

Tight Bounds for Hashing Block Sources*

Kai-Min Chung[†]
Harvard University

Salil Vadhan[‡]
Harvard University

November 14, 2009

Abstract

It is known that if a 2-universal hash function H is applied to elements of a *block source* (X_1, \dots, X_T) , where each item X_i has enough min-entropy conditioned on the previous items, then the output distribution $(H, H(X_1), \dots, H(X_T))$ will be “close” to the uniform distribution. We provide improved bounds on how much min-entropy per item is required for this to hold, both when we ask that the output be close to uniform in statistical distance and when we only ask that it be statistically close to a distribution with small collision probability. In both cases, we reduce the dependence of the min-entropy on the number T of items from $2 \log T$ in previous work to $\log T$, which we show to be optimal. This leads to corresponding improvements to the recent results of Mitzenmacher and Vadhan (SODA ‘08) on the analysis of hashing-based algorithms and data structures when the data items come from a block source.

1 Introduction

A *block source* is a sequence of items $\mathbf{X} = (X_1, \dots, X_T)$ in which each item has at least some k bits of “entropy” conditioned on the previous ones [CG88]. Previous works [CG88, Zuc96, MV08] have analyzed what happens when one applies a 2-universal hash function to each item in such a sequence, establishing results of the following form:

Block-Source Hashing Theorems (informal): *If (X_1, \dots, X_T) is a block source with k bits of “entropy” per item and H is a random hash function from a 2-universal family mapping to $m \ll k$ bits, then $(H(X_1), \dots, H(X_T))$ is “close” to the uniform distribution.*

*An extended abstract of this paper will appear in *RANDOM ‘08* [CV08].

[†]Work done when visiting U.C. Berkeley, supported by US-Israel BSF grant 2002246 and NSF grant CNS-0430336.

[‡]Work done when visiting U.C. Berkeley, supported by the Miller Institute for Basic Research in Science, a Guggenheim Fellowship, US-Israel BSF grant 2006060, and ONR grant N00014-04-1-0478.

In this paper, we prove new results of this form, achieving improved (in some cases, optimal) bounds on how much entropy k per item is needed to ensure that the output is close to uniform, as a function of the other parameters (the output length m of the hash functions, the number T of items, and the “distance” from the uniform distribution). But first we discuss the two applications that have motivated the study of Block-Source Hashing Theorems.

1.1 Applications of Block-Source Hashing

Randomness Extractors. A *randomness extractor* is an algorithm that extracts almost-uniform bits from a source of biased and correlated bits, using a short *seed* of truly random bits as a catalyst [NZ96]. Extractors have many applications in theoretical computer science and have played a central role in the theory of pseudorandomness. (See the surveys [NT99, Sha04, Vad07].) Block-source Hashing Theorems immediately yield methods for extracting randomness from block sources, where the seed is used to specify a universal hash function. The gain over hashing the entire T -tuple at once is that the blocks may be much shorter than the entire sequence, and thus a much shorter seed is required to specify the universal hash function. Moreover, many subsequent constructions of extractors for general sources (without the block structure) work by first converting the source into a block source and performing block-source hashing.

Analysis of Hashing-Based Algorithms. The idea of hashing has been widely applied in designing algorithms and data structures, including hash tables [Knu98], Bloom filters [BM03], summary algorithms for data streams [Mut03], etc. Given a stream of data items (x_1, \dots, x_T) , we first hash the items into $(H(x_1), \dots, H(x_T))$, and carry out a computation using the hashed values. In the literature, the analysis of a hashing algorithm is typically a worst-case analysis on the input data items, and the best results are often obtained by unrealistically modelling the hash function as a truly random function mapping the items to uniform and independent m -bit strings. On the other hand, for realistic, efficiently computable hash functions (eg., 2-universal or $O(1)$ -wise independent hash functions), the provable performance is sometimes significantly worse. However, such gaps seem to not show up in practice, and even standard 2-universal hash functions empirically seem to match the performance of truly random hash functions. To explain this phenomenon, Mitzenmacher and Vadhan [MV08] have suggested that the discrepancy is due to worst-case analysis, and propose to instead model the input items as coming from a block source. Then Block-Source Hashing Theorems imply that the performance of universal hash functions is close to that of truly random hash functions, provided that each item has enough bits of entropy.

1.2 How Much Entropy is Required?

A natural question about Block-Source Hashing Theorems is: how large does the “entropy” k per item need to be to ensure a certain amount of “closeness” to uniform (where both the entropy and closeness can be measured in various ways). This also has practical significance for the latter motivation regarding hashing-based algorithms, as it corresponds to the amount of entropy we need to assume in data items. In [MV08], they provide bounds on the entropy required for two measures of closeness, and use these as basic tools to bound the required entropy in various applications. The requirement is usually some small constant multiple of $\log T$, where T is the number of items in the source, which can be on the borderline between a reasonable and unreasonable assumption about real-life data. Therefore, it is interesting to pin down the optimal answers to these questions. In what follows, we first summarize the previous results, and then discuss our improved analysis and corresponding lower bounds.

A standard way to measure the distance of the output from the uniform distribution is by *statistical distance*.¹ In the randomness extractor literature, classic results [CG88, ILL89, Zuc96] show that using 2-universal hash functions, $k = m + 2 \log(T/\varepsilon) + O(1)$ bits of min-entropy (or even Renyi entropy)² per item is sufficient for the output distribution to be ε -close to uniform in statistical distance. Sometimes a less stringent closeness requirement is sufficient, where we only require that the output distribution is ε -close to a distribution having “small” *collision probability*³. A result of [MV08] shows that $k = m + 2 \log T + \log(1/\varepsilon) + O(1)$ suffices to achieve this requirement. Using 4-wise independent hash functions, [MV08] further reduce the required entropy to $k = \max\{m + \log T, 1/2(m + 3 \log T + \log(1/\varepsilon))\} + O(1)$.

Our Results. We reduce the entropy required in the previous results, as summarized in Table 1. Roughly speaking, we save an additive $\log T$ bits of min-entropy (or Renyi entropy) for all cases. We show that using universal hash functions, $k = m + \log T + 2 \log 1/\varepsilon + O(1)$ bits per item is sufficient for the output to be ε -close to uniform, and $k = m + \log(T/\varepsilon) + O(1)$ is enough for the output to be ε -close to having small collision probability. Using 4-wise independent hash functions, the entropy k further reduces to $\max\{m + \log T, 1/2(m + 2 \log T + \log 1/\varepsilon)\} + O(1)$. The results hold even if we consider the joint distribution $(H, H(X_1), \dots, H(X_T))$ (corresponding to “strong extractors” in the literature on randomness extractors). Substituting our improved bounds in the analysis of hashing-based algorithms from [MV08], we obtain similar reductions in the min-entropy required for every application with 2-universal hashing. With 4-wise independent hashing, we obtain a slight improvement for Linear Probing, and for the

¹The *statistical distance* of two random variables X and Y is $\Delta(X, Y) = \max_T |\Pr[X \in T] - \Pr[Y \in T]|$, where T ranges over all possible events.

²The *min-entropy* of a random variable X is $H_\infty(X) = \min_x \log(1/\Pr[X = x])$. All of the results mentioned actually hold for the less stringent measure of *Renyi entropy* $H_2(X) = \log(1/\mathbb{E}_{x \leftarrow X}[\Pr[X = x]])$.

³The *collision probability* of a random variable X is $\sum_x \Pr[X = x]^2$. By “small collision probability,” we mean that the collision probability is within a constant factor of the collision probability of uniform distribution.

Setting	Previous Results	Our Results
2-universal hashing ε -close to uniform	$m + 2 \log T + 2 \log(1/\varepsilon)$ [CG88, ILL89, Zuc96]	$m + \log T + 2 \log(1/\varepsilon)$
2-universal hashing ε -close to small cp.	$m + 2 \log T + \log(1/\varepsilon)$ [MV08]	$m + \log T + \log(1/\varepsilon)$
4-wise indep. hashing ε -close to small cp.	$\max\{m + \log T,$ $1/2(m + 3 \log T + \log 1/\varepsilon)\}$ [MV08]	$\max\{m + \log T,$ $1/2(m + 2 \log T + \log(1/\varepsilon))\}$

Table 1: Our Results: Each entry denotes the min-entropy (actually, Renyi entropy) required per item when hashing a block source of T items to m -bit strings to ensure that the output has statistical distance at most ε from uniform (or from having collision probability within a constant factor of uniform). Additive constants are omitted for readability.

other applications, we show that the previous bounds can already be achieved with 2-universal hashing. The results are summarized in Table 2.

Although the $\log T$ improvement seems small, we remark that it could be significant for practical settings of parameter. For example, suppose we want to hash 64 thousand internet traffic flows, so $\log T \approx 16$. Each flow is specified by the 32-bit IP addresses and 16-bit port numbers for the source and destination plus the 8-bit transport protocol, for a total of 104 bits. There is a noticeable difference between assuming that each flow contains $3 \log T \approx 48$ vs. $4 \log T \approx 64$ bits of entropy as they are only 104 bits long, and are very structured.

We also prove corresponding lower bounds showing that our upper bounds are almost tight. Specifically, we show that when the data items have not enough entropy, then the joint distribution $(H, H(X_1), \dots, H(X_T))$ can be “far” from uniform. More precisely, we show that if $k = m + \log T + 2 \log 1/\varepsilon - O(1)$, then there exists a block source (X_1, \dots, X_T) with k bits of min-entropy per item such that the distribution $(H, H(X_1), \dots, H(X_T))$ is ε -far from uniform in statistical distance (for H coming from any hash family). This matches our upper bound up to an additive constant. Similarly, we show that if $k = m + \log T - O(1)$, then there exists a block source (X_1, \dots, X_T) with k bits of min-entropy per item such that the distribution $(H, H(X_1), \dots, H(X_T))$ is 0.99-far from having small collision probability (for H coming from any hash family). This matches our upper bound up to an additive constant in case the statistical distance parameter ε is constant; we also exhibit a specific 2-universal family for which the $\log(1/\varepsilon)$ in our upper bound is nearly tight — it cannot be reduced below $\log(1/\varepsilon) - \log \log(1/\varepsilon)$. Finally, we also extend all of our lower bounds to the case that we only consider distribution of hashed values $(H(X_1), \dots, H(X_T))$, rather than their joint distribution with Y . For this case, the lower bounds are necessarily reduced by a term that depends on the size of the hash family. (For standard constructions of universal hash functions, this amounts to $\log n$ bits of entropy, where n is the bit-length of an individual item.)

