

Tight Bounds for Monotone Minimal Perfect Hashing[□]

Sepehr Assadi
Rutgers University

Martin Farach-Colton
Rutgers University

William Kuszmaul
MIT

Abstract

The monotone minimal perfect hash function (MMPHF) problem is the following indexing problem. Given a set $S = \{s_1, \dots, s_n\}$ of n distinct keys from a universe U of size u , create a data structure D that answers the following query:

$$\text{Rank}(q) = \begin{cases} \text{rank of } q \text{ in } S & q \in S \\ \text{arbitrary answer} & \text{otherwise.} \end{cases}$$

Solutions to the MMPHF problem are in widespread use in both theory and practice.

The best upper bound known for the problem encodes D in $O(n \log \log \log u)$ bits and performs queries in $O(\log u)$ time. It has been an open problem to either improve the space upper bound or to show that this somewhat odd looking bound is tight.

In this paper, we show the latter: any data structure (deterministic or randomized) for monotone minimal perfect hashing of any collection of n elements from a universe of size u requires $(n \log \log \log u)$ expected bits to answer every query correctly.

We achieve our lower bound by dening a graph G where the nodes are the possible u inputs and where two nodes are adjacent if they cannot share the same D . The size of D is then lower bounded by the log of the chromatic number of G . Finally, we show that the fractional chromatic number (and hence the chromatic number) of G is lower bounded by $2^{(n \log \log \log u)}$.

1 Introduction

The monotone minimal perfect hash function (MMPHF) problem is the following indexing problem. Given a set $S = \{s_1, \dots, s_n\}$ of n distinct keys from a universe U of size u , create a data structure D that answers the following query:

$$\text{Rank}(q) = \begin{cases} \text{rank of } q \text{ in } S & q \in S \\ \text{arbitrary answer} & \text{otherwise.} \end{cases}$$

The name MMPHF comes from interpreting the data structure D as a hash function: given a sorted array $A = [a_1, \dots, a_n]$, D is hashing each a_i to its position i . The hash function is minimal, meaning it maps n items to n distinct positions, and monotone, meaning $a_i < a_j \implies D(a_i) < D(a_j)$.

It may seem strange at first glance that D is permitted to return arbitrary answers on negative queries. A key insight, however, is that this relaxation allows for asymptotic improvements in space efficiency: whereas the set S would require

$(n \log(u=n))$ bits to encode, Belazzougui, Boldi, Pagh and Vigna [BBPV09] show that it is possible to construct an MMPHF D using as few as $O(n \log \log \log u)$ bits, while supporting $O(\log u)$ -time queries.

The remarkable space efficiency of MMPHF makes it useful for a variety of practical applications (e.g., in security [BCO11], key-value stores [LFAK11] and information retrieval [Nav14]). A high-performance implementation can be found in the Sux4J library [BV08, BBPV11]. MMPHF has also been widely used in

[□]A full version of the paper is available at: <https://arxiv.org/abs/2207.10556>.

the theory community for the design of space-efficient combinatorial pattern-matching algorithms (see, e.g., [BN14, GNP20, Bel14, BN15, CFP⁺15, BCKM20, BGMP16, GOR10]).

Despite the widespread use of MMPHF, it remains an open question [BBPV09, Bol15, D⁺18] to determine the optimal bounds for solving this problem. The best lower bound achieved so far [BBPV11, D⁺18] is $(n \log \log u)$ bits (which follows immediately from the same lower bound for minimal perfect hashing [Meh82]). Even disregarding applications (and the running time to answer queries), the information-theoretic question as to how many bits a MMPHF requires has been posed as a problem of independent combinatorial interest [D⁺18].

Our result. We fully settle this question by establishing the following result:

Theorem 1 (Formalized in Theorem 2). Any data structure (deterministic or randomized) for monotone minimal perfect hashing of any collection of n elements from a universe of size u requires $(n \log \log u)$ expected bits to answer every query correctly. The lower bound holds whenever u is at least $n^{1+\frac{1}{\log n}}$ and at most $\exp(\exp(\text{poly}(n)))$.

Thus, surprisingly, the $O(n \log \log \log u)$ bound achieved by [BBPV09] is asymptotically optimal. The boundary conditions on u in Theorem 1 are also natural in the following sense. There are two trivial solutions for MMPHF. One encodes the input set S in $O(u)$ bits and the other builds a perfect hash table from elements of S to their rank in $O(n \log n)$ space. When u is very small, say, $u = O(n)$, the first solution uses $O(u) = o(n \log \log \log u)$ bits. And when u is very large, that is when u is even beyond $\exp(\exp(\text{poly}(n)))$, then the $O(n \log n)$ -bit solution uses $o(n \log \log \log u)$ bits. (See also the variable-size bucketing reduction of [BBPV11] which reduces the universe size from u to u/n). Our lower bound in Theorem 1 covers almost the entire range in between.

The lower bound achieved by Theorem 1 is remarkably general: it applies independently of the running time of the data structure; and it applies even to randomized data structures that are permitted to store their random bits for free.

Our techniques. The most intuitive approach toward proving a lower bound of d bits on the size of an MMPHF is to encode a d -bit string into the state of the data structure. This approach is already hindered by the fact that MMPHFs only support positive queries, however. If the user already knows which elements are in the input, then the MMPHF encodes no interesting information—but if the user only has partial information about the input, then the user can only get useful information from a small portion of possible MMPHF queries. The previous

(n) lower bound of [Meh82, BBPV11, D⁺18] addresses this as follows: consider any bit-string $x \in \{0, 1\}^d$ and define:

$$S(x) := \{f_3; 6; \dots; 3d; f_{3i+1} \mid i \in [d]; x_i = 1\} \cup \{f_{3i} \mid i \in [d]; x_i = 0\}.$$

For every $i \in [d]$, first, $3i$ belongs to $S(x)$ and thus is a positive query, and secondly, $\text{Rank}(3i) = 2(i-1) + x_i$. This allows us to recover x from any MMPHF for $S(x)$, proving a lower bound of $d = (n)$ bits for MMPHF on size- n subsets of universe $[3n+1]$. This approach, however, seems to be stuck at proving any $\Omega(n)$ lower bound as these “direct encodings” ignore the delicate interaction between different elements in the input set¹.

To get around these obstacles, we take a different approach to proving Theorem 1. We construct a “conflict graph” G whose vertices are the possible inputs to an MMPHF problem for a fixed n and u . Two vertices are adjacent in G if they cannot have the same MMPHF index, that is, if the vertices share an element but with a different rank. Any MMPHF induces a proper coloring of this graph, where the color of a vertex corresponds to its MMPHF representation. As a result, the chromatic number of the conflict graph is a lower bound on how many different MMPHF representations we must have, implying that some input must have a representation of size at least $\log(G)$ bits. This reduces our task to combinatorial problem of lower bounding (G) .²

The problem of bounding chromatic number of graphs defined over these types of set-systems has a rich history in the discrete math literature; see, e.g. [EH66, FHR92, DLR95, ST11]. For instance, Erdős and Hajnal [EH66]

¹Any lower bound of d bits for a data structure immediately implies an encoding of d -bit strings in the state of the data structure by just assigning one bit-string to each state. This means that there is never a formal proof that one cannot encode a bit-string in a data structure and still prove a lower bound.

²Slightly more care must be taken when bounding the expected size of a MMPHF that is permitted to take different sizes on different inputs.

study shift-graphs that have vertices corresponding to n -element subsets of $[u]$ and edges between vertices $(a_1; a_2; \dots; a_n)$ and $(a_2; \dots; a_n; a_{n+1})$ for all $a_1 < a_2 < \dots < a_{n+1}$. They prove that the chromatic number of the shift-graph is $(1 + o(1)) \log^{(n-1)}(u)$, namely, the $(n-1)$ -th iterated logarithm of u . The shift-graph is a subgraph of our conict graph. Thus, by taking $u = 2^{(n)}$, i.e., the tower of twos of height (n) , we can have $\chi(G) = 2^{!(n)}$, and thus prove an $!(n)$ lower bound for MMPHF on n -subsets of (extremely large) universes of size $u = 2^{(n)}$. This is the starting point of our approach. We now need to dramatically decrease the size of the universe, while also dramatically increasing the bound on the chromatic number by considering the conict graph itself, and not only its shift-subgraph.

To lower bound the chromatic number of the conict graph, we consider the relaxation of this problem via fractional colorings (see Section 2.2). Given that this latter problem can be formulated as a linear program (LP), a natural way for proving a lower bound on its value is to exhibit a feasible dual solution instead³. This corresponds to the following problem: exhibit a distribution on vertices of the graph so that for any independent set, the probability that a vertex sampled from the distribution belongs to the independent set is bounded by p ; this then implies that the fractional chromatic number (and in turn the chromatic number) are lower bounded by $1/p$. The main technical novelty of our work lies in the introduction of a highly non-trivial such distribution and the analysis of this probability bound for each independent set (we postpone the overview of this part to Section 4.1 after we setup the required background). This allows us to lower bound the (fractional) chromatic number of the conict-graph by

$(n \log n)$ when the universe is of size $u = 2^{2^{\text{poly}(n)}}$ which gives an $(n \log \log u)$ lower bound for MMPHF on such universes.

Working with fractional colorings, beside being an immensely helpful analytical tool, has several additional benets for us. Firstly, unlike standard (integral) colorings, fractional colorings admit a natural direct product property over a certain union of graphs; this allows us to extend the lower bound for MMPHF from universes of size doubly exponential in n (which are admittedly not the most interesting setting of parameters), all the way down to universes of size $n^{1+o(1)}$. Secondly, unlike the (integral) chromatic number, which yields a lower bound only on the space of deterministic MMPHFs, we show that lower bounding the fractional chromatic number allows us to prove a lower bound even for randomized MMPHFs that have access to their randomness for free. We believe this technique, namely, dening a proper conict graph and bounding its fractional coloring by exhibiting a feasible dual solution, may be applicable to many other data structure problems and is therefore interesting in its own right.

2 Preliminaries

Notation. For any integer $t > s > 1$, we let $[t] := \{1; \dots; t\}$ and let $[s : t] = \{s; \dots; t\}$. For a tuple $(X_1; \dots; X_t)$, we further dene $X_{<i} := (X_1; \dots; X_{i-1})$ and $X_{i-1} := (X_1; \dots; X_{i-1}; X_{i+1}; \dots; X_t)$.

