begin{aligned} n^{-o(1)} &\leq \mathbb{E}_{\mathcal{M}_Q, Q, S_{<Q}} \left[\left| \mathbb{E}_{(S_Q, X)|\mathcal{M}_Q, Q, S_{<Q}} [\chi_{S_Q}(X)] \right| \right] \\ &\leq \mathbb{E}_{\mathcal{T}, B_Q, Q, S_{<Q}} \left[\left| \mathbb{E}_{(S_Q, X)|\mathcal{T}, B_Q, Q, S_{<Q}} [\chi_{S_Q}(X)] \right| \right] \\ &= \mathbb{E}_Z \left[\left| \mathbb{E}_{(S_Q, X)|Z} [\chi_{S_Q}(X)] \right| \right] \end{aligned}$$

where the inequality follows from Jensen's inequality on the convex function, $|\cdot|$, and since $(\mathcal{T}, B_Q, Q)$ determines $\mathcal{M}_Q$. $\square$

We now show that Claim 3.6 leads to a contradiction.

3.2.1 Small Information on S_Q

First, we would like to show that the min-entropy of S_Q conditioned on Z is large. The statement is not true in general. However, a simple Markov argument allows us to “extract” a good event E_Q . We start with the following standard claim on the mutual information.

Claim 3.7.

$$I(S_Q; \mathcal{T}, S_{<Q}, Q) \leq 4w \cdot t_{tot}$$

Proof. By a standard direct-sum argument, since Q is chosen independently at random and $S_1, \dots, S_m$ are i.i.d., we have

$$\begin{aligned} I(S_Q; \mathcal{T}, S_{<Q}, Q) &= \underbrace{I(S_Q; S_{<Q}, Q)}_{=0} + I(S_Q; \mathcal{T} | S_{<Q}, Q) \\ &= \mathbb{E}_Q [I(S_Q; \mathcal{T} | S_{<Q})] \end{aligned}$$

Then we can bound $\mathbb{E}_Q [I(S_Q; \mathcal{T} | S_{<Q})]$ as

$$\begin{aligned} \mathbb{E}_Q [I(S_Q; \mathcal{T} | S_{<Q})] &= \frac{1}{m} \cdot I(\vec{S}; \mathcal{T}) \\ &\leq \underbrace{\frac{|\Phi|}{m}}_{\leq \frac{n}{m} = n^{-\Omega(1)} = o(1)} + \frac{m |A_q|}{m} \\ &\leq o(1) + 3w \cdot t_{tot} \leq 4w \cdot t_{tot} \end{aligned}$$

where the first inequality follows from $I(\vec{S}; \mathcal{T}) \leq H(\mathcal{T}) \leq |\Phi| + m |A_q|$, completing the proof of the claim. $\square$

Unfortunately, bounds on the mutual information (equivalently, on the KL-divergence) alone do not suffice for our proof. Since we are considering the “small” advantage regime, we need to argue that the min-entropy of S_Q conditioned on Z is large, as in [Ko25a]. We will use the following lemma to extract a “good” event in terms of min-entropy.

Lemma 3.8 (S_Q Min-Entropy). *For every setting of the parameter $\alpha \in (0, \frac{1}{10})$, there exist events E_Q^1 and E_Q^2 such that*

$$H(E_Q^1 | \mathcal{T}, S_{<Q}, Q) = H(E_Q^2 | \mathcal{T}, S_{<Q}, S_Q, Q) = 0$$

and

$$\begin{aligned} \Pr [E_Q^1] &\geq 1 - \alpha, \quad \Pr [E_Q^2 | \mathcal{T}, S_{<Q}, Q, E_Q^1=1] \geq 1 - \alpha \\ H_\infty(S_Q | \mathcal{T}, S_{<Q}, Q, E_Q=1) &\geq n - \frac{8w \cdot t_{tot}}{\alpha^2} \end{aligned}$$

where we denote $E_Q^1=1, E_Q^2=1$ as $E_Q=1$ for short. The min-entropy bound holds for every fixing of $(\mathcal{T}, S_{<Q}, Q)$ consistent with $E_Q=1$; in particular it also holds for the average min-entropy $\tilde{H}_\infty(S_Q | \mathcal{T}, S_{<Q}, Q, E_Q=1)$, which is how we use it in Lemma 3.12.

Proof. First, we define the event E_Q^1 as the set of $(\mathcal{T}, S_{<Q}, Q)$ such that

$$D(S_Q | \mathcal{T}, S_{<Q}, Q || S_Q) \leq \frac{4w \cdot t_{tot}}{\alpha}$$

Due to Claim 3.7,

$$\begin{aligned} I(S_Q; \mathcal{T}, S_{<Q}, Q) &= \mathbb{E}_{\mathcal{T}, S_{<Q}, Q} [D(S_Q | \mathcal{T}, S_{<Q}, Q || S_Q)] \\ &\leq 4w \cdot t_{tot} \end{aligned}$$

Markov's inequality implies $\Pr[E_Q^1] \geq 1 - \alpha$. Conditioned on $E_Q^1=1$, we obtain the event E_Q^2 from Claim 2.16. As $E_Q^1=1$ guarantees

$$D(S_Q | \mathcal{T}, S_{<Q}, Q, E_Q^1=1 || S_Q) = D(S_Q | \mathcal{T}, S_{<Q}, Q || S_Q) \leq \frac{4w \cdot t_{tot}}{\alpha}$$

where the equality holds as E_Q^1 is a deterministic function of $\mathcal{T}, S_{<Q}, Q$, Claim 2.16, applied with $t := 4w \cdot t_{tot}/\alpha$ and with $1/\alpha > 10$ in place of α , yields an event E_Q^2 such that

$$\begin{aligned} \Pr[E_Q^2 | \mathcal{T}, S_{<Q}, Q, E_Q^1=1] &\geq 1 - \alpha \\ H(E_Q^2 | S_Q, \mathcal{T}, S_{<Q}, Q, E_Q^1=1) &= 0 \\ H_\infty(S_Q | \mathcal{T}, S_{<Q}, Q, E_Q^1=1, E_Q^2=1) &\geq n - \frac{8w \cdot t_{tot}}{\alpha^2} \end{aligned}$$

which completes the proof of the lemma. $\square$

3.2.2 Small Information on X

Next, we show that $\tilde{H}_\infty(X | Z, E_Q=1)$ is large. We remark that the conditioning on $E_Q=1$ is needed to derive the contradiction.