Type of Hash Family	Previous Results [MV08]	Our Results
Linear Probing		
2-universal hashing	$4 \log T$	$3 \log T$
4-wise independence	$2.5 \log T$	$2 \log T$
Balanced Allocations with d Choices		
2-universal hashing	$(d + 2) \log T$	$(d + 1) \log T$
4-wise independence	$(d + 1) \log T$	—
Bloom Filters		
2-universal hashing	$4 \log T$	$3 \log T$
4-wise independence	$3 \log T$	—

Table 2: Applications: Each entry denotes the min-entropy (actually, Renyi entropy) required per item to ensure that the performance of the given application is “close” to the performance when using truly random hash functions. In all cases, the bounds omit additive terms that depend on how close a performance is desired, and we restrict to the (standard) case that the size of the hash table is linear in the number of items being hashed. That is, $m = \log T + O(1)$.

Techniques. At a high level, all of the previous analyses for hashing block sources were loose due to summing error probabilities over the T blocks. Our improvements come from avoiding this linear blow-up by choosing more refined measures of error. For example, when we want the output to have small statistical distance from uniform, the classic Leftover Hash Lemma [ILL89] says that min-entropy $k = m + 2 \log(1/\varepsilon_0)$ suffices for a single hashed block to be ε_0 -close to uniform, and then a “hybrid argument” implies that the joint distribution of T hashed blocks is $T\varepsilon_0$ -close to uniform [Zuc96]. Setting $\varepsilon_0 = \varepsilon/T$, this leads to a min-entropy requirement of $k = m + 2 \log(1/\varepsilon) + 2 \log T$ per block. We obtain a better bound, reducing $2 \log T$ to $\log T$, by using *Hellinger distance* to analyze the error accumulation over blocks, and only passing to statistical distance at the end.

For the case where we only want the output to be close to having small collision probability, the previous analysis of [MV08] worked by first showing that the expected collision probability of each hashed block $h(X_i)$ is “small” even conditioned on previous blocks, then using Markov’s Inequality to deduce that each hashed block has small collision probability except with some probability ε_0 , and finally doing a union bound to deduce that all hashed blocks have small collision probability except with probability $T\varepsilon_0$. We avoid the union bound by working with more refined notions of “conditional collision probability,” which enable us to apply Markov’s Inequality on the entire sequence rather than on each block individually.

The starting point for our negative results is the tight lower bound for randomness extractors due to Radhakrishnan and Ta-Shma [RT00]. Their methods show that if the min-entropy parameter k is not large enough, then for any hash family, there exists a (single-block) source

X such that $h(X)$ is “far” from uniform (in statistical distance) for “many” hash functions h . We then take our block source (X_1, \dots, X_T) to consist of T iid copies of X , and argue that the statistical distance from uniform grows sufficiently fast with the number T of copies taken. For example, we show that if two distributions have statistical distance ε , then their T -fold products have statistical distance $\Omega(\min\{1, \sqrt{T} \cdot \varepsilon\})$, strengthening a previous bound of Reyzin [Rey04], who proved a bound of $\Omega(\min\{\varepsilon^{1/3}, \sqrt{T} \cdot \varepsilon\})$.

2 Preliminaries

Notations. All logs are based 2. We use the convention that $N = 2^n$, $K = 2^k$, and $M = 2^m$. We think of a data item X as a random variable over $[N] = \{1, \dots, N\}$, which can be viewed as the set of n -bit strings. A hash function $h : [N] \rightarrow [M]$ hashes an item to a m -bit string. A *hash function family* \mathcal{H} is a multiset of hash functions, and H will usually denote a uniformly random hash function drawn from \mathcal{H} . $U_{[M]}$ denotes the uniform distribution over $[M]$. Let $\mathbf{X} = (X_1, \dots, X_T)$ be a sequence of data items. We use $X_{<i}$ to denote the first $i - 1$ items (X_1, \dots, X_{i-1}) . We refer to X_i as an item or a block interchangeably. Our goal is to study the distribution of hashed sequence $(H, \mathbf{Y}) = (H, Y_1, \dots, Y_T) \stackrel{\text{def}}{=} (H, H(X_1), \dots, H(X_T))$.

Hash Families. The *truly random hash family* \mathcal{H} is the set of all functions from $[N]$ to $[M]$. A hash family \mathcal{H} is *s-universal* if for every sequence of distinct elements $x_1, \dots, x_s \in [N]$, $\Pr_H[H(x_1) = \dots = H(x_s)] \leq 1/M^s$. \mathcal{H} is *s-wise independent* if for every sequence of distinct elements $x_1, \dots, x_s \in [N]$, $H(x_1), \dots, H(x_s)$ are independent and uniform random variables over $[M]$.

Block Sources and Collision Probability. For a random variable X , the collision probability of X is $\text{cp}(X) = \Pr[X = X'] = \sum_x \Pr[X = x]^2$, where X' is an independent copy of X . The *Renyi entropy* $H_2(X) = \log(1/\text{cp}(X))$ can be viewed as a measure of the amount of randomness in X (In the randomness extractor literature, the entropy is measured by *min-entropy* $H_\infty(X) = \min_{x \in \text{supp}(X)} \log(1/\Pr[X = x])$, but using the less stringent measure Renyi entropy makes our results stronger since $H_2(X) \geq H_\infty(X)$.) For an event E , $(X|_E)$ is the random variable defined by conditioning X on E .

Definition 2.1 (Block Sources) *A sequence of random variables (X_1, \dots, X_T) over $[N]^T$ is a block K -source if for every $i \in [T]$, and every $x_{<i}$ in the support of $X_{<i}$, we have $\text{cp}(X_i|X_{<i} = x_{<i}) \leq 1/K$. That is, each item X_i has at least $k = \log K$ bits of Renyi entropy even after conditioning on the previous items.*

Let $\mathbf{X} = (X_1, \dots, X_T)$ be a sequence of random variables over $[M]^T$. We are interested in bounding the overall collision probability $\text{cp}(\mathbf{X})$ by the collision probability of each blocks.

Suppose all X_i 's are independent, then $\text{cp}(\mathbf{X}) = \prod_{i=1}^T \text{cp}(X_i)$. The following lemma generalizes Lemma 4.2 in [MV08], which says that if for every $\mathbf{x} \in \mathbf{X}$, the average collision probability of every block X_i conditioning on $X_{<i} = x_{<i}$ is small, then the overall collision probability $\text{cp}(\mathbf{X})$ is also small. In particular, if \mathbf{X} is a block K -source, then $\text{cp}(\mathbf{X}) \leq 1/K^T$.

Lemma 2.2 *Let $\mathbf{X} = (X_1, \dots, X_T)$ be a sequence of random variables such that for every $\mathbf{x} \in \text{supp}(\mathbf{X})$,*

$$\frac{1}{T} \sum_{i=1}^T \text{cp}(X_i | X_{<i} = x_{<i}) \leq \alpha.$$

Then the overall collision probability satisfies $\text{cp}(\mathbf{X}) \leq \alpha^T$.

Proof. By Arithmetic Mean-Geometric Mean Inequality, the inequality in the premise implies

$$\prod_{i=1}^T \text{cp}(X_i | X_{<i} = x_{<i}) \leq \alpha^T.$$

Therefore, it suffices to prove

$$\text{cp}(\mathbf{X}) \leq \max_{\mathbf{x} \in \text{supp}(\mathbf{X})} \prod_{i=1}^T \text{cp}(X_i | X_{<i} = x_{<i}).$$

We prove it by induction on T . The base case $T = 1$ is trivial. Suppose the lemma is true for $T - 1$. We have

$$\begin{aligned} \text{cp}(\mathbf{X}) &= \sum_{x_1} \Pr[X_1 = x_1]^2 \cdot \text{cp}(X_2, \dots, X_T | X_1 = x_1) \\ &\leq \left(\sum_{x_1} \Pr[X_1 = x_1]^2 \right) \cdot \max_{x_1} \text{cp}(X_2, \dots, X_T | X_1 = x_1) \\ &\leq \text{cp}(X_1) \cdot \max_{x_1} \left(\max_{x_2, \dots, x_T} \prod_{i=2}^T \text{cp}(X_i | X_{<i} = x_{<i}) \right) \\ &= \max_{\mathbf{x}} \prod_{i=1}^T \text{cp}(X_i | X_{<i} = x_{<i}), \end{aligned}$$

as desired. ■

Statistical Distance. The statistical distance is a standard way to measure the distance of two distributions. Let X and Y be two random variables. The *statistical distance* of X and Y is $\Delta(X, Y) = \max_T |\Pr[X \in T] - \Pr[Y \in T]| = (1/2) \cdot \sum_x |\Pr[X = x] - \Pr[Y = x]|$, where T ranges over all possible events. When $\Delta(X, Y) \leq \varepsilon$, we say that X is ε -close to Y . Similarly, if $\Delta(X, Y) \geq \varepsilon$, then X is ε -far from Y . The following standard lemma says that if X has small collision probability, then X is close to uniform in statistical distance.

Lemma 2.3 *Let X be a random variable over $[M]$ such that $\text{cp}(X) \leq (1 + \varepsilon)/M$. Then $\Delta(X, U_{[M]}) \leq \sqrt{\varepsilon}$.*

Conditional Collision Probability. Let (X, Y) be jointly distributed random variables. We can define the conditional Renyi entropy of X conditioning on Y as follows.

Definition 2.4 *The conditional collision probability of X conditioning on Y is $\text{cp}(X|Y) = \mathbb{E}_{y \leftarrow Y}[\text{cp}(X|Y=y)]$. The conditional Renyi entropy is $H_2(X|Y) = \log 1/\text{cp}(X|Y)$.*

The following lemma says that as in the case of Shannon entropy, conditioning can only decrease the entropy.