2.1 Problem Denition and Model of Computation For any integer $n; u > 1$, we let $D(n; u)$ be an MMPHF indexing algorithm for size- n subsets of $[u]$. That is, if $S_{n;u} = \{S \subseteq [u] \text{ s.t. } |S| = n\}$ then for all $S \in S_{n;u}$, $D(S)$ is the MMPHF index for S .

For any xed choice of random bits r , we use D^r to denote the resulting MMPHF with random bits r . Note that for any xed choice of r , D^r is deterministic. For any $S \in S_{n;u}$ and randomness r , dene $d^r(S)$ as the size in bits of the MMPHF index $D^r(S)$. Dene:

$$d(n; u) := \max_{S \in S_{n;u}} \mathbb{E}[d^r(S)] :$$

When n and u are clear, we drop them and refer simply to D and d .

In this denition of size, we are not charging the algorithm for storing its randomness. In other words, the algorithm has access to a tape of random bits chosen independent of the input that it can use for both creating the index as well as answering the queries. Furthermore, we also allow the algorithm unbounded computation time⁴. Thus, the only measure of interest for us is the size of the index. Finally, any deterministic MMPHF in this

³This is an inherently dierent technique than the one used in [EH66] for the shift-graph, as it is known that the fractional chromatic number of the shift-graph is $O(1)$ (see, e.g. [ST11]).

⁴In this (non-uniform) information-theoretical setting, one can remove random bits entirely by increasing the space with $O(\log n + \log \log u)$ bits (see, e.g., Newman's Theorem in communication complexity [New91]), but this extra $O(\log \log u)$ is in

model is simply a randomized MMPHF that ignores its random bits and thus we will only focus on randomized MMPHFs from now on.

2.2 Fractional Colorings A key tool that we use in establishing our lower bound is the notion of a fractional coloring of a graph. We now review the basics of fractional colorings, which we need in our proofs. The results mentioned in this subsection are all standard; see, e.g. [SU11] (we present self-contained proofs of these results in [Appendix A](#) for completeness).

Let $G = (V; E)$ be any undirected graph. A proper coloring of G is any assignment of colors to vertices of G so that no edge is monochromatic. The chromatic number $\chi(G)$ is the minimum number of colors in any proper coloring of G .

Let $\mathcal{I}(G) \subseteq 2^V$ denote the set of all independent sets in G , and for any vertex $v \in V$, denote $\mathcal{I}(G; v)$ as the set of all independent sets that contain the vertex v . A fractional coloring of G is any assignment of $x_I \in [0; 1]^{\mathcal{I}(G)}$ to the independent sets of G satisfying the following constraint:

$$\text{for every vertex } v \in V: \sum_{I \in \mathcal{I}(G; v)} x_I > 1;$$

The value $\|x\|_1$ of a fractional coloring x is given by $\sum_{I \in \mathcal{I}(G)} x_I$. The fractional chromatic number $\chi_f(G)$ is the minimum value of any fractional coloring of G . This quantity can be formalized as a linear program (LP):

$$(2.1) \quad \chi_f(G) := \min_{x \in [0; 1]^{\mathcal{I}(G)}} \sum_{I \in \mathcal{I}(G)} x_I \quad \text{subject to} \quad \sum_{I \in \mathcal{I}(G; v)} x_I > 1 \quad \forall v \in V;$$

Any proper coloring of G with k colors induces a solution x of value k to this LP, where x_I is set to 1 for the independent sets I that correspond to (whole) color classes in the coloring. Thus the LP given by [Eq \(2.1\)](#) is indeed a relaxation of the original coloring problem.

Fact 2.1. For any graph G , $\chi_f(G) \leq \chi(G)$.

It is worth mentioning that at the same time $\chi(G) = O(\log |V(G)| \chi_f(G))$ using the standard randomized rounding argument (we do not use this direction explicitly in our paper).

A primal-dual analysis of the fractional-chromatic-number LP implies the following results. These results are standard but we provide proofs in [Appendix A](#) for completeness.

Proposition 2.2. Let $G_1 = (V_1; E_1)$ and $G_2 = (V_2; E_2)$ be arbitrary graphs. Define $G_1 \sqcup G_2$ as a graph on vertices $V_1 \cup V_2$ and define an edge between vertices $(v_1; v_2)$ and $(w_1; w_2)$ whenever $(v_1; w_1)$ is an edge in G_1 or $(v_2; w_2)$ is an edge in G_2 . Then, $\chi_f(G_1 \sqcup G_2) = \chi_f(G_1) + \chi_f(G_2)$.

[Proposition 2.2](#) allows us to determine χ_f of a product of several graphs by focusing on each individual graph separately.

Proposition 2.3. For any graph $G = (V; E)$,

$$\chi_f(G) = \max_{\text{distribution on } V} \min_{I \in \mathcal{I}(G)} \Pr(v \in I) = 1;$$

[Proposition 2.3](#) provides us with a tool to lower bound χ_f by finding a suitable distribution on the vertices so that no independent set has a significant probability of containing a vertex sampled from this distribution.

3 A Lower Bound for MMPHF via Fractional Colorings

We can now formally state the main theorem of this paper.

general unavoidable (see, e.g. [HT01] and references therein), and can be prohibitive for us when u is sufficiently large. Hence, we still explicitly account for randomized data structures in our lower bound.

Theorem 2 (Formalization of Theorem 1). For any $n; u \in \mathbb{N}^+$ such that $n \geq \frac{\log n}{6} \log u$, and for any MMPHF algorithm $D(n; u)$,

$$d(n; u) = (n \log \log \log u):$$

The rest of the paper presents the proof of Theorem 2. We spend the rest of the section reframing the theorem in terms of the fractional chromatic number of a certain graph associated with the MMPHF problem. We will then show how to lower bound the fractional chromatic number in the next section.

3.1 Conict Graph and its Fractional Chromatic Number Let $m > 1$ be an integer and define $M := 2^{m^2+m}$. Define the graph $G(m) := (V(m); E(m))$ as:

- The vertex set is $V(m) = S_{m;M}$, that is, the size- m subsets of $[M]$. We denote each vertex $v \in V(M)$ by the m -tuple $v := (v_1; \dots; v_m)$ where $0 < v_1 < v_2 < \dots < v_m \leq M$.
- The edge set $E(m)$ is defined as follows. Let $v = (v_1; \dots; v_m)$ and $w = (w_1; \dots; w_m)$ be any two vertices in $V(M)$. Then, there is an edge $(v; w) \in G(m)$ if there exists some pair of indexes $i = j \in [m]$ such that $v_i = w_j$.

We refer to $G(m)$ as the conict graph of m . The following lemma clarifies our interest in this graph by showing that fractional chromatic number of $G(m)$ can be used to lower bound size of any MMPHF (for certain parameters of input).

Lemma 3.1. Let $m > 1$ be an integer and let $M = 2^{m^2+m}$. For any MMPHF $D(m; M)$,

$$d(m; M) \geq (\log_f(G(m)) - 2)/2:$$

Proof. Consider any two vertices $v; w \in G(m)$. If there is an edge between v and w , then there exists an element $z = v_i = w_j; i = j$. Therefore for every choice of randomness r , $D^r(v) = D^r(w)$, because query z must return i on $D^r(v)$ and j on $D^r(w)$. This implies that for every r , the set of vertices v with the same $D^r(v)$ form an independent set in $G(m)$ (and the collection of these sets is a coloring of $G(m)$). We use I^r to denote these independent sets in $G(m)$ for this choice of r .

On the other hand, by Proposition 2.3, there exists a distribution μ on $V(m)$ such that

$$(3.2) \quad \chi(G(m)) = \min_{I \in \mathcal{I}(G(m))} \Pr(v \in I) = 1:$$

Let us use that distribution. Under this distribution, by the definition of d ,

$$d = d(m; M) = \max_{v \in V(m)} \mathbb{E}_r[d^r(v)] \geq \mathbb{E}_v \mathbb{E}_r[d^r(v)] = \mathbb{E}_v \mathbb{E}_r[d^r(v)]:$$

An averaging argument now implies that there exists a choice r of random bits such that

$$\mathbb{E}_v d^r(v) \leq d:$$

By Markov's inequality, with probability at least $1/2$, for v , we have that $d^r(v) \leq 2d$.

Recall that $D^r(v)$ corresponds to an independent set in I_r . Moreover, there can be at most $2^{2d+1} - 2$ independent sets I in I_r such that for all $v \in I$, $d^r(v) \leq 2d$; this is because there are at most $2^{2d+1} - 2$ choices for $D^r(v)$ across all $v \in V(m)$ that can use up to $2d$ bits in their index (as the number of non-empty binary strings of length at most $2d$ is $2^{2d+1} - 2$). Since a random v belongs to one of these $2^{2d+1} - 2$ independent sets with probability at least half, we necessarily have some independent set $I \in I_r$ where

$$\Pr(v \in I) \geq \frac{1}{2(2^{2d+1} - 2)} > \frac{1}{2^{2d+2}}:$$

Plugging in this bound in Eq (3.2), we have,

$$\chi(G(m)) \leq 2^{2d+2};$$

which implies that $d \geq (\log_f(G(m)) - 2)/2$, concluding the proof. ■

Lemma 3.1 reduces our task of proving **Theorem 2** to establishing a lower bound on $(G(m))$. This will be accomplished by the following lemma, which we prove in **Section 4**.

Lemma 3.2. There is an absolute constant $\epsilon > 0$ such that for every sufficiently large $m > 1$,

$$\chi_f(G(m)) > m^\epsilon.$$

By plugging in the lower bound of $\chi_f(G(m))$ from **Lemma 3.2** inside **Lemma 3.1**, we get that for any sufficiently large $n > 1$ and universe size $u = 2^{n^{2+\epsilon}}$, the lower bound on the MMPHF problem is $(n \log \log \log u)$ as $\log n = (\log \log \log u)$ here.

Thus **Lemmas 3.1** and **3.2** can be combined to prove **Theorem 2** modulo a serious caveat: the lower bound only holds for instances of the problem wherein the universe size is larger than doubly exponential in n , which is admittedly not the most interesting setting of the parameters. In the next subsection, we use a simple graph product argument (plus **Proposition 2.2**) to extend this lower bound to the whole range of parameters u considered by **Theorem 2**.

3.2 Extending the MMPHF Lower Bound to Small Universes For every integers $m, k > 1$, define $G(m; k) = (V(m; k); E(m; k))$ as the k -fold conict graph where the vertex set $V(m; k)$ is the set of all size- m subsets of $[k+1 : M+k]$ for $M := 2^{m^{2+\epsilon}}$ defined earlier, and the edge set $E(m; k)$ is defined as in normal conict graphs. (Thus $G(m; 0) = G(m)$.)