Lemma 3.9.

$$\tilde{H}_\infty(X | Z, E_Q=1) \geq n - |\Phi| - 10\alpha - 1$$

Proof. Recall that as X is chosen uniformly at random, regardless of the choice of Q and $\vec{S}$,

$$H_\infty(X | \vec{S}, Q) = n$$

while Lemma 2.14 implies that

$$\tilde{H}_\infty(X | \vec{S}, \Phi, Q) \geq n - |\Phi|.$$

But $\vec{S}, \Phi, Q$ and Alice's independent private randomness determine $A_{\mathcal{Q}_{all}}$. That is, $A_{\mathcal{Q}_{all}}$ is determined by Alice's private randomness and $\vec{S}, \Phi$. Therefore,

$$I(X; A_{\mathcal{Q}_{all}} | \vec{S}, \Phi, Q) = 0.$$

Moreover, $(\vec{S}, \mathcal{T}, Q)$ contains $(S_Q, \mathcal{T}, S_{<Q}, Q)$, which determines both E_Q^1 and E_Q^2 . Therefore,

$$\tilde{H}_\infty(X | \vec{S}, \mathcal{T}, E_Q, Q) \geq n - |\Phi|$$

Then by the definition of average min-entropy we have the following equality.

$$\begin{aligned} 2^{-\tilde{H}_\infty(X|\vec{S}, \mathcal{T}, E_Q, Q)} &= \Pr[E_Q=1] \cdot 2^{-\tilde{H}_\infty(X|\vec{S}, \mathcal{T}, E_Q=1, Q)} \\ &\quad + \Pr[E_Q=0] \cdot 2^{-\tilde{H}_\infty(X|\vec{S}, \mathcal{T}, E_Q=0, Q)} \end{aligned}$$

which implies the following inequality

$$\tilde{H}_\infty(X|\vec{S}, \mathcal{T}, E_Q=1, Q) \geq \log(\Pr[E_Q=1]) + \tilde{H}_\infty(X|\vec{S}, \mathcal{T}, E_Q, Q) \quad (9)$$

The left-hand side of (9) is upper bounded by

$$\tilde{H}_\infty(X|\vec{S}, \mathcal{T}, E_Q=1, Q) \leq \tilde{H}_\infty(X|S_{<Q}, \mathcal{T}, E_Q=1, Q)$$

due to Claim 2.15.

Lemma 3.8 guarantees that $\Pr[E_Q=1] \geq (1 - \alpha)^2 \geq 1 - 2\alpha$. Then the right-hand side of (9) is lower bounded by

$$\begin{aligned} \log(\Pr[E_Q=1]) + \tilde{H}_\infty(X|\vec{S}, \mathcal{T}, E_Q, Q) &\geq \log(1 - 2\alpha) + n - |\Phi| \\ &\geq n - |\Phi| - 10\alpha, \end{aligned}$$

where the last lower bound follows from $\alpha \in (0, 0.1)$, and $\log(1 - 2\alpha) \geq -10\alpha$ in this interval.

Combining the two bounds, we obtain

$$\tilde{H}_\infty(X|S_{<Q}, \mathcal{T}, E_Q=1, Q) \geq n - |\Phi| - 10\alpha. \quad (10)$$

Again applying Lemma 2.14 to (10), we get

$$\tilde{H}_\infty(X|S_{<Q}, \mathcal{T}, B_Q, E_Q=1, Q) \geq n - |\Phi| - 10\alpha - 1,$$

which completes the proof of the lemma. $\square$

3.2.3 Small Correlation between S_Q and X

Finally, we show that S_Q and X are weakly correlated after conditioning on Z and $E_Q=1$. We show the following lemma.

Lemma 3.10.

$$I(S_Q; X|Z, E_Q=1) \leq \frac{1}{(1 - \alpha)^2} \frac{|U|}{m}$$

Proof. We start by substituting our choice of Z into $I(S_Q; X|Z, E_Q=1)$:

$$\begin{aligned} I(S_Q; X|Z, E_Q=1) &\leq I(S_Q; XU|S_{<Q}, \mathcal{T}, B_Q, E_Q=1, Q) \\ &\leq I(S_Q; XU|S_{<Q}, \mathcal{T}, E_Q=1, Q) \end{aligned}$$

where the last inequality follows from Fact 2.9 with

$$I(S_Q; B_Q|XU, S_{<Q}, \mathcal{T}, E_Q=1, Q) = 0,$$

as $U, A_{\mathcal{Q}_{all}}$ and Bob's independent randomness fully determine B_Q . That is, depending on U , either the final output of A_Q is chosen as B_Q , or B_Q is uniformly random in $\{\pm 1\}$. Then we can further upper bound the term of interest as

$$\begin{aligned} I(S_Q; XU|S_{<Q}, \mathcal{T}, E_Q=1, Q) &= I(S_Q; XU|S_{<Q}, \mathcal{T}, E_Q^1=1, E_Q^2=1, Q) \\ &\leq \frac{I(S_Q; XU|S_{<Q}, \mathcal{T}, E_Q^1=1, E_Q^2, Q)}{\Pr[E_Q^2=1|S_{<Q}, \mathcal{T}, E_Q^1=1, Q]} \\ &\leq \frac{I(S_Q; XU|S_{<Q}, \mathcal{T}, E_Q^1=1, Q)}{1 - \alpha} \end{aligned}$$

where the last inequality follows from Fact 2.9 with

$$I(E_Q^2; XU|S_Q, S_{<Q}, \mathcal{T}, E_Q^1=1, Q) \leq H(E_Q^2|S_Q, S_{<Q}, \mathcal{T}, E_Q^1=1, Q) = 0$$

from Lemma 3.8. Now E_Q^1 is fully determined by $S_{<Q}, \mathcal{T}, Q$, and is 1 with probability at least $1 - \alpha$. Therefore

$$I(S_Q; XU|S_{<Q}, \mathcal{T}, E_Q^1=1, Q) \leq \frac{I(S_Q; XU|S_{<Q}, \mathcal{T}, Q)}{1 - \alpha}.$$

Finally, using the standard direct-sum technique and the chain rule for mutual information (Fact 2.7), we get

$$\begin{aligned} I(S_Q; XU|S_{<Q}, \mathcal{T}, Q) &= \frac{1}{m} \sum_{q=1}^m I(S_q; XU|S_{<q}, \mathcal{T}) = \frac{I(\vec{S}; XU|\mathcal{T})}{m} \\ &\leq \frac{I(\vec{S}; XU|\Phi)}{m} \leq \frac{I(\vec{S}; XU)}{m} \\ &= \frac{I(\vec{S}; X) + I(\vec{S}; U|X)}{m} \leq \frac{|U|}{m} \end{aligned}$$

where the first inequality follows from Fact 2.9 with $I(A_{\mathcal{Q}_{all}}; XU|\vec{S}, \Phi) = 0$. This is true as $\vec{S}, \Phi$ and Alice's independent randomness determine Alice's message $A_{\mathcal{Q}_{all}}$. The second inequality holds again from Fact 2.9 with $I(\Phi; \vec{S}|XU) = 0$ as Φ is fully determined by U, X and Bob's independent randomness. The final inequality holds from $I(\vec{S}; X) = 0$ and $I(\vec{S}; U|X) \leq |U|$.

Combining the inequalities, we obtain the desired bound

$$I(S_Q; X|Z, E_Q=1) \leq \frac{1}{(1 - \alpha)^2} \frac{|U|}{m}.$$

□

3.2.4 The Main Contradiction

We will use the following combinatorial lemma from [Ko25a]. In its statement, and in its proof in Appendix A, we abbreviate the conditioning on a fixed pair (a, z) as $\mathcal{E} := (A = a, Z = z)$.

Lemma 3.11 ([Ko25a]). *Let X and Y be distributions over $\{0, 1\}^n$. For any $\gamma \geq 3$, if a random process $Z = z$ and A , which contains z_{out} , satisfy the inequalities*

$$\begin{aligned} \tilde{H}_\infty(Y|A, Z = z) + \tilde{H}_\infty(X|A, Z = z) &\geq n + 2\gamma \\ I(Y; X|A, Z = z) &\leq 2^{-2\gamma}, \end{aligned}$$

then it must be the case that

$$\mathbb{E}_{A|Z=z} [|\mathbb{E}_{Y,X|\mathcal{E}} [\chi_Y(X) \cdot z_{out}]|] < 2^{-\gamma+2}$$

We include the full proof of Lemma 3.11 in Appendix A for completeness. The following lemma completes the proof of Theorem 1.1.