Lemma 2.5 *Let (X, Y, Z) be jointly distributed random variables. We have $\text{cp}(X) \leq \text{cp}(X|Y) \leq \text{cp}(X|Y, Z)$.*

Proof. For the first inequality, we have

$$\begin{aligned}
\text{cp}(X) &= \sum_x \Pr[X = x]^2 \\
&= \sum_{y, y'} \Pr[Y = y] \cdot \Pr[Y = y'] \cdot \left(\sum_x \Pr[X = x|Y = y] \cdot \Pr[X = x|Y = y'] \right) \\
&\leq \sum_{y, y'} \Pr[Y = y] \cdot \Pr[Y = y'] \cdot \\
&\quad \left(\sum_x \Pr[X = x|Y = y]^2 \right)^{1/2} \cdot \left(\sum_x \Pr[X = x|Y = y']^2 \right)^{1/2} \\
&= \mathbb{E}_{y \leftarrow Y} \left[\text{cp}(X|Y = y)^{1/2} \right]^2 \\
&\leq \text{cp}(X|Y)
\end{aligned}$$

For the second inequality, observe that for every y in the support of Y , we have $\text{cp}(X|_{Y=y}) \leq \text{cp}((X|_{Y=y})|(Z|_{Y=y}))$ from the first inequality. It follows that

$$\begin{aligned}
\text{cp}(X|Y) &= \mathbb{E}_{y \leftarrow Y} [\text{cp}(X|_{Y=y})] \\
&\leq \mathbb{E}_{y \leftarrow Y} [\text{cp}((X|_{Y=y})|(Z|_{Y=y}))] \\
&= \mathbb{E}_{y \leftarrow Y} \left[\mathbb{E}_{z \leftarrow (Z|_{Y=y})} [\text{cp}(X|_{Y=y, Z=z})] \right] \\
&= \text{cp}(X|Y, Z)
\end{aligned}$$

■

3 Positive Results: How Much Entropy is Sufficient?

In this section, we present our positive results, showing that the distribution of hashed sequence $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$ is close to uniform when H is a random hash function from a 2-universal hash family, and $\mathbf{X} = (X_1, \dots, X_T)$ has sufficient entropy per block. The new contribution is that we will not need $K = 2^k$ to be as large as in previous works, and so save the required randomness in the block source $\mathbf{X} = (X_1, \dots, X_T)$.

3.1 Small Collision Probability Using 2-universal Hash Functions

Let $H : [N] \rightarrow [M]$ be a random hash function from a 2-universal family \mathcal{H} . We first study the conditions under which $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$ is ε -close to having collision probability $O(1/(|\mathcal{H}| \cdot M^T))$. This requirement is less stringent than (H, \mathbf{Y}) being ε -close to uniform in statistical distance, and so requires less bits of entropy. Mitzenmacher and Vadhan [MV08] show that this guarantee suffices for some hashing applications. They show that $K \geq MT^2/\varepsilon$ is enough to satisfy the requirement. We save a factor of T , and show that in fact, $K \geq MT/\varepsilon$, is sufficient. (Taking logs yields the first entry in Table 1, i.e. it suffices to have Renyi entropy $k = m + \log T + \log(1/\varepsilon)$ per block.) Formally, we prove the following theorem.

Theorem 3.1 *Let $H : [N] \rightarrow [M]$ be a random hash function from a 2-universal family \mathcal{H} . Let $\mathbf{X} = (X_1, \dots, X_T)$ be a block K -source over $[N]^T$. For every $\varepsilon > 0$, the hashed sequence $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$ is ε -close to a distribution $(H, \mathbf{Z}) = (H, Z_1, \dots, Z_T)$ such that*

$$\text{cp}(H, \mathbf{Z}) \leq \frac{1}{|\mathcal{H}| \cdot M^T} \left(1 + \frac{M}{K\varepsilon}\right)^T.$$

In particular, if $K \geq MT/\varepsilon$, then (H, \mathbf{Z}) has collision probability at most $(1 + 2MT/K\varepsilon)/(|\mathcal{H}| \cdot M^T)$.

To analyze the distribution of the hashed sequence (H, \mathbf{Y}) , the starting point is the following version of the Leftover Hash Lemma [BBR85, ILL89], which says that when we hash a random variable X with enough entropy using a 2-universal hash function H , the conditional collision probability of $H(X)$ conditioning on H is small.

Lemma 3.2 (The Leftover Hash Lemma) *Let $H : [N] \rightarrow [M]$ be a random hash function from a 2-universal family \mathcal{H} . Let X be a random variable over $[N]$ with $\text{cp}(X) \leq 1/K$. We have $\text{cp}(H(X)|H) \leq 1/M + 1/K$.*

We now sketch how the hashed block source $\mathbf{Y} = (Y_1, \dots, Y_T) = (H(X_1), \dots, H(X_T))$ is analyzed in [MV08], and how we improve the analysis. The following natural approach is taken in [MV08]. Since the data \mathbf{X} is a block K -source, the Leftover Hash Lemma tells us that for every block $i \in [T]$, if we condition on the previous blocks $X_{<i} = x_{<i}$, then the hashed value $(Y_i|_{X_{<i}=x_{<i}})$ has small conditional collision probability, i.e. $\text{cp}((Y_i|_{X_{<i}=x_{<i}})|H) \leq 1/M + 1/K$. This is equivalent to saying that the average collision probability of $(Y_i|_{X_{<i}=x_{<i}})$ over the choice of the hash function H is small, i.e.,

$$\mathbb{E}_{h \leftarrow H} [\text{cp}(h(X_i)|_{X_{<i}=x_{<i}})] = \text{cp}((Y_i|_{X_{<i}=x_{<i}})|H) \leq \frac{1}{M} + \frac{1}{K}.$$

We can then use a Markov argument to say that for every block, with probability at least $1 - \varepsilon/T$ over $h \leftarrow H$, the collision probability is at most $1/M + T/(K\varepsilon)$. We can then take a union bound to say that for every $\mathbf{x} \in \text{supp}(\mathbf{X})$, at least $(1 - \varepsilon)$ -fraction of hash functions h are good in the sense that $\text{cp}(h(X_i)|_{X_{<i}=x_{<i}})$ is small for all blocks $i = 1, \dots, T$. [MV08] shows that if this condition is true for every $(h, \mathbf{x}) \in \text{supp}(H, \mathbf{X})$, then \mathbf{Y} is a block $(1/M + T/(K\varepsilon))$ -source, and thus the overall collision probability is at most $(1 + MT/K\varepsilon)^T/M^T$. [MV08] also shows how to modify an ε -fraction of the distribution to fix the bad hash functions, and thus complete the analysis.

The problem of the above analysis is that taking a Markov argument for each block, and then taking a union bound incurs a loss of factor T . To avoid this, we want to apply Markov argument only once to the whole sequence. For example, a natural thing to try is to sum over blocks to get

$$\mathbb{E}_{h \leftarrow H} \left[\frac{1}{T} \sum_{i=1}^T \text{cp}(h(X_i)|_{X_{<i}=x_{<i}}) \right] = \frac{1}{T} \sum_{i=1}^T \text{cp}((Y_i|_{X_{<i}=x_{<i}})|H) \leq \frac{1}{M} + \frac{1}{K},$$

and use a Markov argument to deduce that for every $\mathbf{x} \in \text{supp}(\mathbf{X})$, with probability $1 - \varepsilon$ over $h \leftarrow H$, the average collision probability per block satisfies

$$\frac{1}{T} \cdot \sum_{i=1}^T \text{cp}(h(X_i)|_{X_{<i}=x_{<i}}) \leq \frac{1}{M} + \frac{1}{K\varepsilon}.$$

We need to bound the collision probability of \mathbf{Y} using this information. We may try to apply Lemma 2.2, but it requires the information on $(1/T) \sum_i \text{cp}(Y_i|_{Y_{<i}=y_{<i}})$ instead of $(1/T) \sum_i \text{cp}(h(X_i)|_{X_{<i}=x_{<i}})$. That is, Lemma 2.2 requires us to condition on previous *hashed values* $Y_{<i}$, whereas the above argument refers to conditioning on the un-hashed values $X_{<i}$. The difficulty with directly reasoning about the former is that conditioned on the hashed values $Y_{<i}$, the hash function H may no longer be uniform (as it is correlated with $Y_{<i}$) and thus the Leftover Hash Lemma no longer applies.

To get around with the issues, we work with the averaged form of conditional collision probability $\text{cp}(Y_i|H, Y_{<i})$, as from Definition 2.4. Our key observation is that now we can apply Lemma 2.5 to deduce that for every block $i \in [T]$, the conditional collision probability satisfies $\text{cp}(Y_i|H, Y_{<i}) \leq \text{cp}(Y_i|H, X_{<i}) \leq 1/M + 1/K$. Then, by a Markov argument, it follows that with probability $1 - \varepsilon$ over $(h, \mathbf{y}) \leftarrow (H, \mathbf{Y})$, the average collision probability satisfies

$$\frac{1}{T} \sum_{i=1}^T \text{cp}(Y_i|_{(H, Y_{<i})=(h, y_{<i})}) \leq \frac{1}{M} + \frac{1}{K\varepsilon}.$$

We can then modify an ε -fraction of distribution, and apply Lemma 2.2 to complete the analysis.

The following lemma formalizes our claim about that the conditional collision probability of every block of (H, \mathbf{Y}) is small.

Lemma 3.3 *Let $H : [N] \rightarrow [M]$ be a random hash function from a 2-universal family \mathcal{H} . Let $\mathbf{X} = (X_1, \dots, X_T)$ be a block K -source over $[N]^T$. Let $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$. Then $\text{cp}(H) = 1/|\mathcal{H}|$ and for every $i \in [T]$, $\text{cp}(Y_i|H, Y_{<i}) \leq 1/M + 1/K$.*

Proof. $\text{cp}(H) = 1/|\mathcal{H}|$ is trivial since H is the uniform distribution. Fix $i \in [T]$. By the definition of block K -source, for every $x_{<i}$ in the support of $X_{<i}$, $\text{cp}(X_i|_{X_{<i}=x_{<i}}) \leq 1/K$. By the Leftover Hash Lemma, we have $\text{cp}((Y_i|_{X_{<i}=x_{<i}})|(H|_{X_{<i}=x_{<i}})) \leq 1/M + 1/K$ for every $x_{<i}$. It follows that $\text{cp}(Y_i|H, X_{<i}) \leq 1/M + 1/K$. Now, we can think of $(Y_i|H, X_{<i})$ as Y_i first conditioning on $(H, Y_{<i})$, and then further conditioning on $X_{<i}$. By Lemma 2.5, we have

$$\text{cp}(Y_i|H, Y_{<i}) \leq \text{cp}(Y_i|H, Y_{<i}, X_{<i}) = \text{cp}(Y_i|H, X_{<i}) \leq 1/M + 1/K,$$

as desired. ■

We use this to prove Theorem 3.1 as outlined above.