Furthermore, for every integer $m, k > 1$, we define the k -fold conict graph, denoted by $G^k(m)$, as the graph:

$$G^k(m) = (V^k(m); E^k(m)) := G(m; 0) \times G(m; M) \times G(m; 2M) \times \dots \times G(m; (k-1)M);$$

where \times denotes the graph product in **Proposition 2.2**. The direct interpretation of the nodes of $V^k(m)$ is a product of tuples from disjoint ranges, but we can also interpret it as a single tuple of length km . This way, $G^k(m)$ is a subset of the conict graph on km -size subsets of $[kM]$ and it makes sense to compute $D(v)$ for any $v \in V^k(m)$.

Therefore, by **Lemma 3.1**, we again have a lower bound of $(\log_f(G^k(m)))$ for MMPHF on tuples of length $n = km$ from a universe of size $u = kM$.

By **Proposition 2.2**, combined with **Lemma 3.2**, we have,

$$\log_f(G^k(m)) = \sum_{i=1}^k \log_f(G(m; i-1)) = k \log_f(G(m)) > k \log_f(m) = (n \log m); i=1$$

where the second equality is because $\chi_f(G(m; i-1)) = \chi_f(G(m))$ for all $i \in [k]$, as these graphs are all isomorphic to each other. Consider a choice of

$$m = (\log \log n)^{1/\epsilon} \text{ and } k = n/(\log \log n)^{1/\epsilon};$$

which in turn gives us

$$u = k \cdot 2^{m^{2+\epsilon}} \cdot k \cdot 2^{m^3} = \frac{n}{(\log \log n)^{1/\epsilon}} \cdot 2^{p_{\log \log n} \cdot \frac{1}{\epsilon}} \cdot n \cdot 2^{p_{\log n}}.$$

By the above equation, we have a lower bound of $(n \log \log \log u)$ for MMPHF given that in this case, $\log m = (\log \log \log u)$. Thus, so far, we have proven **Theorem 2** on both its boundary cases, namely, when $u = n \cdot 2^{p_{\log n}}$ and when $u = 2^{n^{2+\epsilon}}$. The proof can now be extended to the full range of the parameters in the middle by re-parameterizing k appropriately; see **Appendix B** for the complete argument.

We conclude that in order to finish the proof of **Theorem 2**, we need only establish **Lemma 3.2**.

4 Fractional Chromatic Number of Conict Graphs

In this section, we establish a lower bound on the fractional chromatic number of the conict graph $G(m)$ for any (large enough) $m > 1$, and thereby prove **Lemma 3.2**.

Proposition 2.3 gives us a clear path for proving the lower bound on $f(G(m))$ in **Lemma 3.2**: we can design a distribution on vertices of $V(m)$ and then, for every independent set $I \subseteq I(G(m))$, we can upper bound the probability that v sampled from \mathcal{D} belongs to I . As f in **Proposition 2.3** is maximum over all possible distributions, our distribution provides a lower bound for $f(G(m))$.

To continue, we need the following interpretation of the (maximal) independent sets in $G(m)$.

Observation 4.1. Any maximal independent set I in $G(m)$ can be uniquely identified by a function $f_I : [M] \rightarrow [m]$ such that for every vertex $v = (v_1; \dots; v_m) \in V(m)$, we have $f_I(v_i) = i$.

Proof. Consider any two vertices $v, w \in I$. Since there is no edge between $v = (v_1; \dots; v_m)$ and $w = (w_1; \dots; w_m)$ in $G(m)$, whenever $v_i = w_j$, we necessarily have that $i = j$. Thus, any element of $e \in [M]$ can only appear in a single index $i_e \in [m]$ throughout all vertices $v \in I$ (or does not appear at all in v). We can thus define $f_I(e)$ to be i_e , giving us a function f_I with the desired property.

We now show that f_I uniquely identifies I . Define I^0 as set of vertices $v = (v_1; \dots; v_m) \in V(m)$ satisfying $f_I(v_i) = i$ for all $i \in [m]$. I^0 is an independent set satisfying $I \subseteq I^0$. Since I is assumed to be maximal, it follows that $I = I^0$, meaning that we recover I from f_I . ■

Observation 4.1 allows us to reduce **Lemma 3.2** to the following lemma about m -tuples of increasing integers. Proving **Lemma 4.2** is the main technical contribution of our work.

Lemma 4.2. There is an absolute constant $\epsilon > 0$ such that for any sufficiently large $m > 1$ and $M = 2^{m^{m^2+m}}$, the following is true. There exists a distribution on m -tuples of increasing numbers $X_1 < \dots < X_m$ from $[M]$ such that for any function $f : [M] \rightarrow [m]$,

$$\Pr_{(X_1; \dots; X_m)} \left(\exists i \in [m] : f(X_i) = i \right) \leq m^{-\epsilon}.$$

Before proving **Lemma 4.2**, we show how it implies **Lemma 3.2**.

Proof of Lemma 3.2 (assuming **Lemma 4.2**). Any choice of $(X_1; \dots; X_m)$ in **Lemma 4.2** can be mapped to a unique vertex $v \in V(m)$ and vice versa. Thus, $(X_1; \dots; X_m)$ induces a distribution on vertices $V(m)$: sample $(X_1; \dots; X_m)$ and return the vertex $v = (v_1; \dots; v_m)$ where $v_i = X_i$ for all $i \in [m]$. Moreover, for any maximal independent set $I \subseteq I(G)$, by **Observation 4.1**, the vertex corresponding to $(X_1; \dots; X_m)$ belongs to I if $f_I(X_i) = i$ for all $i \in [m]$. Thus,

$$\Pr_v (v \in I) = \Pr_{(X_1; \dots; X_m)} \left(\exists i \in [m] : f_I(X_i) = i \right) \leq m^{-\epsilon}.$$

As every independent set of $G(m)$ is a subset of some maximal independent set, the upper bound continues to hold for every independent set in $G(m)$.

By **Proposition 2.3**,

$$f(G(m)) > \min_{I \subseteq I(G(m))} \Pr(v \in I)^{-1} > m^\epsilon;$$

concluding the proof. ■

The rest of the section proves **Lemma 4.2**. We start with a high-level overview in **Section 4.1**. We then define the distribution that we will use for the proof of **Lemma 4.2** (**Section 4**) and analyze it to establish **Lemma 4.2** (**Section 4.3**). The probability distribution that we construct in these sections should be viewed intuitively as a "hard" input distribution on inputs to the MMPHF problem (in the spirit of Yao's minimax principle).

4.1 A High-Level Overview of the Proof The proof of **Lemma 4.2** is quite dense and requires both a highly delicate probability distribution and several intricate technical arguments. Thus, before getting into the details of this proof, we provide a (very) high-level overview of the logic behind it. In order to convey the intuition, we omit many details from this subsection, instead limiting ourselves to an informal discussion.

The distribution in **Lemma 4.2** is roughly as follows: we start with a "window" Win_1 which is the interval $[1 : M]$, and then sample X_1 uniformly at random from Win_1 . We then pick window Win_2 to be $[X_1 + 1 : X_1 + w_2]$

for an integer $w_2 > 1$ chosen randomly from a carefully designed distribution. Similarly to before, X_2 will be chosen uniformly from Win_2 . We continue like this by picking a new window $\text{Win}_i = [X_{i-1} + 1 : X_{i-1} + w_i]$ for each $i \in [m]$ by sampling each w_i from a distribution that is constructed based on $(w_1; \dots; w_{i-1})$, and then sampling X_i from Win_i . Note that, by design, we will satisfy $X_1 < X_2 < \dots < X_m$.

The key property that this distribution achieves can be explained informally as follows. For any index $i \in [m]$, there is a recursive partitioning of the window Win_i into "dense" and "sparse" intervals, where an interval $I \subseteq \text{Win}_i$ is dense (with respect to the function f and the index i) if at least an $(1/m)$ fraction of entries $j \in I$ satisfy $f(j) = i$, and otherwise I is sparse. The central property that our distribution ensures is that, if the random choice of X_i places it in a dense interval, then (with very high probability) the final window Win_m will itself end up being dense (i.e., for at least a $(2/m)$ fraction of $j \in \text{Win}_m$, $f(j) = i$).

Establishing this property is quite challenging and involves defining the distribution of w_i 's in a highly non-uniform manner (in terms of their values); this is also the source of the doubly exponential dependence of range M on the number of indices m . We postpone the details on how this property can be achieved to the actual proof and focus on why it is a useful property for us.

The analysis of the distribution now uses the property in a potential-function style argument. For each X_i , it is either sampled from a sparse interval or a dense one. If X_i is sampled from a sparse interval I , then no matter the past iterations, the probability that $f(X_i) = i$ is at most $(2/m)$, since at most $(2/m)$ fraction of I can have value $f(j) = i$ by the definition of it being sparse. On the other hand, if X_i is chosen from a dense interval, then at least a $(2/m)$ fraction of entries of Win_m should be mapped to i by f as well (by our property). Seeing Win_m as a potential function now, we have that this latter step can only happen for $(m/2)$ iterations $i \in [m]$; indeed, each time that this happens for some i , we commit some $(2/m)$ fraction of indices $j \in \text{Win}_m$ to having $f(j) = i$, and these sets of indices must be disjoint. As a result, we have that at least $(m/2)$ iterations $i \in [m]$ sample X_i from a sparse interval. Thus,

$$\Pr(f(X_1) = 1; \dots; f(X_m) = m) \leq \prod_{i: X_i \text{ chosen from a sparse interval}} \Pr(f(X_i) = i \mid f(X_1) = 1; \dots; f(X_{i-1}) = (i-1))$$

$$\leq \prod_{i=1}^m \frac{1}{m} = \frac{1}{m^m} = \frac{1}{(m)^m};$$

as desired for the proof of [Lemma 4.2](#).

The main challenge in formalizing the above argument is the design and analysis of the distribution so that the property discussed above holds. Note also that the property cannot hold deterministically; another challenge is to show that it holds with such high probability that the risk of the property ever failing (across the entire construction) can be ignored.

4.2 The Hard Input Distribution in [Lemma 4.2](#) The distribution is defined as follows.

(i) Let $k = m^m$, $S_0 = k^{m+1}$, and $X_0 = 0$.

(ii) For $i = 1$ to m :

(a) Sample two random numbers Y_i from $[2^{S_{i-1}}]$ and Z_i from $[k-1]$ uniformly at random.

(b) Define the random variables of iteration i as:

$$X_i = X_{i-1} + Y_i \quad \text{and} \quad S_i = S_{i-1} - k^{m-i+1} Z_i;$$

(iii) Return $(X_1; \dots; X_m)$ as the resulting random variables.