Lemma 3.12. *If $m = n^{1+\Omega(1)}$, then no 2.5-round Multiphase Communication Game satisfying the parameters of Corollary 3.5 can exist.*

To see how Lemma 3.12 completes the proof of Theorem 1.1, note that if $t_{tot} = o((\log n / \log(t_u w))^2)$, then Corollary 3.5 yields a 2.5-round protocol with the stated parameters, whereas Lemma 3.12 shows that no such protocol exists. This gives the claimed lower bound on t_{tot} .

Proof of Lemma 3.12. Recall that Corollary 3.5 implies Claim 3.6, i.e.,

$$\mathbb{E}_Z \left[\left| \mathbb{E}_{(S_Q, X)|Z} [\chi_{S_Q}(X)] \right| \right] \geq n^{-o(1)}$$

with the parameters

$$\begin{aligned} |U| &\leq n^{1+o(1)}, \quad |\Phi| \leq n - \frac{n^{2/3}}{4}, \\ |A_q| &= \Theta(w \cdot t_{tot}) \leq o(\log^3(n)). \end{aligned}$$

On the other hand, from Lemma 3.8, Lemma 3.9, and Lemma 3.10, we get the following set of inequalities:

$$\tilde{H}_\infty(X|Z, E_Q=1) \geq n - |\Phi| - 10\alpha - 1 \quad (11)$$

$$\begin{aligned} \tilde{H}_\infty(S_Q|Z, E_Q=1) &\geq n - \frac{8w \cdot t_{tot}}{\alpha^2} - 1 \\ &= n - \Theta\left(\frac{|A_q|}{\alpha^2}\right) - 1 \end{aligned} \quad (12)$$

$$I(S_Q; X|Z, E_Q=1) \leq \frac{1}{(1-\alpha)^2} \frac{|U|}{m} \quad (13)$$

where (12) simply follows from

$$\begin{aligned} \tilde{H}_\infty(S_Q|Z, E_Q=1) &= \tilde{H}_\infty(S_Q|B_Q, \mathcal{T}, S_{<Q}, Q, E_Q=1) \\ &\geq \tilde{H}_\infty(S_Q|\mathcal{T}, S_{<Q}, Q, E_Q=1) - 1 \\ &\geq n - \frac{8w \cdot t_{tot}}{\alpha^2} - 1 \end{aligned}$$

where the first inequality follows from Lemma 2.14 and the second inequality follows from Lemma 3.8.

Adding (11) and (12), and plugging in the bounds for $|U|, |\Phi|, |A_q|$, we obtain

$$\begin{aligned} \tilde{H}_\infty(X|Z, E_Q=1) + \tilde{H}_\infty(S_Q|Z, E_Q=1) &\geq \frac{n^{2/3}}{4} + \left(n - \frac{\log^3(n)}{\alpha^2} \right) \\ I(S_Q; X|Z, E_Q=1) &\leq O\left(\frac{1}{(1-\alpha)^2} \frac{n^{1+o(1)}}{m} \right) \end{aligned}$$

upon which we can invoke Lemma 3.11. By setting $\alpha := n^{-1/5}$, and $m := n^{1+\Omega(1)}$, our γ for Lemma 3.11 becomes $\gamma := \Theta(\log n)$. In particular, we get

$$\begin{aligned} \tilde{H}_\infty(X|Z, E_Q=1) + \tilde{H}_\infty(S_Q|Z, E_Q=1) &\geq n + n^{\Theta(1)} \\ I(S_Q; X|Z, E_Q=1) &\leq 2^{-\Theta(\log n)} \end{aligned}$$

Here, we remark that (13) is the bottleneck for γ . This then implies

$$\mathbb{E}_{Z, E_Q=1} \left[\left| \mathbb{E}_{(S_Q, X)|Z, E_Q=1} [\chi_{S_Q}(X)] \right| \right] < n^{-\Omega(1)}.$$

However, this contradicts Claim 3.6, due to the following chain of inequalities.

$$\begin{aligned} n^{-o(1)} &\leq \mathbb{E}_Z \left[\left| \mathbb{E}_{(S_Q, X)|Z} [\chi_{S_Q}(X)] \right| \right] \\ &\leq \mathbb{E}_Z \left[\Pr[E_Q=1|Z=z] \cdot \left| \mathbb{E}_{(S_Q, X)|Z=z, E_Q=1} [\chi_{S_Q}(X)] \right| \right] \\ &\quad + \underbrace{\mathbb{E}_Z \left[\Pr[E_Q=0|Z=z] \cdot \left| \mathbb{E}_{(S_Q, X)|Z=z, E_Q=0} [\chi_{S_Q}(X)] \right| \right]}_{\leq \Pr[E_Q=0]} \\ &\leq \mathbb{E}_Z \left[\left| \mathbb{E}_{(S_Q, X)|Z, E_Q=1} [\chi_{S_Q}(X)] \right| \right] + \Pr[E_Q=0] \\ &\leq \mathbb{E}_{Z, E_Q=1} \left[\left| \mathbb{E}_{(S_Q, X)|Z, E_Q=1} [\chi_{S_Q}(X)] \right| \right] \\ &\quad + \|Z - Z|_{E_Q=1}\|_1 + \Pr[E_Q=0] \end{aligned} \tag{14}$$

Here the underbraced term is at most $\Pr[E_Q=0]$ since $|\chi_{S_Q}(X)| \leq 1$, and the last inequality follows from the standard fact $|\mathbb{E}_X[f(X)] - \mathbb{E}_Y[f(Y)]| \leq \|f\|_\infty \cdot \|X - Y\|_1$. Write $\pi := \Pr[E_Q=1]$, $Z_1 := Z|_{E_Q=1}$ and $Z_0 := Z|_{E_Q=0}$, so that $Z = \pi Z_1 + (1 - \pi)Z_0$. Then $\|Z - Z_1\|_1$ can be upper bounded by

$$\begin{aligned} \|Z - Z_1\|_1 &= \|\pi Z_1 + (1 - \pi)Z_0 - \pi Z_1 - (1 - \pi)Z_1\|_1 \\ &\leq (1 - \pi) \cdot \|Z_0 - Z_1\|_1 \\ &\leq 2 \cdot (1 - \pi) = 2 \cdot \Pr[E_Q=0]. \end{aligned} \tag{15}$$

Combining (14) with (15) yields the final contradiction:

$$\begin{aligned} n^{-o(1)} &\leq (14) \\ &\leq \mathbb{E}_{Z, E_Q=1} \left[\left| \mathbb{E}_{(S_Q, X)|Z, E_Q=1} [\chi_{S_Q}(X)] \right| \right] \\ &\quad + 3 \cdot \Pr[E_Q=0] \\ &\leq n^{-\Omega(1)} + O(\alpha) \leq n^{-\Omega(1)} \end{aligned} \tag{16}$$

The final inequality on $\Pr[E_Q=0] = 1 - \Pr[E_Q=1]$ follows from Lemma 3.8 as

$$\Pr[E_Q=1] \geq (1 - \alpha)^2 \geq 1 - 2\alpha.$$

With $\alpha := n^{-1/5}$, (16) states $n^{-o(1)} \leq n^{-\Omega(1)}$, a contradiction. □

Acknowledgments

The author thanks Huacheng Yu and anonymous reviewers for carefully reading an earlier draft of this paper and providing valuable feedback.