Proof of Theorem 3.1: By Lemma 3.3, for every $i \in [T]$, we have

$$\mathbb{E}_{(h, \mathbf{y}) \leftarrow (H, \mathbf{Y})} [\text{cp}(Y_i|_{(H, Y_{<i})=(h, y_{<i})})] = \text{cp}(Y_i|H, Y_{<i}) \leq \frac{1}{M} + \frac{1}{K}.$$

By linearity of expectation, the average conditional collision probability is also small.

$$\mathbb{E}_{(h, \mathbf{y}) \leftarrow (H, \mathbf{Y})} \left[\frac{1}{T} \sum_{i=1}^T \text{cp}(Y_i|_{(H, Y_{<i})=(h, y_{<i})}) \right] \leq \frac{1}{M} + \frac{1}{K}.$$

Note that the collision probability of a random variable over $[M]$ is at least $1/M$. Thus, Markov's inequality implies that with probability at least $1 - \varepsilon$ over $(h, \mathbf{y}) \leftarrow (H, \mathbf{Y})$,

$$\frac{1}{T} \sum_{i=1}^T \text{cp}(Y_i |_{(H, Y_{<i})=(h, y_{<i})}) \leq \frac{1}{M} + \frac{1}{K\varepsilon} = \frac{1}{M} \cdot \left(1 + \frac{M}{K\varepsilon}\right). \quad (1)$$

In Lemma 3.4 below, we show how to fix the *bad* (h, \mathbf{y}) 's by modifying at most ε -fraction of the distribution. Formally, Lemma 3.4 says that there exists a distribution $(H, \mathbf{Z}) = (H, Z_1, \dots, Z_T)$ such that (H, \mathbf{Y}) is ε -close to (H, \mathbf{Z}) , and for every $(h, \mathbf{z}) \leftarrow (H, \mathbf{Z})$,

$$\frac{1}{T} \sum_{i=1}^T \text{cp}(Z_i |_{(H, Z_{<i})=(h, z_{<i})}) \leq \frac{1}{M} \cdot \left(1 + \frac{M}{K\varepsilon}\right).$$

Applying Lemma 2.2 on $(\mathbf{Z}|_{H=h})$ for every $h \in \text{supp}(\mathcal{H})$, we have

$$\text{cp}(H, \mathbf{Z}) = \frac{1}{|\mathcal{H}|} \cdot \mathbb{E}_{h \leftarrow H} [\text{cp}(\mathbf{Z}|_{H=h})] \leq \frac{1}{|\mathcal{H}| \cdot M^T} \cdot \left(1 + \frac{M}{K\varepsilon}\right)^T.$$

■

Lemma 3.4 *Let $(H, \mathbf{Y}) = (H, Y_1, \dots, Y_T)$ be jointly distributed random variables over $\mathcal{H} \times [M]^T$ such that with probability at least $1 - \varepsilon$ over $(h, \mathbf{y}) \leftarrow (H, \mathbf{Y})$, the average conditional collision probability satisfies*

$$\frac{1}{T} \cdot \sum_{i=1}^T \text{cp}(Y_i |_{(H, Y_{<i})=(h, y_{<i})}) \leq \frac{1}{M} + \alpha.$$

Then there exists a distribution $(H, \mathbf{Z}) = (H, Z_1, \dots, Z_T)$ such that (H, \mathbf{Z}) is ε -close to (H, \mathbf{Y}) , and for every $(h, \mathbf{z}) \leftarrow (H, \mathbf{Z})$, we have

$$\frac{1}{T} \cdot \sum_{i=1}^T \text{cp}(Z_i |_{(H, Z_{<i})=(h, z_{<i})}) \leq \frac{1}{M} + \alpha.$$

Furthermore, the marginal distribution of H is unchanged.

Proof. We define the distribution (H, \mathbf{Z}) as follows.

- Sample $(h, \mathbf{y}) \leftarrow (H, \mathbf{Y})$.
- If $(1/T) \cdot \sum_{i=1}^T \text{cp}(Y_i |_{(H, Y_{<i})=(h, y_{<i})}) \leq 1/M + \alpha$, then output (h, \mathbf{y}) .

- Otherwise, let $j \in [T]$ be the least index such that

$$\frac{1}{T} \sum_{i=1}^j \left(\text{cp}(Y_i |_{(H, Y_{<i})=(h, y_{<i})}) - \frac{1}{M} \right) \leq \alpha \text{ and } \frac{1}{T} \sum_{i=1}^{j+1} \left(\text{cp}(Y_i |_{(H, Y_{<i})=(h, y_{<i})}) - \frac{1}{M} \right) > \alpha$$

- Choose $w_{j+1}, \dots, w_T \leftarrow U_{[M]}$, and output $(h, y_1, \dots, y_j, w_{j+1}, \dots, w_T)$.

It is easy to check that (i) (H, \mathbf{Z}) is well-defined, (ii) (H, \mathbf{Y}) is ε -close to (H, \mathbf{Z}) , (iii) for every $(h, \mathbf{z}) \in (H, \mathbf{Z})$, $(1/T) \cdot \sum_{i=1}^T \text{cp}(Z_i |_{(H, Z_{<i})=(h, z_{<i})}) \leq 1/M + \alpha$, and (iv) the marginal distribution of H is unchanged. \blacksquare

3.2 Small Collision Probability Using 4-wise Independent Hash Functions

As discussed in [MV08], using 4-wise independent hash functions $H : [N] \rightarrow [M]$ from \mathcal{H} , we can further reduce the required randomness in the data $\mathbf{X} = (X_1, \dots, X_T)$. [MV08] shows that in this case, $K \geq MT + \sqrt{2MT^3/\varepsilon}$ is enough for the hashed sequence (H, \mathbf{Y}) to be ε -close to having collision probability $O(1/|\mathcal{H}| \cdot M^T)$. As discussed in the previous subsection, by avoiding using union bounds, we show that $K \geq MT + \sqrt{2MT^2/\varepsilon}$ suffices. (Taking logs yields the second entry in Table 1, i.e. it suffices to have Renyi entropy $k = \max\{m + \log T, (1/2) \cdot (m + 2 \log T + \log(1/\varepsilon))\} + O(1)$ per block.) Formally, we prove the following theorem.

Theorem 3.5 *Let $H : [N] \rightarrow [M]$ be a random hash function from a 4-wise independent family \mathcal{H} . Let $\mathbf{X} = (X_1, \dots, X_T)$ be a block K -source over $[N]^T$. For every $\varepsilon > 0$, the hashed sequence $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$ is ε -close to a distribution $(H, \mathbf{Z}) = (H, Z_1, \dots, Z_T)$ such that*

$$\text{cp}(H, \mathbf{Z}) \leq \frac{1}{|\mathcal{H}| \cdot M^T} \left(1 + \frac{M}{K} + \sqrt{\frac{2M}{K^2\varepsilon}} \right)^T.$$

In particular, if $K \geq MT + \sqrt{2MT^2/\varepsilon}$, then (H, \mathbf{Z}) has collision probability at most $(1 + \gamma)/(|\mathcal{H}| \cdot M^T)$ for $\gamma = 2 \cdot (MT + \sqrt{2MT^2/\varepsilon})/K$.

The improvement of Theorem 3.5 over Theorem 3.1 comes from that when we use 4-wise independent hash families, we have a concentration result on the conditional collision probability for each block, via the following lemma.

Lemma 3.6 ([MV08]) *Let $H : [N] \rightarrow [M]$ be a random hash function from a 4-wise independent family \mathcal{H} , and X a random variable over $[N]$ with $\text{cp}(X) \leq 1/K$. Then we have*

$$\text{Var}_{h \leftarrow H}[\text{cp}(h(X))] \leq \frac{2}{MK^2}.$$

We can then replace the application of Markov's Inequality in the proof of Theorem 3.1 by Chebychev's Inequality to get stronger result. Formally, we prove the following lemma, which suffices to prove Theorem 3.5.

Lemma 3.7 *Let $H : [N] \rightarrow [M]$ be a random hash function from a 4-wise independent family \mathcal{H} . Let $\mathbf{X} = (X_1, \dots, X_T)$ be a block K -source over $[N]^T$. Let $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$. Then with probability at least $1 - \varepsilon$ over $(h, \mathbf{y}) \leftarrow (H, \mathbf{Y})$,*

$$\frac{1}{T} \sum_{i=1}^T \text{cp}(Y_i |_{(H, Y_{<i})=(h, y_{<i})}) \leq \frac{1}{M} \cdot \left(1 + \frac{M}{K} + \sqrt{\frac{2M}{K^2 \varepsilon}} \right).$$

Theorem 3.5 follows immediately by composing Lemma 3.7, 3.4, and 2.2 in the same way as the proof of Theorem 3.1.