To avoid ambiguity, we use lower case letters $(s_i; x_i; y_i; z_i)$ to denote realizations of random variables $(S_i; X_i; Y_i; Z_i)$ for $i \in [m]$.

We have the following basic observation on the range of numbers created in this distribution.

Observation 4.3. Every choice of $(X_1; \dots; X_m)$ and $(S_1; \dots; S_m)$ satisfy the following properties:

- (i) "Monotonicity": for all $i \in [m]$, $X_i \geq X_{i-1}$ and $S_i \leq S_{i-1} - m^m$ (and S_i, X_i are integers).
- (ii) "Boundedness": for every $i \in [m]$, $X_m \leq X_i + (m-i)2^{S_i}$ and $S_m \geq S_i - k^{m-i+1} > 0$.

Proof. Monotonicity of X_i 's holds as Y_i 's are positive. Monotonicity for S_i 's holds because Z_i 's are positive and $k^{m-i+1} > k^{m-m+1} = k = m^m$, meaning that we always have $S_i \leq S_{i-1} - m^m$.

For part (ii), we have,

$$X_m = X_i + \sum_{j=i+1}^m Y_j \leq X_i + \sum_{j=i+1}^m 2^{S_{j-1}} \leq X_i + (m-i)2^{S_i};$$

which proves the boundedness of X_i 's. For S_i 's,

$$\left(\text{as } \prod_{j=i}^m k^{-j} \leq \prod_{j=i}^m k^{-j} = k^{-i+1} (k-1)^{-1} \right)$$

$$S_m = S_i - \sum_{j=i+1}^m k^{m-j+1} Z_j \geq S_i - k^m (k-1) \prod_{j=i}^m k^{-j} \geq S_i - k^{m-i+1};$$

Finally, by this bound, we have $S_m \geq S_0 - k^{m+1} > 0$ as $S_0 = k^{m+1}$. ■

When discussing $(X_1; \dots; X_m)$, we will also need some further denitions:

- For any realization $(s_{<i}; x_{<i})$, we dene the window of iteration $i \in [m]$, $\text{Win}_i := \text{Win}_i(s_{<i}; x_{<i})$, as the support of the random variable X_i conditioned on $(s_{<i}; x_{<i})$, i.e.,

$$\text{Win}_i := \text{Win}_i(s_{<i}; x_{<i}) = [x_{i-1} + 1 : x_{i-1} + 2^{S_{i-1}}];$$

Notice that $|\text{Win}_i(s_{<i}; x_{<i})| = 2^{S_{i-1}}$ and Win_i is determined by $(s_{<i}; x_{<i})$.

- Similarly, for any xed choice of $(s_{<i}; x_{<i})$, consider the following numbers:

$$(4.3) \quad w_{i;j} := 2^{S_{i-1} - j(k^{m-i+1})} \quad \text{for all } j \in \{0; \dots; k-1\};$$

This way, $j \in \text{Win}_{i+1}(s_{<i}; x_{<i})$ is chosen uniformly at random from $\{w_{i;1}; \dots; w_{i;k-1}\}$ (depending solely on the choice of $Z_i \in [k-1]$ which also determines S_i). Moreover, the ratio of $w_{i;j}$ and $w_{i;j+1}$ is xed for any $j \in \{0; \dots; k-1\}$ and we dene this quantity as

$$(4.4) \quad r_i := 2^{k^{m-i+1}} = \frac{w_{i;j}}{w_{i;j+1}} \quad \text{for any } j \in \{0; \dots; k-1\};$$

Observation 4.4. For any xed $(s_{<i}; x_{<i})$, the supports of random variables $j \in \text{Win}_{i+1}; \dots; j \in \text{Win}_m$ are subsets of the interval $[2^{m-m} w_{i;Z_i+1} : w_{i;Z_i}]$.

Proof. By denition,

$$j \in \text{Win}_{i+1} \implies j = 2^{S_i} = 2^{S_{i-1} - k^{m-i+1} Z_i} = w_{i;Z_i};$$

Moreover, by **Observation 4.3**, for any $j \in \{i+1; \dots; m\}$, we have $j \in \text{Win}_j \subseteq j \in \text{Win}_{i+1}$. Thus each of these windows can have length at most $w_{i;Z_i}$, proving the upper bound side.

For the lower bound, for any $j \in \{i+1; \dots; m\}$, we have,

$$\begin{aligned} \text{(by parts (i); (ii) of Observation 4.3)} \quad j \in \text{Win}_j &\geq j \in \text{Win}_m = 2^{S_m-1} > 2^{S_i - k^{m-i+1} + m^m} \\ &= 2^{m^m} 2^{S_i} 2^{-k^{m-i+1}} = 2^{m^m} w_{i;Z_i} r_i^{-1} = 2^{m^m} w_{i;Z_i+1}; \end{aligned}$$

This concludes the proof. ■

We need one final definition for now:

- For the function $f : [M] \rightarrow [m]$, we define the density of index $i \in [m]$ in f over a window Win , denoted by $\text{density}_f(\text{Win}; i)$, as

$$\text{density}_f(\text{Win}; i) := \frac{|\{j \in \text{Win} : f(j) = i\}|}{|\text{Win}|},$$

namely, the fraction of entries of the window that are equal to i .

Observation 4.5. For any choice of $(s_{<i}; x_{<i})$, we have,

$$\Pr(f(X_i) = i \mid s_{<i}; x_{<i}) = \text{density}_f(\text{Win}_i(s_{<i}; x_{<i}); i):$$

Proof. Conditioned on $(s_{<i}; x_{<i})$, X_i is chosen uniformly at random from $\text{Win}_i(s_{<i}; x_{<i})$. The observation therefore follows from the definition of $\text{density}_f(\text{Win}_i(s_{<i}; x_{<i}); i)$: ■

4.3 Analysis of the Hard Distribution { Proof of Lemma 4.2 We prove Lemma 4.2 by individually considering each iteration in the distribution.

Lemma 4.6. For any iteration $i \in [m-1]$ and conditioned on any choice of $(s_{<i}; x_{<i})$, at least one of the following two conditions is true:

$$(i) \Pr(f(X_i) = i \mid s_{<i}; x_{<i}) \geq \frac{101}{m} \quad \text{or} \quad (ii) \Pr(\text{density}_f(\text{Win}_m; i) < \frac{2}{m} \mid s_{<i}; x_{<i}) < \frac{1}{k^{1/3}}.$$

The guarantee in Lemma 4.6 does not apply to the last iteration (omitted for technical reasons).

The main bulk of this section is to prove Lemma 4.6. We then show at the end of the section that this lemma easily implies Lemma 4.2. To continue, we need some definitions.

Definition 4.7. The window-tree of iteration $i \in [m]$ for $(s_{<i}; x_{<i})$, denoted by $T_i := T(s_{<i}; x_{<i})$, is the following rooted tree with $k+1$ levels (the root is at level 0):

- (i) Every non-leaf node v of the tree has r_v many child-nodes.
- (ii) Every node v at a level $\ell \in \{0, \dots, k\}$ is associated with a window $\text{Win}(v)$ of length $w_{i,\ell}$.
- (iii) The root r is associated with the window $\text{Win}(r) := \text{Win}_i(s_{<i}; x_{<i})$. The windows associated with child-nodes of a node v at level ℓ partition $\text{Win}(v)$ of length $w_{i,\ell}$ into equal-size windows of length $w_{i,\ell+1}$ (recall that $r_v = w_{i,\ell}/w_{i,\ell+1}$ child-nodes). Moreover, the left most child-node receives the window in the partition with the smallest starting point, the next child-node on the right receives the next window with smallest part, and so on.
- (iv) The density of a node v with respect to any function $f : [M] \rightarrow [m]$ is defined as

$$\text{density}_f(v) := \text{density}_f(\text{Win}(v); i):$$

One way we use the window-tree in our analysis is to consider the process of sampling X_i (which is uniform over $\text{Win}_i(s_{<i}; x_{<i})$ at this stage) as traversing the window-tree via a root-to-leaf path. This is formalized in the following observation.

Observation 4.8. The distribution of X_i conditioned on $(s_{<i}; x_{<i})$ can be alternatively seen as: (i) Sample a root-to-leaf path v_0, v_1, \dots, v_k where v_0 is the root of T_i and where each $v_{\ell+1}$ is a child-node of v_ℓ chosen uniformly at random; then, (ii) sample X_i uniformly at random from $\text{Win}(v_k)$. We refer to v_0, v_1, \dots, v_k as the sampling path of X_i .

Proof. X_i is distributed uniformly over Win_i and leaf-nodes of T_i form an equipartition of Win_i . ■

In addition, we define a pruning procedure for any window-tree T_i as follows.

Definition 4.9. Fix a function $f : [M] \rightarrow [m]$ and a window-tree T_i for some $i \in [m]$. We say that a node $v \in T_i$ is sparse if

$$\text{density}_f(v) \leq \frac{100}{m}.$$

Consider the following procedure for pruning T_i : Start from the root down to the leaf-nodes and prune any sparse node, as well as the whole subtree rooted at that node. We refer to a sparse node that was pruned on its own (i.e., any node that is sparse and has no sparse ancestors) as a directly pruned node and to other pruned nodes as indirectly pruned.

Finally, for $\ell \in \{0, \dots, k\}$, define p_ℓ as the fraction of directly pruned nodes at level ℓ of the tree over all level- ℓ nodes that are not indirectly pruned.

It is worth noting that pruning is deterministic conditioned on $(s_{<i}; x_{<i})$.

With these definitions, we can now start proving **Lemma 4.6**. This will be done by considering some different cases handled by the following claims. The first (and easiest) case is when most nodes of the window-tree are pruned, in which case we achieve property (i) of **Lemma 4.6**.

Claim 4.10 (Case I: "Many Directly Pruned Nodes"). Suppose

$$\sum_{\ell=0}^k (1 - p_\ell) \leq \frac{1}{m}.$$

Then, for any choice of $(s_{<i}; x_{<i})$,

$$\Pr_{X_i}(f(X_i) = i \mid s_{<i}; x_{<i}) \leq \frac{101}{m}.$$

Proof. Let W_{rem} denote the subset of Win_i that remains after removing windows of all pruned leaf-nodes from Win_i . We have that

$$|W_{\text{rem}}| = \frac{\# \text{ leaf-nodes of } T_i \text{ that are not pruned}}{\# \text{ leaf-nodes of } T_i} |\text{Win}_i| = \sum_{\ell=0}^k (1 - p_\ell) |\text{Win}_i| \leq \frac{|\text{Win}_i|}{m},$$

where the second equality is because at each level ℓ of the tree, the number of not pruned nodes drops by a factor of $(1 - p_\ell)$ by the definition of p_ℓ .