References

- [BL15] Joshua Brody and Kasper Green Larsen. Adapt or Die: Polynomial Lower Bounds for Non-Adaptive Dynamic Data Structures. *THEORY OF COMPUTING*, 11:19, 2015.
- [BM13] Mark Braverman and Ankur Moitra. An information complexity approach to extended formulations. In *Proceedings of the forty-fifth annual ACM symposium on Theory of Computing*, pages 161–170, Palo Alto California USA, June 2013. ACM. URL: <https://dl.acm.org/doi/10.1145/2488608.2488629>, doi:10.1145/2488608.2488629.
- [Bra22] Mark Braverman. Communication and information complexity. *Proc. Int. Cong. Math. 2022*, 2022.
- [CEEP16] Arkadev Chattopadhyay, Jeff Edmonds, Faith Ellen, and Toniann Pitassi. Upper and Lower Bounds on the Power of Advice. *SIAM Journal on Computing*, 45(4):1412–1432, January 2016. URL: <https://epubs.siam.org/doi/10.1137/15M1031862>, doi:10.1137/15M1031862.
- [CGL15] Raphael Clifford, A. Grønlund, and K. G. Larsen. New Unconditional Hardness Results for Dynamic and Online Problems. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1089–1107, October 2015. doi:10.1109/FOCS.2015.71.
- [CT06] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, New York, NY, USA, 2006.
- [DGW19] Zeev Dvir, Alexander Golovnev, and Omri Weinstein. Static data structure lower bounds imply rigidity. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, pages 967–978, Phoenix, AZ, USA, June 2019. Association for Computing Machinery. doi:10.1145/3313276.3316348.
- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. *SIAM Journal on Computing*, 38(1):97–139, January 2008. URL: <http://epubs.siam.org/doi/10.1137/060651380>, doi:10.1137/060651380.
- [Dru12] Andrew Drucker. Limitations of Lower-Bound Methods for the Wire Complexity of Boolean Operators. In *2012 IEEE 27th Conference on Computational Complexity*, pages 170–180, June 2012. doi:10.1109/CCC.2012.39.
- [FS89] Michael L. Fredman and Michael E. Saks. The Cell Probe Complexity of Dynamic Data Structures. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 345–354, 1989.
- [JRS09] Rahul Jain, Jaikumar Radhakrishnan, and Pranab Sen. A property of quantum relative entropy with an application to privacy in quantum communication. *Journal of the ACM*, 56(6):1–32, September 2009. URL: <https://dl.acm.org/doi/10.1145/1568318.1568323>, doi:10.1145/1568318.1568323.
- [JS10] S. Jukna and G. Schnitger. Circuits with arbitrary gates for random operators, 2010. Version Number: 1. URL: <https://arxiv.org/abs/1004.5236>, doi:10.48550/ARXIV.1004.5236.

- [Juk12] Stasys Jukna. *Boolean function complexity: advances and frontiers*, volume 27. Springer Science & Business Media, 2012.
- [Juk13] Stasys Jukna. Complexity of Linear Boolean Operators. *Foundations and Trends® in Theoretical Computer Science*, 9(1):1–123, 2013. URL: <http://www.nowpublishers.com/articles/foundations-and-trends-in-theoretical-computer-science/TCS-063>, doi:10.1561/04000000063.
- [KLMS26] Michal Koucký, Bruno Loff, Tulasimohan Molli, and Michael E. Saks. The Natural Proofs Barrier against Data-Structure Lower-Bounds. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, pages 1338–1343, Salt Lake City UT USA, June 2026. ACM. URL: <https://dl.acm.org/doi/10.1145/3798129.3800843>, doi:10.1145/3798129.3800843.
- [Ko25a] Young Kun Ko. Lower Bounds for Linear Operators, 2025. Version Number: 1. URL: <https://arxiv.org/abs/2509.02730>, doi:10.48550/ARXIV.2509.02730.
- [Ko25b] Young Kun Ko. Unifying the Landscape of Super-Logarithmic Dynamic Cell-Probe Lower Bounds, 2025. Version Number: 1. URL: <https://arxiv.org/abs/2510.17717>, doi:10.48550/ARXIV.2510.17717.
- [KPI25] Oliver Korten, Toniann Pitassi, and Russell Impagliazzo. Stronger Cell Probe Lower Bounds via Local PRGs. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2480–2504, Sydney, Australia, December 2025. IEEE. URL: <https://ieeexplore.ieee.org/document/11369100/>, doi:10.1109/FOCS63196.2025.00130.
- [KW20] Young Kun Ko and Omri Weinstein. An Adaptive Step Toward the Multiphase Conjecture. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 752–761, Durham, NC, USA, November 2020. IEEE. URL: <https://ieeexplore.ieee.org/document/9317945/>, doi:10.1109/FOCS46700.2020.00075.
- [Lar12] Kasper Green Larsen. The Cell Probe Complexity of Dynamic Range Counting. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing, STOC '12*, pages 85–94, New York, NY, USA, 2012. ACM. URL: <http://doi.acm.org/10.1145/2213977.2213987>, doi:10.1145/2213977.2213987.
- [LWY20] Kasper Green Larsen, Omri Weinstein, and Huacheng Yu. Crossing the Logarithmic Barrier for Dynamic Boolean Data Structure Lower Bounds. *SIAM Journal on Computing*, 49(5):STOC18–323–STOC18–367, January 2020. URL: <https://epubs.siam.org/doi/10.1137/18M1198429>, doi:10.1137/18M1198429.
- [LY25] Kasper Green Larsen and Huacheng Yu. Super-Logarithmic Lower Bounds for Dynamic Graph Problems. *SIAM Journal on Computing*, pages FOCS23–42–FOCS23–69, February 2025. URL: <https://epubs.siam.org/doi/10.1137/24M1638215>, doi:10.1137/24M1638215.
- [MNSW98] Peter Bro Miltersen, Noam Nisan, Shmuel Safra, and Avi Wigderson. On Data Structures and Asymmetric Communication Complexity. *J. Comput. Syst. Sci.*, 57(1):37–49, 1998. doi:10.1006/jcss.1998.1577.

- [Pat07] Mihai Patrascu. Lower Bounds for 2-dimensional Range Counting. In *Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing*, STOC '07, pages 40–46, New York, NY, USA, 2007. ACM. URL: <http://doi.acm.org/10.1145/1250790.1250797>, doi:10.1145/1250790.1250797.
- [Pat10] Mihai Patrascu. Towards polynomial lower bounds for dynamic problems. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 603–610. ACM, 2010. doi:10.1145/1806689.1806772.
- [PD06] Mihai Patrascu and Erik D. Demaine. Logarithmic Lower Bounds in the Cell-Probe Model. *SIAM J. Comput.*, 35(4):932–963, April 2006. URL: <http://dx.doi.org/10.1137/S0097539705447256>, doi:10.1137/S0097539705447256.
- [PTW08] Rina Panigrahy, Kunal Talwar, and Udi Wieder. A Geometric Approach to Lower Bounds for Approximate Near-Neighbor Search and Partial Match. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 414–423, Philadelphia, PA, USA, October 2008. IEEE. URL: <http://ieeexplore.ieee.org/document/4690975/>, doi:10.1109/FOCS.2008.68.
- [PTW10] Rina Panigrahy, Kunal Talwar, and Udi Wieder. Lower Bounds on Near Neighbor Search via Metric Expansion. In *FOCS 2010*, pages 805–814, 2010.
- [Pă11] Mihai Pătraşcu. Unifying the Landscape of Cell-Probe Lower Bounds. *SIAM J. Comput.*, 40(3):827–847, 2011.
- [Sie04] Alan Siegel. On Universal Classes of Extremely Random Constant-Time Hash Functions. *SIAM J. Comput.*, 33(3):505–543, 2004. URL: <http://dx.doi.org/10.1137/S0097539701386216>, doi:10.1137/S0097539701386216.
- [Tho13] Mikkel Thorup. Mihai PăTraşCu: Obituary and Open Problems. *SIGACT News*, 44(1):110–114, March 2013. URL: <http://doi.acm.org/10.1145/2447712.2447737>, doi:10.1145/2447712.2447737.
- [Vio26] Emanuele Viola. *MATHEMATICS OF THE IMPOSSIBLE*. Cambridge University Press, 1 edition, June 2026.
- [Wil17] Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *ICM 2018*, 2017.
- [Yao81] Andrew Chi-Chih Yao. Should Tables Be Sorted? *J. ACM*, 28(3):615–628, July 1981. URL: <http://doi.acm.org/10.1145/322261.322274>, doi:10.1145/322261.322274.