Proof of Lemma 3.7: Recall that we have

$$\mathbb{E}_{(h, \mathbf{y}) \leftarrow (H, \mathbf{Y})} \left[\frac{1}{T} \sum_{i=1}^T \text{cp}(Y_i |_{(H, Y_{<i})=(h, y_{<i})}) \right] \leq \frac{1}{M} + \frac{1}{K}.$$

Hence, our goal is to upper bound the probability of the value $(1/T) \sum_{i=1}^T \text{cp}(Y_i |_{(H, Y_{<i})=(h, y_{<i})})$ deviating from its mean by $\sqrt{2/MK^2\varepsilon}$. Our strategy is to bound the variance of a properly defined random variable, and then apply Chebychev's Inequality. By Lemma 3.6, the information we get from 4-wise independent hash function is that for every $i \in [T]$, we have

$$\text{Var}_{h \leftarrow H} [\text{cp}(Y_i |_{(H, X_{<i})=(h, x_{<i})})] \leq \frac{2}{MK^2} \quad \forall x_{<i} \in \text{supp}(X_{<i}) \quad (2)$$

Fix $i \in [T]$, let us try to bound the variance of the i -th block. There are two issues to take care of. Firstly, the variance we have is conditioning on $X_{<i}$ instead of $Y_{<i}$. Secondly, even when conditioning on $X_{<i}$, it is possible that the variance is

$$\text{Var}_{(h, \mathbf{x}) \leftarrow (H, \mathbf{X})} [\text{cp}(Y_i |_{(H, X_{<i})=(h, x_{<i})})] = \Omega\left(\frac{1}{K^2}\right) \gg \frac{2}{MK^2}.$$

The reason is that conditioning on different $X_{<i} = x_{<i}$, the collision probability of $(Y_i |_{X_{<i}=x_{<i}})$ may have different expectation over $h \leftarrow \mathcal{H}$. Thus, we have to subtract the mean first. Let us define

$$f(h, x_{<i}) = \text{cp}(Y_i |_{(H, X_{<i})=(h, x_{<i})}) - \mathbb{E}_{h \leftarrow H} [\text{cp}(Y_i |_{(H, X_{<i})=(h, x_{<i})})]$$

Now, for every $x_{<i} \in \text{supp}(X_{<i})$, $f(H, x_{<i})$ has mean 0, and variance $\leq 2/MK^2$. It follows that

$$\text{Var}_{(h, \mathbf{x}) \leftarrow (H, \mathbf{X})} [f(h, x_{<i})] \leq \frac{2}{MK^2}.$$

We now deal with the issue of conditioning on $X_{<i}$ versus $Y_{<i}$. Let us define

$$g(h, y_{<i}) = \mathbf{E}_{x_{<i} \leftarrow (X_{<i} |_{(H, Y_{<i}) = (h, y_{<i})})} [f(h, x_{<i})].$$

We claim that

$$\text{cp}(Y_i |_{(H, Y_{<i}) = (h, y_{<i})}) \leq \frac{1}{M} + \frac{1}{K} + g(h, y_{<i}).$$

Indeed, by Lemma 2.5 and the definition of f and g ,

$$\begin{aligned} & \text{cp}(Y_i |_{(H, Y_{<i}) = (h, y_{<i})}) \\ & \leq \text{cp}((Y_i |_{(H, Y_{<i}) = (h, y_{<i})}) | (X_i |_{(H, Y_{<i}) = (h, y_{<i})})) \\ & = \mathbf{E}_{x_{<i} \leftarrow (X_{<i} |_{(H, Y_{<i}) = (h, y_{<i})})} [\text{cp}(Y_i |_{(H, X_{<i}) = (h, x_{<i})})] \\ & = \mathbf{E}_{x_{<i} \leftarrow (X_{<i} |_{(H, Y_{<i}) = (h, y_{<i})})} \left[f(h, x_{<i}) + \mathbf{E}_{h \leftarrow H} [\text{cp}(Y_i |_{(H, X_{<i}) = (h, x_{<i})})] \right] \\ & \leq g(h, y_{<i}) + \frac{1}{M} + \frac{1}{K} \end{aligned}$$

Also note that $g(H, Y_{<i})$ has mean 0 and small variance:

$$\begin{aligned} \mathbf{E}_{(h, y_{<i}) \leftarrow (H, Y_{<i})} [g(h, y_{<i})] &= \mathbf{E}_{(h, \mathbf{x}) \leftarrow (H, \mathbf{X})} [f(h, x_{<i})] = 0, \\ \text{Var}_{(h, y_{<i}) \leftarrow (H, Y_{<i})} [g(h, y_{<i})] &\leq \text{Var}_{(h, \mathbf{x}) \leftarrow (H, \mathbf{X})} [f(h, x_{<i})] \leq \frac{2}{MK^2}. \end{aligned}$$

The above argument holds for every block $i \in [T]$. Taking average over blocks, we get

$$\begin{aligned} \mathbf{E}_{(h, \mathbf{y}) \leftarrow (H, \mathbf{Y})} \left[\frac{1}{T} \sum_{i=1}^T g(h, y_{<i}) \right] &= 0, \\ \text{Var}_{(h, \mathbf{y}) \leftarrow (H, \mathbf{Y})} \left[\frac{1}{T} \sum_{i=1}^T g(h, y_{<i}) \right] &\leq \frac{2}{MK^2}, \text{ and} \\ \frac{1}{T} \sum_{i=1}^T \text{cp}(Y_i |_{(H, Y_{<i}) = (h, y_{<i})}) &\leq \frac{1}{M} + \frac{1}{K} + \left(\frac{1}{T} \sum_{i=1}^T g(h, y_{<i}) \right). \end{aligned}$$

Finally, we can apply Chebychev's Inequality to random variable $(1/T) \cdot \sum_i g(H, Y_{<i})$ to get the desired result: with probability $1 - \varepsilon$ over $(h, \mathbf{y}) \leftarrow (H, \mathbf{Y})$,

$$\frac{1}{T} \sum_{i=1}^T \text{cp}(Y_i |_{(H, Y_{<i}) = (h, y_{<i})}) \leq \frac{1}{M} \cdot \left(1 + \frac{M}{K} + \sqrt{\frac{2M}{K^2 \varepsilon}} \right).$$

■

3.3 Statistical Distance to Uniform Distribution

Let $H : [N] \rightarrow [M]$ be a random hash function from a 2-universal family \mathcal{H} . Let $\mathbf{X} = (X_1, \dots, X_T)$ be a block K -source over $[N]^T$. In this subsection, we study the statistical distance between the distribution of hashed sequence $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$ and the uniform distribution $(H, U_{[M]^T})$. Classic results of [CG88, ILL89, Zuc96] show that if $K \geq MT^2/\varepsilon^2$, then (H, \mathbf{Y}) is ε -close to uniform. The proof idea is as follows. The Leftover Hash Lemma together with Lemma 2.3 tells us that the joint distribution of hash function and a hashed value $(H, Y_i) = (H, H(X_i))$ is $\sqrt{M/K}$ -close to uniform $U_{[M]}$ even conditioning on the previous blocks $X_{<i}$. One can then use a hybrid argument to show that the distance grows linearly with the number of blocks, so (H, \mathbf{Y}) is $T \cdot \sqrt{M/K}$ -close to uniform. Taking $K \geq MT^2/\varepsilon^2$ completes the analysis.

We save a factor of T , and show that in fact, $K = MT/\varepsilon^2$ is sufficient. (Taking logs yields the third entry in Table 1, i.e. it suffices to have Renyi entropy $k = m + \log T + 2 \log(1/\varepsilon)$ per block.) Formally, we prove the following theorem.

Theorem 3.8 *Let $H : [N] \rightarrow [M]$ be a random hash function from a 2-universal family \mathcal{H} . Let $\mathbf{X} = (X_1, \dots, X_T)$ be a block K -source over $[N]^T$. For every $\varepsilon > 0$ such that $K > MT/\varepsilon^2$, the hashed sequence $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$ is ε -close to uniform $(H, U_{[M]^T})$.*

Recall that the previous analysis goes by passing to statistical distance first, and then measuring the growth of distance using statistical distance. This incurs a quadratic dependency of K on T . Since without further information, the hybrid argument is tight, to save a factor of T , we have to measure the increase of distance over blocks in another way, and pass to statistical distance only in the end. It turns out that the *Hellinger distance* (cf., [GS02]) is a good measure for our purposes:

Definition 3.9 (Hellinger distance) *Let X and Y be two random variables over $[M]$. The Hellinger distance between X and Y is*

$$d(X, Y) \stackrel{\text{def}}{=} \left(\frac{1}{2} \sum_i (\sqrt{\Pr[X=i]} - \sqrt{\Pr[Y=i]}) \right)^{1/2} = \sqrt{1 - \sum_i \sqrt{\Pr[X=i] \cdot \Pr[Y=i]}}.$$

Like statistical distance, Hellinger distance is a distance measure for distributions, and it takes value in $[0, 1]$. The following standard lemma says that the two distance measures are closely related. We remark that the lemma is tight in both directions even if Y is the uniform distribution.

Lemma 3.10 (cf., [GS02]) *Let X and Y be two random variables over $[M]$. We have*

$$d(X, Y)^2 \leq \Delta(X, Y) \leq \sqrt{2} \cdot d(X, Y).$$

In particular, the lemma allows us to upper-bound the statistical distance by upper-bounding the Hellinger distance. Since our goal is to bound the distance to uniform, it is convenient to introduce the following definition.

Definition 3.11 (Hellinger Closeness to Uniform) *Let X be a random variable over $[M]$. The Hellinger closeness of X to uniform $U_{[M]}$ is*

$$C(X) \stackrel{\text{def}}{=} \frac{1}{M} \sum_i \sqrt{M \cdot \Pr[X = i]} = 1 - d(X, U_{[M]})^2.$$

Note that $C(X, Y) = C(X) \cdot C(Y)$ when X and Y are independent random variables, so the Hellinger closeness is well-behaved with respect to products (unlike statistical distance). By Lemma 3.10, if the Hellinger closeness $C(X)$ is close to 1, then X is close to uniform in statistical distance. Recall that collision probability behaves similarly. If the collision probability $\text{cp}(X)$ is close to $1/M$, then X is close to uniform. In fact, by the following normalization, we can view the collision probability as the 2-norm of X , and the Hellinger closeness as the 1/2-norm of X .

Let $f(i) = M \cdot \Pr[X = i]$ for $i \in [M]$. In terms of $f(\cdot)$, the collision probability is $\text{cp}(X) = (1/M^2) \cdot \sum_i f(i)^2$, and Lemma 2.3 says that if the “2-norm” $M \cdot \text{cp}(X) = \mathbb{E}_i[f(i)^2] \leq 1 + \varepsilon$ where the expectation is over uniform $i \in [M]$, then $\Delta(X, U) \leq \sqrt{\varepsilon}$. Similarly, Lemma 3.10 says that if the “1/2-norm” $C(X) = \mathbb{E}_i[\sqrt{f(i)}] \geq 1 - \varepsilon$, then $\Delta(X, U) \leq \sqrt{\varepsilon}$.