Let DP denote the set of all nodes in the tree T_i that were directly pruned. Note that the windows $\text{Win}()$ for $v \in \text{DP}$ partition $\text{Win}_i \setminus W_{\text{rem}}$. This implies that

(by the definition of $\text{density}_f()$ function)

$$\text{density}_f(\text{Win}_i; i) = \frac{1}{|\text{Win}_i|} |W_{\text{rem}}| \text{density}_f(W_{\text{rem}}; i) + \sum_{v \in \text{DP}} \text{density}_f(v) |\text{Win}(v)| \quad (1)$$

$\text{density}_f(v) \leq \frac{100}{m}$ by the definition of sparsity, and $\text{density}_f(W_{\text{rem}}; i) \leq 1$

$$\begin{aligned} &\leq \frac{1}{|\text{Win}_i|} |W_{\text{rem}}| + \sum_{v \in \text{DP}} \frac{100}{m} |\text{Win}(v)| \\ &\quad (\text{as } |W_{\text{rem}}| = |\text{Win}_i| \leq \frac{1}{m} \text{ as established above, and } \sum_{v \in \text{DP}} |\text{Win}(v)| \leq |\text{Win}_i|) \\ &\leq \frac{1}{m} + \frac{100}{m} = \frac{101}{m}; \end{aligned}$$

By **Observation 4.5**, we have,

$$\Pr_{X_i}(f(X_i) = i \mid s_{<i}; x_{<i}) = \text{density}_f(\text{Win}_i; i) \leq \frac{101}{m};$$

concluding the proof. \blacksquare

We now consider the complementary case, while also taking the randomness of Z_i into account. Recall that Z_i is uniform over $[k-1]$ and that $j \text{Win}_{i+1} j = w_{i;Z_i}$. For any fixed realization z_i of Z_i , recall the sampling-path-based process of sampling X_i outlined in **Observation 4.8**. Consider the rest z_i vertices in this path, namely, $o; \dots; z_{i-1}$ that start from the root and end at a level z_{i-1} node of T_i .

Define $E(s_{<i}; x_{<i}; z_i; X_i)$ to be the event that none of the nodes in $o; \dots; z_{i-1}$ are pruned. Event $E(s_{<i}; x_{<i}; z_i; X_i)$ depends only on the choice of X_i (to traverse the root-to-leaf path), and is conditioned on $s_{<i}; x_{<i}$ (which determine the window-tree T_i) and z_i (which determines the level of the tree that we focus on). To avoid clutter, when it is clear from the context, we refer to this event simply by E_i .

We partition the remaining cases based on whether or not the event E_i happens.

Claim 4.11 (Case II: \A Pruned Node on the Sampling Path"). Fix any choice of z_i and $(s_{<i}; x_{<i})$. In the case that the event E_i does not happen, we have,

$$\Pr_{X_i}(f(X_i) = i \mid s_{<i}; x_{<i}; z_i; \overline{E_i}) \leq \frac{100}{m}.$$

Proof. After conditioning on $(s_{<i}; x_{<i}; z_i)$, the event E_i is only a function of the sampling process of X_i outlined in **Observation 4.8**. Assuming E_i does not happen, we know that there exists a unique node j on the path $o; \dots; z_{i-1}$ such that j is sparse and is directly pruned. By additionally conditioning on the subpath $o; \dots; j$, we have that X_i is chosen uniformly at random from $\text{Win}(j)$ at this point. Thus,

$$\Pr_{X_i}(f(X_i) = i \mid s_{<i}; x_{<i}; z_i; \overline{E_i})$$

(as these subpaths partition all possible choices for E_i to not happen)

$$= \sum_{\substack{j=0 \\ (o; \dots; j): \\ \text{directly pruned}}}^{z_{i-1}} \Pr(f(X_i) = i \wedge (o; \dots; j) \text{ is on the sampling path} \mid s_{<i}; x_{<i}; z_i; \overline{E_i})$$

(as X_i is chosen uniformly from $\text{Win}(j)$ under these conditions)

$$= \sum_{\substack{j=0 \\ (o; \dots; j): \\ \text{directly pruned}}}^{z_{i-1}} \Pr((o; \dots; j) \text{ is on the sampling path} \mid s_{<i}; x_{<i}; z_i; \overline{E_i}) \frac{|\text{Win}(j)|}{|\text{Win}(j)|} \frac{f(j)}{|\text{Win}(j)|}$$

(by the definition of density_f)

$$= \sum_{\substack{j=0 \\ (o; \dots; j): \\ \text{directly pruned}}}^{z_{i-1}} \Pr((o; \dots; j) \text{ is on the sampling path} \mid s_{<i}; x_{<i}; z_i; \overline{E_i}) \text{density}(\text{Win}(j); i)$$

(as j needs to be sparse to be directly pruned)

$$\leq \sum_{\substack{j=0 \\ (o; \dots; j): \\ \text{directly pruned}}}^{z_{i-1}} \Pr((o; \dots; j) \text{ is on the sampling path} \mid s_{<i}; x_{<i}; z_i; \overline{E_i}) \frac{100}{m}$$

This can now be further upper bounded by $100/m$ as the probability terms are summing over all disjoint events that can lead to $\overline{E_i}$ (conditioned on this event) and thus add up to one. ■

Finally, we have the following case which handles the situation when E_i happens. The following claim is the heart of the proof.

Claim 4.12 (Case III: \No Pruned Nodes on the Sampling Path"). Fix any choice of z_i and $(s_{<i}; x_{<i})$. In the case that the event E_i happens, we have,

$$\Pr_{X_i}(\text{density}_f(\text{Win}_m; i) < \frac{2}{m} \mid s_{<i}; x_{<i}; z_i; E_i) \leq 4 \cdot p_{z_i} + p_{z_i+1} + \frac{m}{r_i} \Pr$$

Proof. Throughout this proof, we always condition on $s_{<i}; x_{<i}; z_i$; and $E_i = E(s_{<i}; x_{<i}; z_i; X_i)$ and so may not mention this explicitly in the probability terms. This is the information we have so far:

- None of the nodes $o_{i-1}; \dots; z_{i-1}$ on the sampling path are pruned as we conditioned on the event E_i (although z_{i-1} is still a random variable and is not fixed yet just by these conditions).
- Window Win_m is going to have size at least $2^{m-m} w_{i;z_i+1}$ and at most $w_{i;z_i}$ by **Observation 4.4**.

- By **Observation 4.3**,

$$(\text{by the definition of } w_{i;z_i} = 2^{S_i}) \quad X_m \leq X_i + (m-i) 2^{S_i} = X_i + (m-i) w_{i;z_i} :$$

- Win_m starts at X_{m-1} and ends at $X_{m-1} + j\text{Win}_m$. We can think of the process of sampling Win_m as first sampling $j\text{Win}_m$, then sampling the offset $O_{i;m} := X_{m-1} - X_i = \sum_{j=i+1}^{m-1} Y_j$ conditioned on $j\text{Win}_m$, and then sampling X_i conditioned on $O_{i;m}$ and $j\text{Win}_m$.
- We further have that X_i conditioned on $O_{i;m}$ and $j\text{Win}_m$ is still uniform over $\text{Win}(z_{i-1})$. This is because $j\text{Win}_m$ is only a function of $Z_{i+1}; \dots; Z_m$, and $X_{m-1} - X_i$ is only a function of $Y_{i+1}; \dots; Y_{m-1}$, while X_i is only a function of Y_i ; namely, Y_i is independent of $Y_{i+1}; \dots; Y_m$ and $Z_{i+1}; \dots; Z_m$ and is chosen uniformly from $[2^{S_i}]$ (recall that $i < m$ in this lemma).

In the following, we condition on any fixed choice of offset $o_{i;m}$ for $O_{i;m}$ and on $j\text{Win}_m$. We have already established that

$$(4.5) \quad 2^{m-m} w_{i;z_i+1} \leq j\text{Win}_m \leq w_{i;z_i} \quad \text{and} \quad o_{i;m} \leq (m-i) w_{i;z_i} :$$

Moreover, the distribution of Win_m conditioned on $o_{i;m}; j\text{Win}_m$ (and $s_{<i}; x_{<i}; z_i; E_i$ that we always condition on in this proof), is $X_i + o_{i;m}$ for X_i chosen randomly from $\text{Win}(z_{i-1})$. Also, given that $o_{i;m} \leq (m-i) w_{i;z_i}$ while $j\text{Win}_m(z_{i-1}) = w_{i;z_i+1} = r_i w_{i;z_i}$ and $r_i = 2^{k-m-i+1} > 2^k$ as $i \leq m$, the distribution of X_i and $X_i + o_{i;m}$ are quite close to each other modulo a negligible factor. Thus, for intuition, we can think of X_i itself as the distribution of starting point for Win_m in this context (although we will of course take this difference into account explicitly in the proof). We now use this information to prove the claim. To simplify the exposition, we are going to separate the analysis based on level z_i and level z_{i+1} of the window-tree.

Analysis on level z_i of the window-tree. Firstly, since $j\text{Win}_m \leq w_{i;z_i}$, and each node at level z_i of the window-tree T_i has a window of length $w_{i;z_i}$, we get that Win_m intersects with windows of at most two consecutive nodes at level z_i of T_i , which are solely determined by the choice of X_i . We use $z_1(X_i)$ and $z_2(X_i)$ to denote these two nodes with z_1 being the one where the starting point of Win_m , namely, $X_i + o_{i;m}$, lies in, and $z_2(X_i)$ being the one containing the endpoint $X_i + o_{i;m} + j\text{Win}_m$ (note that it is possible that $z_2 = z_1$).