A Omitted Proofs of Claims and Lemmas

To make the presentation self-contained, we include the omitted proofs of claims and lemmas.

Claim 2.15.

$$\tilde{H}_{\infty}(A|B, C) \leq \tilde{H}_{\infty}(A|B)$$

Proof. We first show the following inequality,

$$\begin{aligned} \mathbb{E}_{c \sim C|B=b} \left[\max_a \Pr[A = a|B = b, C = c] \right] \\ \geq \max_a \Pr[A = a|B = b], \end{aligned} \quad (17)$$

which is a standard fact about the max function. Let $a^* := \arg \max_a \Pr[A = a|B = b]$. Then

$$\Pr[A = a^*|B = b] = \mathbb{E}_{c \sim C|B=b} [\Pr[A = a^*|B = b, C = c]] \leq \mathbb{E}_{c \sim C|B=b} \left[\max_a \Pr[A = a|B = b, C = c] \right]$$

With (17) established, we are ready to prove the claim. Recall that

$$\tilde{H}_\infty(A|B, C) := -\log \left(\mathbb{E}_{b, c \sim B, C} \left[\max_a \Pr[A = a|B = b, C = c] \right] \right)$$

Taking expectation over B on both sides of (17), we have

$$\mathbb{E}_{b, c \sim B, C} \left[\max_a \Pr[A = a|B = b, C = c] \right] \geq \mathbb{E}_{b \sim B} \left[\max_a \Pr[A = a|B = b] \right]$$

Therefore, we get

$$\begin{aligned} \tilde{H}_\infty(A|B, C) &= -\log \left(\mathbb{E}_{b, c \sim B, C} \left[\max_a \Pr[A = a|B = b, C = c] \right] \right) \\ &\leq -\log \left(\mathbb{E}_{b \sim B} \left[\max_a \Pr[A = a|B = b] \right] \right) \\ &= \tilde{H}_\infty(A|B) \end{aligned}$$

□

Claim 2.16 ([JRS09]). *Let $\mathcal{D}$ be a distribution and $\mathcal{U}$ the uniform distribution over some $\mathcal{S} \subseteq \{0, 1\}^n$. If $D(\mathcal{D}||\mathcal{U}) \leq t$ with $t > 1$, then for every $\alpha \geq 4$ there exists an event E such that*

$$\begin{aligned} \mathcal{D}(E) &\geq 1 - \frac{1}{\alpha} \\ H_\infty(\mathcal{D}|_E) &\geq \log |\mathcal{S}| - 2\alpha t \end{aligned}$$

Proof. We partition the domain according to $\mathcal{D}(x)$. Let E denote the set of x such that

$$\log \frac{\mathcal{D}(x)}{\mathcal{U}(x)} < 2\alpha t - 1,$$

and let $p := 1 - \mathcal{D}(E)$ and $q := 1 - \mathcal{U}(E)$, so that $q \leq 2^{-(2\alpha t - 1)} p$. By the data-processing inequality (or log-sum inequality) for the indicator of E ,

$$\begin{aligned} t &\geq D(\mathcal{D}||\mathcal{U}) \geq p \log \frac{p}{q} + (1 - p) \log \frac{1 - p}{1 - q} \\ &\geq p(2\alpha t - 1) - p \log e, \end{aligned}$$

where the last inequality uses $1 - q \leq 1$ and $(1 - p) \log(1 - p) \geq -p \log e$. As $\alpha t \geq 4 \geq 1 + \log e$, we get $p \leq t/(2\alpha t - 1 - \log e) \leq t/(\alpha t) = 1/\alpha$. In particular $\mathcal{D}(E) \geq 1 - 1/\alpha \geq 1/2$, so for any x in the support of $\mathcal{D}|_E$,

$$\mathcal{D}|_E(x) = \frac{\mathcal{D}(x)}{\mathcal{D}(E)} \leq 2 \cdot \mathcal{D}(x) < 2^{2\alpha t} \cdot \mathcal{U}(x) = 2^{-\log |\mathcal{S}| + 2\alpha t},$$

which then gives $H_\infty(\mathcal{D}|_E) \geq \log |\mathcal{S}| - 2\alpha t$. □

Towards the proof of Lemma 3.11, we need the following well-known fact about the Hadamard matrix.

Claim A.1 (Lindsey's Lemma). *Let H be a $2^n \times 2^n$ Hadamard matrix. Let μ and ν be distributions over $\{0, 1\}^n$, written as vectors. Then*

$$|\mu^T H \nu| \leq \|\mu\|_2 \|\nu\|_2 \cdot 2^{n/2}$$

Also, Fact 2.12 implies the following simple proposition about the ℓ_2^2 norm of the distribution versus its min-entropy.

Proposition A.2.

$$\mathbb{E}_{b \sim B} \left[\sum_a \Pr[A = a | B = b]^2 \right] \leq 2^{-\tilde{H}_\infty(A|B)}$$

Proof.

$$\sum_a \Pr[A = a | B = b]^2 = 2^{-H_2(A|B=b)} \leq 2^{-H_\infty(A|B=b)}$$

Therefore,

$$\mathbb{E}_{b \sim B} \left[\sum_a \Pr[A = a | B = b]^2 \right] \leq \mathbb{E}_{b \sim B} \left[2^{-H_\infty(A|B=b)} \right] = 2^{-\tilde{H}_\infty(A|B)}$$

where the last equality follows from the definition of average min-entropy. $\square$

Now we are ready to prove Lemma 3.11. We prove an equivalent statement. Recall that we abbreviate the conditioning on a fixed pair (a, z) as $\mathcal{E} := (A = a, Z = z)$.