We now discuss our approach to prove Theorem 3.8. We want to show that (H, \mathbf{Y}) is close to uniform. All we know is that the conditional collision probability $\text{cp}(Y_i | H, Y_{<i})$ is close to $1/M$ for every block. If all blocks are independent, then the overall collision probability $\text{cp}(H, \mathbf{Y})$ is small, and so (H, \mathbf{Y}) is close to uniform. However, this is not true without independence, since 2-norm tends to over-weight heavy elements. In contrast, the 1/2-norm does not suffer this problem. Therefore, our approach is to show that small conditional collision probability implies large Hellinger closeness. Formally, we have the following lemma.

Lemma 3.12 *Let $\mathbf{X} = (X_1, \dots, X_T)$ be jointly distributed random variables over $[M_1] \times \dots \times [M_T]$ such that $\text{cp}(X_i | X_{<i}) \leq \alpha_i / M_i$ for every $i \in [T]$. Then the Hellinger closeness satisfies*

$$C(\mathbf{X}) \geq \sqrt{\frac{1}{\alpha_1 \dots \alpha_T}}.$$

With this lemma, the proof of Theorem 3.8 is immediate.

Proof of Theorem 3.8: By Lemma 3.3, $\text{cp}(H) = 1/|\mathcal{H}|$, and $\text{cp}(Y_i | H, Y_{<i}) \leq (1 + M/K)/M$ for every $i \in [T]$. By Lemma 3.12, the Hellinger closeness satisfies $C(H, \mathbf{Y}) \geq (1 + M/K)^{-T/2} \geq 1 - MT/2K$ (recall that $K \geq MT/\varepsilon^2$). It follows by Lemma 3.10 that

$$\Delta((H, \mathbf{Y}), (H, U_{[M]^T})) \leq \sqrt{2} \cdot d((H, \mathbf{Y}), (H, U_{[M]^T})) = \sqrt{2} \cdot \sqrt{1 - C(H, \mathbf{Y})} \leq \sqrt{MT/K} \leq \varepsilon. \quad \blacksquare$$

We proceed to prove Lemma 3.12. The main idea is to use Hölder's inequality to relate two different norms. We recall Hölder's inequality first.

Lemma 3.13 (Hölder's inequality [Dur04])

- Let F, G be two non-negative functions from $[M]$ to \mathbb{R} , and $p, q > 0$ satisfying $1/p + 1/q = 1$. Let x be a uniformly random index over $[M]$. We have

$$\mathbb{E}_x[F(x) \cdot G(x)] \leq \mathbb{E}_x[F(x)^p]^{1/p} \cdot \mathbb{E}_x[G(x)^q]^{1/q}.$$

- In general, let F_1, \dots, F_n be non-negative functions from $[M]$ to \mathbb{R} , and $p_1, \dots, p_n > 0$ satisfying $1/p_1 + \dots + 1/p_n = 1$. We have

$$\mathbb{E}_x[F_1(x) \cdots F_n(x)] \leq \mathbb{E}_x[F_1(x)^{p_1}]^{1/p_1} \cdots \mathbb{E}_x[F_n(x)^{p_n}]^{1/p_n}.$$

Proof of Lemma 3.12: We prove it by induction on T . The base case $T = 1$ is already non-trivial. Let X be a random variable over $[M]$ with $\text{cp}(X) \leq \alpha/M$, we need to show that the Hellinger closeness $C(X) \geq \sqrt{1/\alpha}$. Recall the normalization we mentioned before. Let $f(x) = M \cdot \Pr[X = x]$ for every $x \in [M]$. In terms of $f(\cdot)$, we want to show that $\mathbb{E}_x[f(x)^2] \leq \alpha$ implies $\mathbb{E}_x[\sqrt{f(x)}] \geq \sqrt{1/\alpha}$. Note that $\mathbb{E}_x[f(x)] = 1$. We now apply Hölder's inequality with $F = f^{2/3}$, $G = f^{1/3}$, $p = 3$, and $q = 3/2$. We have

$$\mathbb{E}_x[f(x)] \leq \mathbb{E}_x[f(x)^2]^{1/3} \cdot \mathbb{E}_x[f(x)^{1/2}]^{2/3},$$

which implies

$$C(X) = \mathbb{E}_x[\sqrt{f(x)}] \geq \mathbb{E}_x[f(x)]^{3/2} / \mathbb{E}_x[f(x)^2]^{1/2} \geq \sqrt{1/\alpha}.$$

Suppose the lemma is true for $T - 1$, we show that it is true for T . Let $f(x) = M_1 \cdot \Pr[X_1 = x]$. To apply the induction hypothesis, we consider the conditional random variables $(X_2, \dots, X_T |_{X_1=x})$ for every $x \in [M_1]$. For every $x \in [M_1]$ and $j = 2, \dots, T$, we define $g_j(x) = M_j \cdot \text{cp}((X_j |_{X_1=x}) | (X_2, \dots, X_{j-1} |_{X_1=x}))$ to be the "normalized" conditional collision probability. By induction hypothesis, we have $C(X_2, \dots, X_T |_{X_1=x}) \geq \sqrt{g_2(x) \cdots g_T(x)}$ for every $x \in [M_1]$. It follows that

$$C(\mathbf{X}) = \mathbb{E}_x[\sqrt{f(x)} \cdot C(X_2, \dots, X_T |_{X_1=x})] \geq \mathbb{E}_x[\sqrt{f(x)/g_2(x) \cdots g_T(x)}].$$

We use Hölder's inequality twice to show that $\mathbb{E}_x[\sqrt{f(x)/g_2(x) \cdots g_T(x)}] \geq \sqrt{1/\alpha_1 \cdots \alpha_T}$. Let us first summarize the constraints we have. By definition, we have $\mathbb{E}_x[f(x)^2] \leq \alpha_1$. Fix

$j \in \{2, \dots, T\}$. Note that

$$\begin{aligned}
& \text{cp}(X_j | X_{<j}) \\
&= \mathbb{E}_{x \leftarrow X_1} [\text{cp}((X_j | X_1=x) | (X_2, \dots, X_{j-1} | X_1=x))] \\
&= \mathbb{E}_{x \leftarrow X_1} [g_j(x)/M_j] \\
&= \mathbb{E}_{x \leftarrow U_{[M]}} [f(x)g_j(x)]/M_j
\end{aligned}$$

It follows that $\mathbb{E}_x[f(x)g_j(x)] \leq \alpha_j$ for $j = 2, \dots, T$. Now, we apply the second version of Hölder's Inequality with $F_1 = (f/g_2 \cdots g_T)^{1/2}$, $F_j = (fg_j)^{1/(T+1)}$ for $j = 2, \dots, T$, $p_1 = 2/(T+1)$, and $p_j = 1/(T+1)$ for $j = 2, \dots, T$, which gives

$$\mathbb{E}_x [f(x)^{T/(T+1)}] \leq \mathbb{E}_x \left[\sqrt{f(x)/g_2(x) \cdots g_T(x)} \right]^{2/(T+1)} \cdot \mathbb{E}_x [f(x)g_2(x)]^{1/(T+1)} \cdots \mathbb{E}_x [f(x)g_T(x)]^{1/(T+1)},$$

so

$$\begin{aligned}
\mathbb{E}_x \left[\sqrt{f(x)/g_2(x) \cdots g_T(x)} \right] &\geq \mathbb{E}_x \left[f(x)^{T/(T+1)} \right]^{(T+1)/2} \cdot \prod_{j=2}^T \mathbb{E}_x [f(x)g_j(x)]^{-1/2} \\
&\geq \mathbb{E}_x \left[f(x)^{T/(T+1)} \right]^{(T+1)/2} \cdot \sqrt{1/\alpha_2 \cdots \alpha_T}.
\end{aligned}$$

It remains to lower bound the first term by $\sqrt{1/\alpha_1}$. We apply Hölder again with $F = f^{2/(T+2)}$, $G = f^{T/(T+2)}$, $p = T+2$, and $q = (T+2)/(T+1)$, which gives

$$\mathbb{E}_x [f(x)] \leq \mathbb{E}_x [f(x)^2]^{1/(T+2)} \cdot \mathbb{E}_x \left[f(x)^{T/(T+1)} \right]^{(T+1)/(T+2)},$$

so

$$\mathbb{E}_x \left[f(x)^{T/(T+1)} \right]^{(T+1)/2} \geq \mathbb{E}_x [f(x)]^{(T+2)/2} / \mathbb{E}_x [f(x)^2]^{1/2} \geq \sqrt{1/\alpha_1}.$$

Combining the inequalities, we have $C(\mathbf{X}) \geq \sqrt{1/\alpha_1 \cdots \alpha_T}$. ■

4 Negative Results: How Much Entropy is Necessary?

In this section, we provide lower bounds on the entropy needed for the data items. We show that if K is not large enough, then for every hash family \mathcal{H} , there exists a block K -source $\mathbf{X} = (X_1, \dots, X_T)$ such that the hashed sequence $\mathbf{Y} = (H(X_1), \dots, H(X_T))$ do not satisfy the desired closeness requirements to uniform (possibly in conjunction with the hash function H).

4.1 Lower Bound for Statistical Distance to Uniform Distribution

Let us first consider the requirement for the joint distribution of (H, \mathbf{Y}) being ε -close to uniform. When there is only one block, this is exactly the requirement for a “strong extractor”. The lower bound in the extractor literature, due to Radhakrishnan and Ta-Shma [RT00] shows that $K \geq \Omega(M/\varepsilon^2)$ is necessary, which is tight up to a constant factor. Our goal is to show that when hashing T blocks, the value of K required for each block increases by a factor of T . Intuitively, each block will produce some error (i.e., the hashed value is not close to uniform), and the overall error will accumulate over the blocks, so we need to inject more randomness per block to reduce the error. Indeed, we use this intuition to show that $K \geq \Omega(MT/\varepsilon^2)$ is necessary for the hashed sequence to be ε -close to uniform, matching the upper bound in Theorem 3.8. Note that the lower bound holds even for a truly random hash family. Formally, we prove the following theorem.