We prove that with high probability, neither of these nodes are pruned. Let us focus on $z_1(X_i)$ first (the analysis is almost identical for $z_2(X_i)$ and we can then apply a union bound). For any $i' \in \{0; \dots; k-1\}$, let $P(i')$ (resp. $\text{DP}(i')$) denote the set of pruned (resp. directly pruned) nodes at level i' of T_i ; similarly, for a node $z \in T_i$, let $P(z)$ (resp. $\text{DP}(z)$) denote the set of child-nodes of z that are pruned (resp. directly pruned). For any fixed choice of z_{i-1} on the sampling path of X_i ,

(as z_1 is in level z_i and $P(z_i)$ is the set of all pruned nodes of this level)

$$\Pr(z_1(X_i) \text{ is pruned} \mid z_{i-1}) = \frac{\sum_{X_i} \Pr(z_1(X_i) \text{ is pruned} \mid z_{i-1}, X_i)}{\sum_{X_i} \Pr(z_1(X_i) \text{ is pruned} \mid z_{i-1}, X_i)}$$

(by partitioning the nodes in level z_i between child-nodes of z_{i-1} and remaining ones)

$$= \frac{\sum_{X_i} \Pr(z_1(X_i) \text{ is pruned} \mid z_{i-1}, X_i)}{\sum_{X_i} \Pr(z_1(X_i) \text{ is pruned} \mid z_{i-1}, X_i)} + \frac{\sum_{X_i} \Pr(z_1(X_i) \text{ is pruned} \mid z_{i-1}, X_i)}{\sum_{X_i} \Pr(z_1(X_i) \text{ is pruned} \mid z_{i-1}, X_i)}$$

$$(4.6) \quad \leq jP(z_{i-1}) + \frac{1}{r_i} + (m-i) \frac{1}{r_i}$$

where the last inequality holds because of the following reasoning. Firstly, the probability that $x_1(X_i)$ is equal to any fixed node at level z_i is at most $1/r_i$. This is because

$$\Pr(x_1(X_i) = j_{z_i-1}) = \Pr(X_i + o_{i,m} \in \text{Win}(j_{z_i-1})) \leq \frac{j \cdot \text{Win}(j_{z_i-1})}{j \cdot \text{Win}(j_{z_i-1})} = \frac{1}{r_i}$$

because X_i is chosen uniformly from $\text{Win}(j_{z_i-1})$, and $j \cdot \text{Win}(j_{z_i-1}) = j \cdot \text{Win}(j_{z_i-1}) = r_i$ as j_{z_i-1} is at level z_i . This immediately implies the first term in the RHS of Eq (4.6). For the second term, for $x_1(X)$ to intersect with a node not in the subtree of j_{z_i-1} , we need to have $X_i + o_{i,m} \notin \text{Win}(j_{z_i-1})$, while we know $X_i \in \text{Win}(j_{z_i-1})$. As $o_{i,m} \leq (m-i)w_{i,z}$ by Eq (4.5), and any node at level z_i has a window of length $w_{i,z}$, we get that there are most $(m-i)$ choices of x_1 outside child-nodes of j_{z_i-1} that can also become $x_1(X_i)$. The second part of RHS in Eq (4.6) now follows from this and the upper bound of $1/r_i$ on the probability of each node.

Finally, by taking the expectation over the choice of j_{z_i-1} ,

(by the law of total probability, over the choice of j_{z_i-1} in the sampling path)

$$\Pr(x_1(X_i) \text{ is pruned}) = \mathbb{E}_{j_{z_i-1}} \Pr(x_1(X_i) \text{ is pruned} | j_{z_i-1})$$

(by Eq (4.6))

$$\begin{aligned} & \leq \mathbb{E}_{j_{z_i-1}} \left[\frac{j \cdot P(j_{z_i-1})}{r_i} + \frac{(m-i)}{r_i} \right] \\ & = p_{z_i} + \frac{(m-i)}{r_i}; \end{aligned}$$

where in the final equality, we used the fact that j_{z_i-1} is chosen from non-pruned nodes (by conditioning on E_i), and thus $j \cdot P(j_{z_i-1}) = r_i$ is the fraction of pruned nodes over all not indirectly pruned nodes at level z_i , which by definition is p_{z_i} .

Doing the same exact analysis, we can bound the probability that $x_2(X_i)$ is pruned also as

$$\Pr(x_2(X_i) \text{ is pruned}) \leq p_{z_i} + \frac{(m-i) + 1}{r_i};$$

where the +1 term in the RHS compared to the one for x_1 comes from the fact that $x_2(X_i)$ can have $(m-i+1)$ choices outside subtree of j_{z_i-1} (because we are now considering $X_i + o_{i,m} + j \cdot \text{Win}_m \leq X_i + (m-i+1)w_{i,z}$ instead). By the union bound on the probabilities for $x_1(X_i)$ and $x_2(X_i)$,

$$(4.7) \quad \Pr(\text{either of } x_1(X_i) \text{ or } x_2(X_i) \text{ is pruned}) \leq 2p_{z_i} + 2 \frac{m-i+1}{r_i}$$

Analysis on level $z_i + 1$ of the window-tree. For the next step, let $x_1(X_i); \dots; x_t(X_i)$ denote the child-nodes of $x_1(X_i)$ and $x_2(X_i)$ such that $\text{Win}(j(X_i))$ is entirely contained in Win_m . Again, the choice of $x_1; \dots; x_t$ is only a function of X_i . Moreover, since $j \cdot \text{Win}_m > 2^m w_{i,z+1}$ by Eq (4.5), while the window of each node at level $z_i + 1$ is of size $w_{i,z+1}$, we have that $t > 2^m - 2$ always. We now bound the probability that each j is (directly) pruned, for $j \in [t]$. This part of the analysis is quite similar to that of level z_i with only minor changes.

For any choice of $x_1(X_i)$ and $x_2(X_i)$,

(because $\text{Win}_m = \text{Win}(x_1) \cup \text{Win}(x_2)$ and thus j has no choice outside child-nodes of x_1 or x_2)

$$\begin{aligned} \Pr(j(X_i) \text{ is directly pruned} | j_1, j_2) &= \Pr(j(X_i) = j_1, j_2) \\ &= \frac{DP(j_1) \cdot DP(j_2)}{j \cdot DP(j_1) + j \cdot DP(j_2)} \leq \frac{1}{r_i} \end{aligned} \quad (4.8)$$

where we are again going to argue that the probability that $j(X_i)$ is equal to any fixed node is at most $1/r_i$ conditioned on the choice of x_1 and x_2 . For $j(X_i)$ to be equal to a node we need to have that $X_i + o_{i,m} + (j-1)w_{i,z+1} \in \text{Win}(j)$; this is because $j(X_i)$ appears after $(j-1)$ nodes of level $z_i + 1$ that

are fully inside Win_m and each such window has length $w_{i;z_i+1}$ (note that this is a necessary but not a sufficient condition). Thus,

$$\Pr_{\mathbf{X}}(j(X_i) = j_{1;2}) \leq \Pr_{\mathbf{X}_{i;z_i}}(X_i + o_{i;m} + (j-1)w_{i;z_i+1} \in \text{Win}(j_{1;2})) \leq \frac{j \text{Win}(j)}{w_i} = \frac{1}{r_i}$$

where the last inequality is because conditioned on Win_m intersecting with $j_{1;2}$, X_i is chosen uniformly at random from a window of length $w_{i;z_i}$ (equal to length of $\text{Win}(1)$ and $\text{Win}(2)$); the final equality also uses that $j \text{Win}(j) = w_{i;z_i+1} = w_{i;z_i} = r_i$. Hence, we get Eq (4.8).

We can now deduce that

$$\begin{aligned} & \mathbb{E}_{\mathbf{X}} [\# \text{ of } 1(X_i); \dots; t(X_i) \text{ that are directly pruned}] \\ & \text{(by the law of total probability over the choices of } j_{1;2}) \\ &= \mathbb{E}_{j_{1;2}} \mathbb{E}_{\mathbf{X}_i} [\# \text{ of } 1(X_i); \dots; t(X_i) \text{ that are directly pruned } j_{1;2}] \\ (4.9) \quad &= \mathbb{E}_{j_{1;2}} [j \text{DP}(1)j + j \text{DP}(2)j] \frac{1}{r_i}; \quad t \end{aligned}$$

where the last inequality is by Eq (4.8).

Let $\bar{P}(z_i)$ denote the set of not pruned nodes in level z_i and let $\hat{P}(z_i)$ denote the set of nodes in level z_i whose parents are not pruned. Since we are conditioning on E_i , we know that X_i is uniformly random from the interval $[2^{z_i} \text{Win}(j)]$. It follows that $X_{m-1} = X_i + o_{i;m}$ is uniformly random in a range whose size is also $2^{z_i} \text{Win}(j)$. Thus, for any level- z_i node, we have that

$$\Pr[1 = j] = \Pr[X_{m-1} \in \text{Win}(j)] \leq \frac{j \text{Win}(j)}{2^{z_i} \text{Win}(j)} = \frac{1}{j P(z_i)}$$

Summing over the level- (z_i+1) nodes that are directly pruned, we have that

$$\mathbb{E} j \text{DP}(1)j = \sum_i \Pr[1 \text{ is the parent of } j] \leq \sum_D \frac{j P(z_i+1)j}{j \hat{P}(z_i)j} \leq \frac{j P(z_i+1)j}{j P(z_i)j}$$

using the upper bound established above on the probability that 1 is any fixed node. Note that

$$p_{z_i+1} = \frac{j \text{DP}(z_i+1)j}{r_i j P(z_i)j};$$

i.e., the number of directly pruned nodes in level z_i+1 divided by the number of nodes with not pruned parents. Therefore, $\mathbb{E} j \text{DP}(1)j \leq r_i p_{z_i+1}$. By the same reasoning (but applied to 2, which contains the endpoint of X_m), we have that $\mathbb{E} j \text{DP}(2)j \leq r_i p_{z_i+1}$.

Thus, we can use Eq (4.9) to conclude that

$$\mathbb{E}_{\mathbf{X}} [\# \text{ of } 1(X_i); \dots; t(X_i) \text{ that are directly pruned}] \leq 2 p_{z_i+1} t.$$

By Markov's inequality,

$$(4.10) \quad \Pr_{\mathbf{X}}(\text{more than } t=2 \text{ of } 1(X_i); \dots; t(X_i) \text{ are directly pruned}) \leq 4 p_{z_i+1}.$$

Finally, by considering the possibility that at least one of 1 or 2 could be pruned also we have,

$$\begin{aligned} & \Pr_{\mathbf{X}}(\text{more than } t=2 \text{ of } 1; \dots; t \text{ are pruned}) \\ & \leq \Pr(\text{more than } t=2 \text{ of } 1(X_i); \dots; t(X_i) \text{ are directly pruned}) \\ & \quad + \Pr(\text{either of } 1 \text{ or } 2 \text{ are pruned}) \leq \\ (4.11) \quad & 4 p_{z_i+1} + 2 p_{z_i} + \frac{2m}{r_i} \end{aligned}$$

by Eq (4.7) and Eq (4.10).