Lemma 3.11. *No setting of a random variable $Z = z$ and A , which contains z_{out} , can simultaneously satisfy all three of the following inequalities for any $\gamma \geq 3$*

$$\mathbb{E}_{A|Z=z} [|\mathbb{E}_{Y,X|\mathcal{E}} [\chi_Y(X) \cdot z_{out}]|] \geq 2^{-\gamma+2} \quad (18)$$

$$\tilde{H}_\infty(Y|A, Z = z) + \tilde{H}_\infty(X|A, Z = z) \geq n + 2\gamma \quad (19)$$

$$I(Y; X|A, Z = z) \leq 2^{-2\gamma} \quad (20)$$

Proof. For the sake of contradiction, suppose such z and A exist. First, as z_{out} is ± 1 , the absolute value of the guess does not change. That is

$$|\mathbb{E}_{Y,X|\mathcal{E}} [\chi_Y(X) \cdot z_{out}]| = |\mathbb{E}_{Y,X|\mathcal{E}} [\chi_Y(X)]| \quad (21)$$

Then we can use the ℓ_1 bound to have

$$|\mathbb{E}_{Y,X|\mathcal{E}} [\chi_Y(X)]| \leq |Y|_{\mathcal{E}} \cdot H \cdot |X|_{\mathcal{E}} \quad (22)$$

$$+ \|Y|_{\mathcal{E}} \times X|_{\mathcal{E}} - (Y, X)|_{\mathcal{E}}\|_1 \quad (23)$$

We bound the expectation of (23). Our KL-divergence term is then equal to the mutual information between Y and X conditioned on $Z = z$. Namely, using the chain rule for the KL-divergence,

$$\begin{aligned} & D((Y, X)|_{\mathcal{E}} || Y|_{\mathcal{E}} \times X|_{\mathcal{E}}) \\ &= \underbrace{D(X|_{\mathcal{E}} || X|_{\mathcal{E}})}_{=0} + \mathbb{E}_{x \sim X|_{\mathcal{E}}} [D(Y|_{X=x, \mathcal{E}} || Y|_{\mathcal{E}})] \\ &= I(Y; X|_{\mathcal{E}}) \end{aligned}$$

Then, due to Pinsker's inequality (Fact 2.6), we have

$$\|Y|_{\mathcal{E}} \times X|_{\mathcal{E}} - (Y, X)|_{\mathcal{E}}\|_1 \leq 2\sqrt{I(Y; X|_{\mathcal{E}})}$$

Taking expectation over A and applying Jensen's inequality,

$$\mathbb{E}_{A|Z=z} [\|Y|_{\mathcal{E}} \times X|_{\mathcal{E}} - (Y, X)|_{\mathcal{E}}\|_1] \leq 2\sqrt{I(Y; X|A, Z=z)} \leq 2^{-\gamma+1} \quad (24)$$

where the last bound follows from (20).

Next, we bound (22). Due to Claim A.1 and the Cauchy-Schwarz inequality,

$$\begin{aligned} \mathbb{E}_{A|Z=z} [|Y|_{\mathcal{E}} \cdot H \cdot X|_{\mathcal{E}}|] &\leq \mathbb{E}_{A|Z=z} \left[2^{n/2} \cdot \|Y|_{\mathcal{E}}\|_2 \cdot \|X|_{\mathcal{E}}\|_2 \right] \\ &\leq 2^{n/2} \cdot \sqrt{\mathbb{E}_{A|Z=z} [\|Y|_{\mathcal{E}}\|_2^2] \cdot \mathbb{E}_{A|Z=z} [\|X|_{\mathcal{E}}\|_2^2]} \end{aligned}$$

Proposition A.2 implies

$$\begin{aligned} \mathbb{E}_{A|Z=z} [\|Y|_{\mathcal{E}}\|_2^2] &\leq 2^{-\tilde{H}_{\infty}(Y|A, Z=z)} \\ \mathbb{E}_{A|Z=z} [\|X|_{\mathcal{E}}\|_2^2] &\leq 2^{-\tilde{H}_{\infty}(X|A, Z=z)} \end{aligned}$$

which would in turn imply

$$\begin{aligned} \sqrt{\mathbb{E}_{A|Z=z} [\|Y|_{\mathcal{E}}\|_2^2] \cdot \mathbb{E}_{A|Z=z} [\|X|_{\mathcal{E}}\|_2^2]} &\leq 2^{-(\tilde{H}_{\infty}(Y|A, Z=z) + \tilde{H}_{\infty}(X|A, Z=z))/2} \\ &\leq 2^{-\frac{n+2\gamma}{2}} = 2^{-n/2} \cdot 2^{-\gamma} \end{aligned}$$

which then implies (22) is upper bounded by

$$(22) \leq 2^{n/2} \cdot 2^{-n/2} \cdot 2^{-\gamma} = 2^{-\gamma} \quad (25)$$

Therefore, we get

$$\mathbb{E}_{A|Z=z} [|\mathbb{E}_{Y, X|_{\mathcal{E}}} [\chi_Y(X) \cdot z_{out}]|] \leq \frac{3}{2^{\gamma}} \quad (26)$$

which contradicts (18). $\square$

B Lifting Theorem

We consider the following “hard” functions, which, roughly speaking, are functions f with small discrepancy under every product distribution of large min-entropy.

Definition B.1 (Hard Functions). *Let $\mathcal{S} \subseteq \{0, 1\}^n$. We say $f : \{0, 1\}^{2n} \rightarrow \{\pm 1\}$ is hard (with respect to $\mathcal{S}$) if there is a constant $c > 0$ such that, for all large enough n , the following holds for all random variables S over $\mathcal{S}$, X over $\{0, 1\}^n$, and A such that S and X are independent conditioned on $A = a$, for every a . If*

$$\tilde{H}_{\infty}(S|A) \geq \log |\mathcal{S}| - \log^3 n, \quad \tilde{H}_{\infty}(X|A) \geq \frac{\sqrt{n}}{3},$$

then

$$\mathbb{E}_{a \sim A} [|\mathbb{E}_{(S, X)|A=a} [f(S, X)]|] \leq n^{-c}.$$

Example: Inner Product Inner Product over $\mathbb{F}_2$ satisfies the definition with $\mathcal{S} = \{0, 1\}^n$ and any constant $c > 0$. By Claim A.1 together with $\|\mu\|_2^2 \leq 2^{-H_\infty(\mu)}$ (Fact 2.12), for every a ,

$$|\mathbb{E}_{(S,X)|A=a}[f(S, X)]| \leq 2^{n/2} \cdot 2^{-H_\infty(S|A=a)/2} \cdot 2^{-H_\infty(X|A=a)/2}.$$

Averaging over a and applying the Cauchy–Schwarz inequality,

$$\begin{aligned} \mathbb{E}_{a \sim A} [|\mathbb{E}_{(S,X)|A=a}[f(S, X)]|] &\leq 2^{n/2} \left(\mathbb{E}_{a \sim A} [2^{-H_\infty(S|A=a)}] \right)^{1/2} \left(\mathbb{E}_{a \sim A} [2^{-H_\infty(X|A=a)}] \right)^{1/2} \\ &= 2^{(n - \tilde{H}_\infty(S|A) - \tilde{H}_\infty(X|A))/2} \leq 2^{(\log^3 n - \sqrt{n}/3)/2}, \end{aligned}$$

which is $2^{-\Omega(\sqrt{n})}$. Disjointness, on the other hand, does not satisfy the definition: sampling and revealing a few coordinates of X already gives a nontrivial advantage (see [BM13]).

A Good Event for Small Divergence The proof of Claim 2.16 also gives the following corollary for the regime in which the divergence is much smaller than 1.