Theorem 4.1 *Let N, M , and T be positive integers and $\varepsilon \in (0, \varepsilon_0)$ a real number such that $N \geq MT/\varepsilon^2$, where $\varepsilon_0 > 0$ is a small absolute constant. Let $H : [N] \rightarrow [M]$ be a random hash function from an hash family \mathcal{H} . Then there exists an integer $K = \Omega(MT/\varepsilon^2)$, and a block K -source $\mathbf{X} = (X_1, \dots, X_T)$ such that $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$ is ε -far from uniform $(H, U_{[M]^T})$ in statistical distance.*

To prove the theorem, we need to find such an \mathbf{X} for every hash family \mathcal{H} . Following the intuition, we find an X that incurs certain error on a single block, and take $\mathbf{X} = (X_1, \dots, X_T)$ to be T i.i.d. copies of X . More precisely, we first find a K -source X such that for $\Omega(1)$ -fraction of hash functions $h \in \mathcal{H}$, $h(X)$ is $\Omega(\varepsilon/\sqrt{T})$ -far from uniform. This step is the same as the lower bound proof for extractors [RT00], which uses the probabilistic method. We pick X to be a random *flat* K -source, i.e., a uniform distribution over a random set of size K , and show that X satisfies the desired property with nonzero probability. The next step is to measure how the error accumulates over independent blocks. Note that for a fixed hash function h , the hashed sequence $(h(X_1), \dots, h(X_T))$ consists of T i.i.d. copies of $h(X)$. Reyzin [Rey04] has shown that the statistical distance increases \sqrt{T} when we have T independent copies for small T . However, Reyzin’s result only shows an increase up to distance $O(\delta^{1/3})$, where δ is the statistical distance of the original random variables. We improve Reyzin’s result to show that the $\Omega(\sqrt{T})$ growth continues until the distance reaches some absolute constant. We then use it to show that the joint distribution (H, \mathbf{Y}) is far from uniform.

The following lemma corresponds to the first step.

Lemma 4.2 *Let N and M be positive integers and $\varepsilon \in (0, 1/4), \delta \in (0, 1)$ real numbers such that $N \geq M/\varepsilon^2$. Let $H : [N] \rightarrow [M]$ be a random hash function from an hash family \mathcal{H} . Then there exists an integer $K = \Omega(\delta^2 M/\varepsilon^2)$, and a flat K -source X over $[N]$ such that with probability at least $1 - \delta$ over $h \leftarrow H$, $h(X)$ is ε -far from uniform.*

Proof. Let $K = \lfloor \min\{\alpha \cdot M/\varepsilon^2, N/2\} \rfloor$ for some α to be determined later. Let X be a random flat K -source over $[N]$. That is, $X = U_S$ where $S \subset [N]$ is a uniformly random size K subset of $[N]$. We claim that for every hash function $h : [N] \rightarrow [M]$,

$$\Pr_S [h(U_S) \text{ is } \varepsilon\text{-far from uniform}] \geq 1 - c \cdot \sqrt{\alpha} \quad (3)$$

for some absolute constant c . Let us assume (3), and prove the lemma first. Since the claim holds for every hash function h ,

$$\Pr_{h \leftarrow H, S} [h(U_S) \text{ is } \varepsilon\text{-far from uniform}] \geq 1 - c \cdot \sqrt{\alpha}.$$

Thus, there exists a flat K -source U_S such that

$$\Pr_{h \leftarrow H} [h(U_S) \text{ is } \varepsilon\text{-far from uniform}] \geq 1 - c \cdot \sqrt{\alpha}.$$

The lemma follows by setting $\alpha = \min\{\delta^2/c^2, 1/32\}$. We proceed to prove (3). It suffices to show that for every $y \in [M]$, with probability at least $1 - c' \cdot \sqrt{\alpha}$ over random U_S , the deviation of $\Pr[h(U_S) = y]$ from $1/M$ is at least $4\varepsilon/M$, where c' is another absolute constant. That is,

$$\Pr_S \left[\left| \Pr[h(U_S) = y] - \frac{1}{M} \right| \geq \frac{4\varepsilon}{M} \right] \geq 1 - c' \cdot \sqrt{\alpha}. \quad (4)$$

Again, let us see why (4) is sufficient to prove (3) first. Let us call $y \in [M]$ is *bad* for S if

$$\left| \Pr[h(U_S) = y] - \frac{1}{M} \right| \geq \frac{4\varepsilon}{M}.$$

Since Inequality (4) holds for every $y \in [M]$, we have

$$\Pr_{S,y} [y \text{ is bad for } S] \geq 1 - c' \cdot \sqrt{\alpha},$$

where y is uniformly random over $[M]$. It follows that

$$\Pr_S [\text{at least } 1/2\text{-fraction of } y \text{ are bad for } S] \geq 1 - 2c' \cdot \sqrt{\alpha}$$

Observe that if at least $1/2$ -fraction of y are bad for S , then $\Delta(h(X), U_{[M]}) \geq \varepsilon$. Inequality (3) follows by setting $c = 2c'$.

It remains to prove (4). Let $T = h^{-1}(y)$. We have $\Pr_S[h(U_S) = y] = |S \cap T|/|S|$. Thus, recall that $K \leq \alpha M/\varepsilon^2$, (4) follows from inequality

$$\Pr_S \left[\left| |S \cap T| - \frac{K}{M} \right| < \frac{4K\varepsilon}{M} \right] \leq c' \cdot \sqrt{\frac{K\varepsilon^2}{M}},$$

which follows by the claim below by setting $L = K/M$, and $\beta = 4\varepsilon\sqrt{K/M}$ (Working out the parameters, we have $c' = 4c''$, $\varepsilon < 1/4$ implies $\beta < \sqrt{L}$, and $\alpha \leq 1/32$ implies $\beta < 1$.)

Claim 4.3 Let $N, K > 1$ be positive integers such that $N > 2K$, and $L \in [0, K/2]$, $\beta \in (0, \min\{1, \sqrt{L}\})$ real numbers. Let $S \subset [N]$ be a random subset of size K , and $T \subset [N]$ be a fixed subset of arbitrary size. We have

$$\Pr_S \left[\left| |S \cap T| - L \right| \leq \beta \sqrt{L} \right] \leq c'' \cdot \beta,$$

for some absolute constant c'' .

Intuitively, the probability in the claim is maximized when the set T has size NL/K so that $L = \mathbb{E}_S[|S \cap T|]$, and the claim follows by observing that in this case, the distribution has deviation $\Theta(\sqrt{L})$, and each possible outcome has probability $O(\sqrt{1/L})$. The formal proof of the claim is in Appendix A and is proved by expressing the probability in terms of binomial coefficients, and estimating them using Stirling formula. ■

The next step is to measure the increase of statistical distance over independent random variables.

Lemma 4.4 Let X and Y be random variables over $[M]$ such that $\Delta(X, Y) \geq \varepsilon$. Let $\mathbf{X} = (X_1, \dots, X_T)$ be T i.i.d. copies of X , and let $\mathbf{Y} = (Y_1, \dots, Y_T)$ be T i.i.d. copies of Y . We have

$$\Delta(\mathbf{X}, \mathbf{Y}) \geq \min\{\varepsilon_0, c\sqrt{T} \cdot \varepsilon\},$$

where ε_0, c are absolute constants.

We defer the proof of the above lemma to Appendix B.

Proof of Theorem 4.1: The absolute constant ε_0 in the theorem is a half of the ε_0 in Lemma 4.4. By Lemma 4.2 there is a flat K -source such that for $1/2$ -fraction of hash functions $h \in \mathcal{H}$, $h(X)$ is $(2\varepsilon/c\sqrt{T})$ -far from uniform, for $K = \Omega((1/2)^2 M / (2\varepsilon/c\sqrt{T})^2) = \Omega(MT/\varepsilon^2)$. We set $\mathbf{X} = (X_1, \dots, X_T)$ to be T independent copies of X . Consider a hash function h such that $h(X)$ is $(2\varepsilon/c\sqrt{T})$ -far from uniform. By Lemma 4.4, $(h(X_1), \dots, h(X_T))$ is 2ε -far from uniform. Note that this holds for $1/2$ -fraction of hash function h . It follows that

$$\Delta((H, \mathbf{Y}), (H, U_{[M]^T})) = \mathbb{E}_{h \leftarrow H} \left[\Delta((h(X_1), \dots, h(X_T)), U_{[M]^T}) \right] \geq \frac{1}{2} \cdot 2\varepsilon = \varepsilon. \quad \blacksquare$$

4.2 Lower Bound for Small Collision Probability

In this subsection, we prove lower bounds on the entropy needed per item to ensure that the sequence of hashed values is close to having small collision probability. Since this requirement is less stringent than being close to uniform, less entropy is needed from the source. The interesting setting in applications is to require the hashed sequence $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$

to be ε -close to having collision probability $O(1/(|\mathcal{H}| \cdot M^T))$. Recall that in this setting, instead of requiring $K \geq MT/\varepsilon^2$, $K \geq \Omega(MT/\varepsilon)$ is sufficient for 2-universal hash functions (Theorem 3.1), and $K \geq \Omega(MT + T\sqrt{M/\varepsilon})$ is sufficient for 4-wise independent hash functions (Theorem 3.5). The main improvement from 2-universal to 4-wise independent hashing is the better dependency on ε . Indeed, it can be shown that if we use truly random hash functions, we can reduce the dependency on ε to $\log(1/\varepsilon)$. Since we are now proving lower bounds for arbitrary hash families, we focus on the dependency on M and T . Specifically, our goal is to show that $K = \Omega(MT)$ is necessary. More precisely, we show that when $K \ll MT$, it is possible for the hashed sequence (H, \mathbf{Y}) to be .99-far from any distribution that has collision probability less than $100/(|\mathcal{H}| \cdot M^T)$.

We use the same strategy as in the previous subsection to prove this lower bound. Fixing a hash family \mathcal{H} , we take T independent copies (X_1, \dots, X_T) of the worst-case X found in Lemma 4.2, and show that $(H, H(X_1), \dots, H(X_T))$ is far from having small collision probability. The new ingredient is to show that when we have T independent copies, and $K \ll MT$, then $(h(X_1), \dots, h(X_T))$ is very far from uniform (say, 0.99-far) for many $h \in \mathcal{H}$. We then argue that in this case, we can not reduce the collision probability of $(h(X_1), \dots, h(X_T))$ by changing a small fraction of distribution, which implies the overall distribution (H, \mathbf{Y}) is far from any distribution (H', \mathbf{Z}) with small collision probability. Formally, we prove the following theorem.