Concluding the proof. Let us now condition on the event that at least $t=2$ of nodes $1; \dots; t$ are not pruned, namely, the complement of the event in Eq (4.11). Given that Win_m can have non-empty intersection with at most two other level- $(z_i + 1)$ nodes beside $1; \dots; t$, and that non-pruned nodes are all dense, conditioned on the above event, we have,

$$\text{density}_f(\text{Win}_m; i) > \frac{(t=2) \cdot 100=m}{t+2} > \frac{100}{3m} > \frac{2}{m};$$

as $t > 2^{m^m} - 2$. Thus, by Eq (4.11), we have,

$$\Pr_{\text{density}_f}(\text{Win}_m; i) \leq \frac{2}{m} \cdot \frac{6}{2p_{z_i} + 4p_{z_i+1} + \frac{2m}{r_i}} < \frac{4}{r_i} \cdot \frac{p_{z_i} + p_{z_i+1} + \frac{m}{r_i}}{m};$$

concluding the proof. ■

Claims 4.10 to 4.12 now cover all possible cases and allow us to prove Lemma 4.6.

Proof of Lemma 4.6. Fix the tree T_i and consider its pruning process. If $\prod_{j=0}^{Q_k} (1 - p_j) \leq 1/m$, we achieve the first condition of the lemma by Claim 4.10 and are thus done. Now consider the complement case. In this case, we have,

$$\frac{1}{m} < \prod_{j=0}^{Q_k} (1 - p_j) \leq \exp \left(- \sum_{j=0}^{Q_k} p_j \right);$$

which implies that $\sum_{j=0}^{Q_k} p_j \leq \ln m$. Recall that the choice of Z_i in the distribution is uniform over $[k-1]$ regardless of conditioning on $(s_{<i}; x_{<i})$. Since Z_i is chosen uniformly from $[k-1]$, we have,

$$\mathbb{E}_{Z_i} [p_{z_i} + p_{z_i+1}] \leq \frac{1}{k-1} \sum_{j=1}^{Q_k} p_j + \frac{1}{k-1} \sum_{j=2}^{Q_k} p_j \leq \frac{2}{k-1} \sum_{j=0}^{Q_k} p_j \leq \frac{2 \ln m}{(k-1)};$$

By Markov bound, we have,

$$\Pr_{Z_i} [p_{z_i} + p_{z_i+1} > \frac{4}{1-k} \frac{\ln m}{k^{1=3}}] \leq \frac{1}{m};$$

We can now condition on any choice z_i of Z_i such that $p_{z_i} + p_{z_i+1} \leq (4 \ln m)/k^{1=2}$. At this point, either event E_i does not happen, in which case, by Claim 4.11, we again obtain condition (i) of the lemma; or the event E_i happens, which by Claim 4.12 and the choice of r_i in Eq (4.4) implies

$$\Pr_{x_i} [\text{density}_f(\text{Win}_m; i) \leq \frac{2}{m} \cdot \frac{j(s_{<i}; x_{<i})}{4} \leq \frac{4 \ln m}{k^{1=2}} + \frac{m}{k^{1=2}} \cdot \frac{1}{k^m}] \leq \frac{1}{m};$$

as $i \leq m-1$ and thus $m=2^{k^{m-i}} \leq m=2^{k-1}=k^{1=3}$, as $k = m^m$. Taking the union bound over the above two events, we also obtain condition (ii) of the lemma. ■

Finally, we use this lemma to conclude the proof of Lemma 4.2.

Proof of Lemma 4.2. Let $T_1; T_2 \subseteq [m]$ denote, respectively, the iterations in which condition (i) or condition (ii) of Lemma 4.6 happens. Note that T_1 and T_2 are random variables over the randomness of S_i 's and X_i 's. We first claim that with high probability $|T_2| \leq m=2$. This is because for any iteration $i \in T_2$ and any choice of $(s_{<i}; x_{<i})$ of prior iterations, by Lemma 4.6,

$$\Pr_{x_i} [\text{density}_f(\text{Win}_m; i) \leq \frac{2}{m} \cdot \frac{j(s_{<i}; x_{<i})}{m}] \leq \frac{1}{k^{1=3}};$$

A union bound on at most m choices for indices on T_2 then implies that with probability at least $1 - m^{-k^{1=3}}$, we have $\text{density}_f(\text{Win}_m; i) > \frac{2}{m}$ for all $i \in T_2$. But then conditioned on this event, the size of T_2 cannot be $m=2$ or

larger as otherwise Win_m contains $m=2$ disjoint sets each of which contains than a $2=m$ fraction of the window, which is a contradiction. Thus,

$$(\text{as } k = m^m) \quad \Pr(jT_2j > m=2) \leq \frac{m}{k^{1=3}} \frac{1}{k^{1=4}}.$$

We condition on the complement of this event in the following, namely, that $jT_2j < m=2$. Let $i_1; \dots; i_{m=2}$ denote the first $m=2$ indices of T_1 which by the conditioning on the size of T_2 is well defined. We have,

$$\Pr(\text{for all } j \in [m=2]: f(X_{i_j}) = i_j) = \prod_{j \in [m=2]} \Pr(f(X_{i_j}) = i_j \mid f(X_{i_1}) = i_1; \dots; f(X_{i_{j-1}}) = i_{j-1})$$

(since these are type (i) iterations and we can apply condition (i) of Lemma 4.6)

$$\leq \frac{101^{m=2}}{m} :$$

Putting these two together, combined with the value of $k = m^m$, implies that,

$$\Pr_{(X_1; \dots; X_m)}((8i \in [m] : f(X_i) = i) \leq \frac{1}{k^{1=4}} + \frac{101^{m=2}}{m} \leq m^{-m};$$

for some constant $\epsilon > 0$ (taking $\epsilon = 1=100$ certainly suffices). This concludes the proof. \blacksquare

Acknowledgements We are thankful to the anonymous reviewers of SODA 2023 for many helpful comments on the presentation of the paper.

William Kuszmaul was partially sponsored by the United States Air Force Research Laboratory and the United States Air Force Artificial Intelligence Accelerator under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

Additionally, Sepehr Assadi was supported in part by a NSF CAREER grant CCF-2047061, a Google Research gift, and a Fulcrum award from Rutgers Research Council. Martin Farach-Colton was supported by NSF grants CCF-2118620 and CCF-2106999. William Kuszmaul was supported by an NSF GRFP fellowship and a Hertz fellowship.

References

- [BBPV09] Djamel Belazzougui, Paolo Boldi, Rasmus Pagh, and Sebastiano Vigna. Monotone minimal perfect hashing: searching a sorted table with $O(1)$ accesses. In SODA, pages 785{794. SIAM, 2009. 1, 2
- [BBPV11] Djamel Belazzougui, Paolo Boldi, Rasmus Pagh, and Sebastiano Vigna. Theory and practice of monotone minimal perfect hashing. Journal of Experimental Algorithmics (JEA), 16:3{1, 2011. 1, 2
- [BCKM20] Djamel Belazzougui, Fabio Cunial, Juha Karkkainen, and Veli Mäkinen. Linear-time string indexing and analysis in small space. ACM Transactions on Algorithms (TALG), 16(2):1{54, 2020. 2
- [BCO11] Alexandra Boldyreva, Nathan Chenette, and Adam O’Neill. Order-preserving encryption revisited: Improved security analysis and alternative solutions. In Annual Cryptology Conference, pages 578{595. Springer, 2011. 1
- [Bel14] Djamel Belazzougui. Linear time construction of compressed text indices in compact space. In Proceedings of the forty-sixth Annual ACM Symposium on Theory of Computing, pages 148{193, 2014. 2
- [BGMP16] Djamel Belazzougui, Travis Gagie, Veli Mäkinen, and Marco Previtali. Fully dynamic de bruijn graphs. In International symposium on string processing and information retrieval, pages 145{152. Springer, 2016. 2
- [BN14] Djamel Belazzougui and Gonzalo Navarro. Alphabet-independent compressed text indexing. ACM Transactions on Algorithms (TALG), 10(4):1{19, 2014. 2
- [BN15] Djamel Belazzougui and Gonzalo Navarro. Optimal lower and upper bounds for representing sequences. ACM Transactions on Algorithms (TALG), 11(4):1{21, 2015. 2

- [Bol15] Paolo Boldi. Minimal and monotone minimal perfect hash functions. In Giuseppe F. Italiano, Giovanni Pighizzini, and Donald Sannella, editors, *Mathematical Foundations of Computer Science 2015 - 40th International Symposium, MFCS 2015*, Milan, Italy, August 24–28, 2015, Proceedings, Part I, volume 9234 of *Lecture Notes in Computer Science*, pages 3{17. Springer, 2015. 2
- [BV08] P Boldi and S Vigna. Sux4j 1.0. 2008. 1
- [CFP⁺15] Raphaël Cliord, Allyx Fontaine, Ely Porat, Benjamin Sach, and Tatiana Starikovskaya. Dictionary matching in a stream. In *Algorithms-ESA 2015*, pages 361{372. Springer, 2015. 2
- [D⁺18] Martin Dietzfelbinger et al. 4.3 space complexity of monotone minimal perfect hashing. *Dagstuhl Reports*, Vol. 7, Issue 5 ISSN 2192-5283, page 19, 2018. 2
- [DLR95] Dwight Duus, Hannon Lefmann, and Vojtech Rödl. Shift graphs and lower bounds on ramsey numbers $rk(l; r)$. *Discrete Mathematics*, 137(1-3):177{187, 1995. 2
- [EH66] Paul Erdős and András Hajnal. On chromatic number of graphs and set-systems. *Acta Math. Acad. Sci. Hungar*, 17(61-99):1, 1966. 2, 3
- [FHRT92] Zoltan Füredi, Peter Hajnal, Vojtech Rödl, and William T Trotter. Interval orders and shift graphs. 1992. 2
- [GNP20] Travis Gagie, Gonzalo Navarro, and Nicola Prezza. Fully functional sux trees and optimal text searching in bwt-runs bounded space. *Journal of the ACM (JACM)*, 67(1):1{54, 2020. 2
- [GOR10] Roberto Grossi, Alessio Orlandi, and Rajeev Raman. Optimal trade-offs for succinct string indexes. In *International Colloquium on Automata, Languages, and Programming*, pages 678{689. Springer, 2010. 2
- [HT01] Torben Hagerup and Torsten Tholey. Efficient minimal perfect hashing in nearly minimal space. In Afonso Ferreira and Horst Reichel, editors, *STACS 2001, 18th Annual Symposium on Theoretical Aspects of Computer Science*, Dresden, Germany, February 15-17, 2001, Proceedings, volume 2010 of *Lecture Notes in Computer Science*, pages 317{326. Springer, 2001. 4
- [LFAK11] Hyeontaek Lim, Bin Fan, David G Andersen, and Michael Kaminsky. Silt: A memory-efficient, high-performance key-value store. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*, pages 1{13, 2011. 1
- [Meh82] Kurt Mehlhorn. On the program size of perfect and universal hash functions. In *23rd Annual Symposium on Foundations of Computer Science (sfcs 1982)*, pages 170{175. IEEE, 1982. 2
- [Nav14] Gonzalo Navarro. Spaces, trees, and colors: The algorithmic landscape of document retrieval on sequences. *ACM Computing Surveys (CSUR)*, 46(4):1{47, 2014. 1
- [New91] Ilan Newman. Private vs. common random bits in communication complexity. *Inf. Process. Lett.*, 39(2):67{71, 1991. 3
- [ST11] Gabor Simonyi and Gábor Tardos. On directed local chromatic number, shift graphs, and borsuk-like graphs. *Journal of Graph Theory*, 66(1):65{82, 2011. 2, 3
- [SU11] Edward R Scheinerman and Daniel H Ullman. *Fractional graph theory: a rational approach to the theory of graphs*. Courier Corporation, 2011. 4, 19

Appendix

A Proofs of Standard Results in Fractional Coloring

We prove [Propositions 2.2](#) and [2.3](#) here for completeness. These proofs are standard; see, e.g. [\[SU11\]](#). We start by presenting the dual view of fractional colorings that is the key to these proofs.