Corollary B.2 (of the proof of Claim 2.16). *Let S be uniform over a finite set $\mathcal{S}$, and let Y be a random variable with $I(S; Y) \leq \varepsilon \leq 1/8$. Let E be the event that the chosen pair (s, y) of (S, Y) satisfies $\Pr[S = s|Y = y] \leq 2/|\mathcal{S}|$. Then*

$$\Pr[E] \geq 1 - \sqrt{2\varepsilon} \geq \frac{1}{2}, \quad \tilde{H}_\infty(S|Y, E=1) \geq \log |\mathcal{S}| - 2.$$

Proof. Denote by μ_y the distribution of S conditioned on $Y = y$, and by $\mathcal{U}$ the uniform distribution over $\mathcal{S}$, which is the original distribution of S . We follow the proof of Claim 2.16 for each $\mathcal{D} = \mu_y$, with threshold 1: conditioned on $Y = y$, the event E is the set

$$E_y := \left\{ s : \log \frac{\mu_y(s)}{\mathcal{U}(s)} \leq 1 \right\} = \left\{ s : \mu_y(s) \leq \frac{2}{|\mathcal{S}|} \right\}.$$

The main change is the bound on $\Pr[E]$, for which, as the divergence is now below 1, we use Pinsker’s inequality. On the complement $\mathcal{A}_y := \mathcal{S} \setminus E_y$, we have $\mu_y > 2\mathcal{U}$, so that

$$\mu_y(\mathcal{A}_y) \leq 2(\mu_y(\mathcal{A}_y) - \mathcal{U}(\mathcal{A}_y)) \leq 2\|\mu_y - \mathcal{U}\|_{TV}.$$

Averaging over y , and using Pinsker’s inequality (Fact 2.6), Jensen’s inequality and $\log e \geq 1$,

$$\begin{aligned} \Pr[E=0] = \mathbb{E}_y [\mu_y(\mathcal{A}_y)] &\leq 2\mathbb{E}_y \left[\sqrt{\frac{D(\mu_y||\mathcal{U})}{2 \log e}} \right] \\ &\leq \sqrt{2\mathbb{E}_y [D(\mu_y||\mathcal{U})]} = \sqrt{2I(S; Y)} \leq \sqrt{2\varepsilon} \leq \frac{1}{2}. \end{aligned}$$

For the min-entropy, as in the proof of Claim 2.16, for every y with $\mu_y(E_y) > 0$ and every $s \in E_y$,

$$\Pr[S = s|Y = y, E=1] = \frac{\mu_y(s)}{\mu_y(E_y)} \leq \frac{2}{|\mathcal{S}| \mu_y(E_y)}.$$

Conditioned on $E=1$, the value y has probability $\Pr[Y = y] \mu_y(E_y) / \Pr[E]$, so averaging over y gives

$$2^{-\tilde{H}_\infty(S|Y, E=1)} \leq \sum_y \frac{\Pr[Y = y] \mu_y(E_y)}{\Pr[E]} \cdot \frac{2}{|\mathcal{S}| \mu_y(E_y)} \leq \frac{2}{|\mathcal{S}| \Pr[E]} \leq \frac{4}{|\mathcal{S}|}. \quad \square$$

We are now ready to prove the following general lifting theorem, which extends Theorem 1.1 to every hard function.

Theorem B.3 (Lifting Theorem). *Let f be hard (Definition B.1) and let $m = n^{1+\Omega(1)}$. Then any data structure for the Multiphase Problem for f whose update time satisfies $t_u = n^{o(1)}$ must have*

$$t_{tot} \geq \Omega \left(\left(\frac{\log n}{\log(t_u w)} \right)^2 \right).$$

Proof. The proof follows that of Lemma 3.12, except that Z contains the transcripts of only a random set of $k \ll m$ queries, as in [KW20], and Definition B.1 replaces Lemma 3.11, which is specific to Inner Product. The protocol does not know which queries are in the set, so the transcripts of the other queries in the set carry little information about S_Q and Corollary B.2 applies. The transcript of Q itself costs only $|A_Q|$ bits of average min-entropy by Lemma 2.14.

Setup Suppose, for contradiction, that $t_{tot} = o((\log n / \log(t_u w))^2)$, and let Π be the 2.5-round protocol of Corollary 3.5, so that

$$|U| \leq n^{1+o(1)}, \quad |\Phi| \leq n - \frac{n^{2/3}}{4}, \quad |A_Q| \leq o(\log^3(n)),$$

and $\text{adv}_{\vec{S}}(G_f, \Pi) \geq n^{-o(1)}$ for every $\vec{S}$. As the hard distribution, we select each S_q independently and uniformly from $\mathcal{S}$, and X uniformly at random. We now define the random variables used in the rest of the proof.

- *The random set of queries.* Let $\sigma \in \text{Sym}([m])$ be a uniformly random permutation and $j \in [k]$ a uniformly random index, independent of each other and of everything else, where

$$k := \sqrt{nm}.$$

We need $k \gg |U|$ for the correlation bound and $k \ll m/|A_Q|$ for the information bound below; as $m = n^{1+\Omega(1)}$, this k satisfies both, and $k \leq m/2$ for large enough n . The random set of queries is $\sigma([k])$, and we write

$$Q := \sigma(j), \quad J' := \sigma([k] \setminus \{j\})$$

for the query we analyze and for the other queries in the set, respectively.

- *The conditioning.* Let

$$\begin{aligned} Z &:= (\mathcal{T}_\sigma, S_{\sigma(<j)}, B_Q, \sigma, j), & \mathcal{T}_\sigma &:= (\Phi, A_{\sigma(1)}, \dots, A_{\sigma(k)}), \\ Y &:= (\Phi, A_{J'}, S_{\sigma(<j)}, \sigma, j), & S_{\sigma(<j)} &:= (S_{\sigma(1)}, \dots, S_{\sigma(j-1)}). \end{aligned}$$

That is, Z is the Z of Lemma 3.12 with the transcripts restricted to the random set $\sigma([k])$, ordered by σ , and Y is Z without the messages A_Q and B_Q of Q itself.

- *The good event.* Let E be the event of Corollary B.2 for $S = S_Q$ and Y , that is, the event that the chosen pair (s, y) of (S_Q, Y) satisfies $\Pr[S_Q = s | Y = y] \leq 2/|\mathcal{S}|$. As in Section 3.2, we identify E with its indicator random variable.

The permutation plays two roles: Q is a uniformly random query, and $\sigma([k])$ is a uniformly random set of queries containing Q that is independent of the protocol.

Information from the Other Queries We first show that the transcripts of the other queries in the random set reveal little information about S_Q .

Claim B.4.

$$I(S_Q; Y) \leq \frac{|\Phi| + k|A_q|}{m - k + 1} \leq \frac{2n}{m} + 2\sqrt{\frac{n}{m}} \cdot |A_q| = n^{-\Omega(1)}. \quad (27)$$

Proof. The argument is essentially from [KW20], using the standard direct-sum technique. The intuition is as follows. Together, Φ and the transcripts $A_{J'}$ consist of at most $|\Phi| + k|A_q|$ bits, and they may reveal information about any of the m coordinates of $\vec{S}$. However, the protocol does not know σ , so it cannot direct this information towards S_Q : given J' , the query Q is uniform over the $m - k + 1$ queries outside J' , and as the coordinates of $\vec{S}$ are independent, the information splits among these queries, so that S_Q receives only a $1/(m - k + 1)$ fraction of it on average.

Formally, as S_Q is uniform over $\mathcal{S}$ for every fixing of (σ, j) , and as $S_{\sigma(<j)}$ is part of $S_{J'}$, writing $Y_{\max} := (\Phi, A_{J'}, S_{J'})$,

$$I(S_Q; Y) \leq I(S_Q; Y_{\max} | \sigma, j) = \mathbb{E}_{J'} \left[\frac{1}{m - k + 1} \sum_{q \notin J'} I(S_q; Y_{\max}) \right],$$

where the equality holds as, conditioned on J' , the query Q is uniform over the $m - k + 1$ queries outside J' . For every fixed J' ,

$$\sum_{q \notin J'} I(S_q; Y_{\max}) \leq I(S_{[m] \setminus J'}; Y_{\max}) = I(S_{[m] \setminus J'}; \Phi, A_{J'} | S_{J'}) \leq |\Phi| + k|A_q|.$$

Here the equality holds as $S_{[m] \setminus J'}$ is independent of $S_{J'}$, and the first inequality is the chain rule over the queries outside J' , together with

$$I(S_q; Y_{\max} | S_T) = I(S_q; Y_{\max}, S_T) \geq I(S_q; Y_{\max}) \quad \text{for every set } T \not\ni q \text{ of queries,}$$

as S_q is independent of S_T . This proves the first inequality in (27), and the second one uses $|\Phi| \leq n$, $k = \sqrt{nm}$ and $m - k \geq m/2$. $\square$

Conditioning on a Good Event By Claim B.4, Corollary B.2 applies to $S = S_Q$ and Y with $\varepsilon := I(S_Q; Y) \leq 1/8$, so that

$$\Pr[E] \geq \frac{1}{2}, \quad \eta := \Pr[E=0] \leq \sqrt{2I(S_Q; Y)} = n^{-\Omega(1)}.$$

Conditioned on $E=1$, we obtain the analogues of Lemma 3.8, Lemma 3.9 and Lemma 3.10.