The dual view of fractional colorings. Given that $\chi_f(G)$ is defined as a solution to an LP, we can use duality to also express $\chi_f(G)$ via the following LP:

$$(A.1) \quad \chi_f(G) := \max_{y \in [0,1]^{V(G)}} \sum_{v \in G} y_v \quad \text{subject to} \quad \sum_{v \in I} y_v \leq 1 \quad \forall I \in \mathcal{I}(G);$$

This LP is a fractional relaxation of the clique number of G , namely, the size of the largest clique in G (since, in any integral solution to this LP, the y -values that are 1 must be on the vertices of a clique). Interestingly, although the chromatic number and clique size are not duals, their relaxations are.

Proposition (Restatement of [Proposition 2.2](#)). Let $G_1 = (V_1; E_1)$ and $G_2 = (V_2; E_2)$ be arbitrary graphs. Define $G_1 \sqcup G_2$ as a graph on vertices $V_1 \cup V_2$ and define an edge between vertices $(v_1; v_2)$ and $(w_1; w_2)$ whenever $(v_1; w_1)$ is an edge in G_1 or $(v_2; w_2)$ is an edge in G_2 . Then, $\chi_f(G_1 \sqcup G_2) = \chi_f(G_1) + \chi_f(G_2)$.

Proof of [Proposition 2.2](#). We first prove that

$$(A.2) \quad \chi_f(G_1 \sqcup G_2) \geq \chi_f(G_1) + \chi_f(G_2);$$

Let $y^1 \in [0,1]^{V_1}$ and $y^2 \in [0,1]^{V_2}$ be optimal solutions to the dual LP given by [Eq \(A.1\)](#) for G_1 and G_2 , respectively. Consider the assignment $y \in [0,1]^{V_1 \cup V_2}$ where $y_{u_1; u_2} = y_{u_1}^1 + y_{u_2}^2$. We clearly have that

$$\sum_{(u_1; u_2) \in V_1 \cup V_2} y_{u_1; u_2} = \sum_{u_1 \in V_1} y_{u_1}^1 + \sum_{u_2 \in V_2} y_{u_2}^2 = \chi_f(G_1) + \chi_f(G_2);$$

We now argue that y is also a valid solution to the dual LP given by [Eq \(A.1\)](#) for $G_1 \sqcup G_2$. Fix any independent set $I \in \mathcal{I}(G_1 \sqcup G_2)$. By the definition of the product, we know that I can be written as a product set, namely, $I = I_1 \sqcup I_2$ for $I_1 \in \mathcal{I}(G_1)$ and $I_2 \in \mathcal{I}(G_2)$. Thus,

$$\sum_{(u_1; u_2) \in I} y_{u_1; u_2} = \sum_{u_1 \in I_1} y_{u_1}^1 + \sum_{u_2 \in I_2} y_{u_2}^2 \leq 1 + 1 = 2;$$

where the inequality is by the constraint of dual LP for y^1 and y^2 each. Thus, y is a solution to the dual LP for $G_1 \sqcup G_2$, proving [Eq \(A.2\)](#).

We now prove that

$$(A.3) \quad \chi_f(G_1 \sqcup G_2) \leq \chi_f(G_1) + \chi_f(G_2);$$

using the primal LP instead. Let $x^1 \in [0,1]^{V_1}$ and $x^2 \in [0,1]^{V_2}$ be optimal solutions to primal LP from [Eq \(2.1\)](#) for G_1 and G_2 , respectively. Consider the assignment $x \in [0,1]^{V_1 \cup V_2}$ where $x_i = x_i^1 + x_i^2$, using the fact from the previous part that $I = I_1 \sqcup I_2$ for $I_1 \in \mathcal{I}(G_1)$ and $I_2 \in \mathcal{I}(G_2)$.

We again clearly have that

$$\sum_{(u_1; u_2) \in I(G_1 \sqcup G_2)} x_i = \sum_{i \in I_1(G_1)} x_i^1 + \sum_{i \in I_2(G_2)} x_i^2 \leq 1 + 1 = 2;$$

so it remains to prove that x is a valid solution to the primal LP from Eq (2.1) for $G_1 \cup G_2$. Fix any vertex $(u_1; u_2) \in V_1 \times V_2$ and consider all independent sets $I_1 \subseteq V_1$ that contain u_1 and $I_2 \subseteq V_2$ that contain u_2 . Then, $I_1 \cup I_2$ is also an independent set in $G_1 \cup G_2$ that contains $(u_1; u_2)$. Thus,

$$\sum_{(u_1; u_2) \in I_1 \cup I_2} x_{(u_1; u_2)} = \sum_{u_1 \in I_1} x_{u_1} + \sum_{u_2 \in I_2} x_{u_2} \geq \sum_{u_1 \in I_1} x_{u_1}^{(1)} + \sum_{u_2 \in I_2} x_{u_2}^{(2)} \geq 1 = 1;$$

where the inequality is by the constraint of primal LP from Eq (2.1) for $x^{(1)}$ and $x^{(2)}$ each. Thus, x is a solution to the primal LP from Eq (2.1) for $G_1 \cup G_2$, proving Eq (A.3). ■

Proposition (Restatement of Proposition 2.3). For any graph $G = (V; E)$,

$$f(G) = \max_{\text{distribution on } V} \min_{I \subseteq V} \Pr(v \in I) = 1;$$

Proof of Proposition 2.3. Let μ be any distribution on $V(G)$ and define $b := \max_{I \subseteq V(G)} \Pr(v \in I)^{-1}$. Create $y \in \mathbb{R}^{V(m)}$ such that $y_v = b \cdot \mu(v)$ for every vertex $v \in V(m)$ where $\mu(v)$ is the probability of vertex v under the distribution μ . We claim that y is a feasible dual solution in Eq (A.1).

For every independent set $I \subseteq V(G)$,

$$\sum_{v \in I} y_v = b \sum_{v \in I} \mu(v) = b \Pr(v \in I) \leq 1;$$

by the definition of b . Thus y is a feasible dual solution. Moreover,

$$\sum_{v \in V(G)} y_v = b \sum_{v \in V(G)} \mu(v) = b;$$

As the dual LP in Eq (A.1) is a maximization LP, we have that $f(G) \geq b = \max_{I \subseteq V(G)} \Pr(v \in I)^{-1}$, for any distribution μ on the vertices.

Conversely, let y be any optimal solution to the dual LP and let $c := \sum_{v \in V} y_v$. Define a distribution μ on the vertices V by setting $\mu(v) = y_v/c$. For any independent set $I \subseteq V(G)$, we have,

$$\Pr(v \in I) = \sum_{v \in I} \mu(v) = \sum_{v \in I} y_v/c \leq 1/c;$$

where the inequality is because y is a feasible dual solution. Thus, there exists a distribution μ such that $f(G) = c \leq \max_{I \subseteq V(G)} \Pr(v \in I)^{-1}$.

Combining these two parts concludes the proof. ■

B Covering The Full Range of the Universe Size

We now generalize the proof of Theorem 2 to the full parameter range specified in the theorem. Consider u and n satisfying

$$n2^{2^{\frac{p}{\log \log n}}} \leq u \leq 2^{n^2 + n}.$$

Notice that, on the lower-bound side, we are actually covering a slightly larger range (and therefore proving a slightly stronger result) than required to establish Theorem 2.

Set

$$m = (\log \log u)^{1/6} \quad \text{and} \quad k = n/m = n/(\log \log u)^{1/6}.$$

Note that the setting of k implicitly requires that $(\log \log u)^{1/6} \leq n$, which follows from the fact that $(\log \log u)^{1/6} \leq (n^2 + n)^{1/6} \leq n$.

The k -fold conict graph $G^k(m)$ has $\log_f(G^k(m)) = (n \log m) = (n \log \log \log u)$ as already argued in [Section 3.2](#). To complete the proof, we must establish that the graph $G^k(m)$ has vertices that are subsets of a universe whose size u_0 satisfies $u_0 \leq u$. Solving for u^0 , we have that

$$u^0 = kM = \frac{n}{(\log \log u)^{1/6}} 2^{m^{m^2+m}} n 2^{m^3=2} n 2^{p_{\log \log u=2}};$$

On the other hand, $u \leq n 2^{p_{\log \log u}}$. It follows that

$$\frac{u}{n} \leq \frac{2^{p_{\log \log u}}}{2^{p_{\log \log u=2}}} = 1;$$

which completes the proof of [Theorem 2](#) for any choice of u between $n 2^{p_{\log \log u}}$ and 2^{n^2+n} .

Finally, we remark that the term 2^{n^2+n} in the upper bound is not tight and can be replaced by any other $2^{poly(n)}$ term; this is simply because for any $u = 2^{poly(n)}$, $\log \log \log u = (\log n)$ and thus for any larger universe size u also, we can simply focus on the smallest 2^{n^2+n} numbers in the universe and still obtain the same asymptotic lower bound. The lower bound term is also not tight and can be replaced with $n 2^{(\log \log n)^c}$ for any constant $c \geq 2$ (0; 1=2) by the same argument.