Claim B.5. *For large enough n ,*

$$\tilde{H}_{\infty}(S_Q | Z, E=1) \geq \log |\mathcal{S}| - \log^3 n, \quad (28)$$

$$\tilde{H}_{\infty}(X | Z, E=1) \geq \frac{\sqrt{n}}{3}, \quad (29)$$

$$I(S_Q; X | Z, E=1) \leq \frac{2|U|}{k} = n^{-\Omega(1)}. \quad (30)$$

Proof. For (28), recall that Z consists of Y , A_Q and B_Q , where (A_Q, B_Q) takes at most $2^{|A_Q|+1}$ values. So Lemma 2.14, applied conditioned on $E=1$, and Corollary B.2 give

$$\tilde{H}_\infty(S_Q|Z, E=1) \geq \tilde{H}_\infty(S_Q|Y, E=1) - |A_Q| - 1 \geq \log |\mathcal{S}| - |A_Q| - 3 \geq \log |\mathcal{S}| - \log^3 n,$$

as $|A_Q| = o(\log^3 n)$.

For (29), we follow the proof of Lemma 3.9. Let R_A denote Alice's private randomness and $W := (\vec{S}, \Phi, R_A, \sigma, j)$. Then W determines every A_q , hence Y and A_Q , and hence also E , which is determined by S_Q and Y . By the definition of average min-entropy (as in (9)) and $\Pr[E] \geq 1/2$, then Claim 2.15, and then Lemma 2.14 for B_Q and for Φ ,

$$\begin{aligned} \tilde{H}_\infty(X|Z, E=1) &\geq \tilde{H}_\infty(X|Z, E) - 1 \geq \tilde{H}_\infty(X|W, B_Q) - 1 \geq \tilde{H}_\infty(X|\vec{S}, R_A, \sigma, j) - |\Phi| - 2 \\ &= n - |\Phi| - 2 \geq \frac{n^{2/3}}{4} - 2 \geq \frac{\sqrt{n}}{3}. \end{aligned}$$

For (30), we follow the proof of Lemma 3.10 in three steps. First, we remove the conditioning on E , at the cost of a factor $1/\Pr[E] \leq 2$:

$$I(S_Q; X|Z, E=1) \leq \frac{I(S_Q; XU|Z, E)}{\Pr[E]} \leq 2 I(S_Q; XU|Z).$$

The first inequality holds as $I(S_Q; XU|Z, E)$ is the average of $I(S_Q; XU|Z, E=b)$ over $b \in \{0, 1\}$, weighted by $\Pr[E=b]$, and as adding U can only increase the information. The second follows from Fact 2.9 with $I(XU; E|S_Q, Z) = 0$, as E is determined by S_Q and Z . Next, we drop B_Q from Z :

$$I(S_Q; XU|Z) \leq I(S_Q; XU|\mathcal{T}_\sigma, S_{\sigma(<j)}, \sigma, j).$$

This again follows from Fact 2.9, now with $I(S_Q; B_Q|XU, \mathcal{T}_\sigma, S_{\sigma(<j)}, \sigma, j) = 0$, as B_Q is determined by U , A_Q and Bob's randomness, which is independent of $(\vec{S}, R_A, \sigma, j)$ given (X, U, Φ) . Finally, as j is uniform over $[k]$, the chain rule (Fact 2.7) over the queries $\sigma(1), \dots, \sigma(k)$, in this order, gives the direct-sum bound

$$\begin{aligned} I(S_Q; XU|\mathcal{T}_\sigma, S_{\sigma(<j)}, \sigma, j) &= \frac{1}{k} \mathbb{E}_\sigma \left[\sum_{i=1}^k I(S_{\sigma(i)}; XU|\mathcal{T}_\sigma, S_{\sigma(<i)}) \right] = \frac{1}{k} \mathbb{E}_\sigma [I(S_{\sigma([k])}; XU|\mathcal{T}_\sigma)] \\ &\leq \frac{1}{k} \mathbb{E}_\sigma [I(\vec{S}; XU|\mathcal{T}_\sigma)] \leq \frac{I(\vec{S}; XU|\Phi)}{k} \leq \frac{|U|}{k}, \end{aligned}$$

where the first inequality follows from the chain rule, and the last two are as in the proof of Lemma 3.10. Combining the three steps,

$$I(S_Q; X|Z, E=1) \leq \frac{2|U|}{k} \leq \frac{2n^{1+o(1)}}{\sqrt{nm}} = n^{-\Omega(1)}. \quad \square$$

Contradiction We conclude the proof of the theorem by deriving a contradiction as in the proofs of Claim 3.6 and Lemma 3.12. As Q is uniform over $[m]$ and independent of $\vec{S}$, X and the randomness of Π ,

$$\mathbb{E}[B_Q \cdot f(S_Q, X)] = \mathbb{E}_{\vec{S}} [\text{adv}_{\vec{S}}(G_f, \Pi)] \geq n^{-o(1)}.$$

As $|f| \leq 1$ and B_Q is determined by Z , restricting to E loses at most η :

$$n^{-o(1)} \leq \mathbb{E}_{Z|E=1} \left[\left| \mathbb{E}_{(S_Q, X)|Z, E=1} [f(S_Q, X)] \right| \right] + O(\eta). \quad (31)$$

For z in the support of Z conditioned on $E=1$, consider the product of the two conditional marginals,

$$R_z := S_Q|_{Z=z, E=1} \times X|_{Z=z, E=1}.$$

As R_z has the same marginals, (28) and (29) let us apply Definition B.1 with A distributed as Z conditioned on $E=1$ and (S, X) distributed as R_z given $A = z$, which gives

$$\mathbb{E}_{Z|E=1} [\|\mathbb{E}_{R_Z}[f]\|] \leq n^{-c}.$$

On the other hand, as $|f| \leq 1$, Pinsker's inequality (Fact 2.6), Jensen's inequality and (30) give

$$\begin{aligned} \mathbb{E}_{Z|E=1} \left[\left| \mathbb{E}_{(S_Q, X)|Z, E=1} [f(S_Q, X)] - \mathbb{E}_{R_Z}[f] \right| \right] &\leq \mathbb{E}_{Z|E=1} [\|(S_Q, X)|_{Z, E=1} - R_Z\|_1] \\ &\leq \sqrt{2I(S_Q; X|Z, E=1)} = n^{-\Omega(1)}. \end{aligned}$$

Combining the last two bounds with (31) and $\eta = n^{-\Omega(1)}$,

$$n^{-o(1)} \leq n^{-c} + n^{-\Omega(1)},$$

which is a contradiction. □