Theorem 4.5 *Let N, M , and T be positive integers such that $N \geq MT$. Let $\delta \in (0, 1)$ and $\alpha > 1$ be real numbers such that $\alpha < \delta^3 \cdot e^{T/32}/128$. Let $H : [N] \rightarrow [M]$ be a random hash function from a hash family \mathcal{H} . There exists an integer $K = \Omega(\delta^2 MT / \log(\alpha/\delta))$, and a block K -source $\mathbf{X} = (X_1, \dots, X_T)$ such that $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$ is $(1 - \delta)$ -far from any distribution (H', \mathbf{Z}) with $\text{cp}(H', \mathbf{Z}) \leq \alpha/(|\mathcal{H}| \cdot M^T)$.*

Think of α and δ as constants. Then the theorem says that $K = \Omega(MT)$ is necessary for the hashed sequence $(H, H(X_1), \dots, H(X_T))$ to be close to having small collision probability, matching the upper bound in Theorem 3.1. In the previous proof, we used Lemma 4.4 to measure the increase of distance over blocks. However, the lemma can only measure the progress up to some small constant. It is known that if the number of copies T is larger than $\Omega(1/\varepsilon^2)$, where ε is the statistical distance of original copy, then the statistical distance goes to 1 exponentially fast. Formally, we use the following lemma.

Lemma 4.6 ([SV99]) *Let X and Y be random variables over $[M]$ such that $\Delta(X, Y) \geq \varepsilon$. Let $\mathbf{X} = (X_1, \dots, X_T)$ be T i.i.d. copies of X , and let $\mathbf{Y} = (Y_1, \dots, Y_T)$ be T i.i.d. copies of Y . We have*

$$\Delta(\mathbf{X}, \mathbf{Y}) \geq 1 - e^{-T\varepsilon^2/2}.$$

We remark that Lemma 4.4 and 4.6 are incomparable. In the parameter range of Lemma 4.4, Lemma 4.6 only gives $\Delta(\mathbf{X}, \mathbf{Y}) \geq \Omega(T\varepsilon^2)$ instead of $\Omega(\sqrt{T}\varepsilon)$. To argue that the overall

distribution is far from having small collision probability, we introduce the following notion of nonuniformity.

Definition 4.7 Let X be a random variable over $[M]$ with probability mass function p . X is (δ, β) -nonuniform if for every function $q : [M] \rightarrow \mathbb{R}$ such that $0 \leq q(x) \leq p(x)$ for all $x \in [M]$, and $\sum_x q(x) \geq \delta$, the function satisfies

$$\sum_{x \in [M]} q(x)^2 > \beta/M.$$

Intuitively, a distribution X over $[M]$ is (δ, β) -nonuniform means that even if we remove $(1 - \delta)$ -fraction of probability mass from X , the ‘‘collision probability’’ remains greater than β/M . In particular, X is $(1 - \delta)$ -far from any random variable Y with $\text{cp}(Y) \leq \beta/M$.

Lemma 4.8 Let X be a random variable over $[M]$. If X is $(1 - \eta)$ -far from uniform, then X is $(2\sqrt{\beta \cdot \eta}, \beta)$ -nonuniform for every $\beta \geq 1$.

Proof. Let p be the probability mass function of X , and $q : [M] \rightarrow \mathbb{R}$ be a function such that $0 \leq q(x) \leq p(x)$ for every $x \in [M]$, and $\sum_x q(x) \geq 2\sqrt{\beta \cdot \eta}$. Our goal is to show that $\sum_x q(x)^2 > \beta/M$. Let $T = \{x \in [M] : p(x) \geq 1/M\}$. Note that

$$\Delta(X, U_{[M]}) = \Pr[X \in T] - \Pr[U_{[M]} \in T] \geq 1 - \eta.$$

This implies $\Pr[X \in T] \geq 1 - \eta$, and $\mu(T) = \Pr[U_{[M]} \in T] \leq \eta$. Now,

$$\sum_{x \in T} q(x) \geq 2\sqrt{\beta \cdot \eta} - \Pr[X \notin T] \geq 2\sqrt{\beta \cdot \eta} - \eta > \sqrt{\beta \cdot \eta},$$

and $\mu(T) \leq \eta$ implies

$$\sum_{x \in [M]} q(x)^2 \geq \sum_{x \in T} q(x)^2 \geq \frac{(\sum_{x \in T} q(x))^2}{|T|} > \frac{\beta}{M}.$$

■

We are ready to prove Theorem 4.5.

Proof of Theorem 4.5: By Lemma 4.2 with $\varepsilon = \sqrt{2 \ln(128\alpha/\delta^3)/T} < 1/4$, there is a flat K -source X such that for $(1 - \delta/4)$ -fraction of hash function $h \in \mathcal{H}$, $h(X)$ is ε -far from uniform, for $K = \Omega((\delta/4)^2 M/\varepsilon^2) = \Omega(\delta^2 MT/\log(\alpha/\delta))$. We set $\mathbf{X} = (X_1, \dots, X_T)$ to be T independent copies of X . Consider a hash function h such that $h(X)$ is ε -far from uniform. By Lemma 4.6, $(h(X_1), \dots, h(X_T))$ is $(1 - \eta)$ -far from uniform, for $\eta = e^{-\varepsilon^2 T/2} = \delta^3/128\alpha$. By Lemma 4.8, $(h(X_1), \dots, h(X_T))$ is $(\delta/4, 2\alpha/\delta)$ -nonuniform for $(1 - \delta/4)$ -fraction of hash functions h . By the first statement of Lemma 4.9 below, this implies that (H, \mathbf{Y}) is $(1 - \delta)$ -far from any distribution (H', \mathbf{Z}) with collision probability $\alpha/(|\mathcal{H}| \cdot M^T)$. ■

Lemma 4.9 *Let (H, Y) be a joint distribution over $\mathcal{H} \times [M]$ such that the marginal distribution H is uniform over \mathcal{H} . Let $\varepsilon, \delta, \alpha$ be positive real numbers.*

1. *If $Y|_{H=h}$ is $(\delta/4, 2\alpha/\delta)$ -nonuniform for at least $(1 - \delta/4)$ -fraction of $h \in \mathcal{H}$, then (H, Y) is $(1 - \delta)$ -far from any distribution (H', Z) with $\text{cp}(H', Z) \leq \alpha/(|\mathcal{H}| \cdot M)$.*
2. *If $Y|_{H=h}$ is $(0.1, 2\alpha/\varepsilon)$ -nonuniform for at least 2ε -fraction of $h \in \mathcal{H}$, then (H, Y) is ε -far from any distribution (H', Z) with $\text{cp}(H', Z) \leq \alpha/(|\mathcal{H}| \cdot M)$.*

Proof. We introduce the following notations first. For every $h \in \mathcal{H}$, we define $q_h : [M] \rightarrow \mathbb{R}$ by

$$q_h(y) = \min\{\Pr[(H, Y) = (h, y)], \Pr[(H', Z) = (h, y)]\}$$

for every $y \in [M]$. We also define $f : \mathcal{H} \rightarrow \mathbb{R}$ by

$$f(h) = \sum_{y \in [M]} q_h(y) \leq \Pr[H = h] = \frac{1}{|\mathcal{H}|}.$$

For the first statement, let (H', Z) be a random variable over $\mathcal{H} \times [M]$ that is $(1 - \delta)$ -close to (H, Y) . We need to show that $\text{cp}(H', Z) > \alpha/(|\mathcal{H}| \cdot M)$. Note that $\sum_h f(h) = 1 - \Delta((H, Y), (H', Z)) \geq \delta$. So there are at least a $(3\delta/4)$ -fraction of hash functions h with $f(h) \geq (\delta/4)/|\mathcal{H}|$. At least a $(3\delta/4) - (\delta/4) = \delta/2$ -fraction of h satisfy both $f(h) \geq (\delta/4)/|\mathcal{H}|$ and $Y|_{H=h}$ is $(\delta/4, 2\alpha/\delta)$ -nonuniform. By the definition of nonuniformity, for each such h , we have

$$\sum_{y \in [M]^T} (|\mathcal{H}| \cdot q_h(y))^2 > \frac{2\alpha}{\delta \cdot M}.$$

Therefore,

$$\text{cp}(H', Z) \geq \sum_{h, y} q_h(y)^2 > \left(\frac{\delta}{2} \cdot |\mathcal{H}|\right) \cdot \frac{2\alpha}{\delta \cdot |\mathcal{H}|^2 M} = \frac{\alpha}{|\mathcal{H}| \cdot M}.$$

Similarly, for the second statement, let (H', Z) be a random variable over $\mathcal{H} \times [M]$ that is ε -close to (H, Y) . We need to show that $\text{cp}(H', Z) > \alpha/(|\mathcal{H}| \cdot M)$. Note that $\sum_h f(h) = 1 - \Delta((H, Y), (H', Z)) \geq 1 - \varepsilon$. So there are at least a $1 - \varepsilon/0.9$ -fraction of h with $f(h) \geq 0.1/|\mathcal{H}|$. At least a $2\varepsilon - \varepsilon/0.9 > \varepsilon/2$ -fraction of hash functions satisfy both $f(h) \geq 0.1/|\mathcal{H}|$ and $Y|_{H=h}$ is $(0.1, 2\alpha/\varepsilon)$ -nonuniform. By Lemma 4.8, for each such h , we have

$$\sum_{y \in [M]} (|\mathcal{H}| \cdot q_h(y))^2 > \frac{2\alpha}{\varepsilon \cdot M}.$$

Therefore,

$$\text{cp}(H', Z) \geq \sum_{h, y} q_h(y)^2 > \left(\frac{\varepsilon}{2} \cdot |\mathcal{H}|\right) \cdot \frac{2\alpha}{\varepsilon \cdot |\mathcal{H}|^2 M} = \frac{\alpha}{|\mathcal{H}| \cdot M}.$$

■

4.3 Lower Bounds for the Distribution of Hashed Values Only

We can extend our lower bounds to the distribution of hashed sequence $\mathbf{Y} = (H(X_1), \dots, H(X_T))$ along (without H) for both closeness requirements, at the price of losing the dependency on ε and incurring some dependency on the size of the hash family. Let $2^d = |\mathcal{H}|$ be the size of the hash family. The dependency on d is necessary. Intuitively, the hashed sequence \mathbf{Y} contains at most $T \cdot m$ bits of entropy, and the input (H, X_1, \dots, X_T) contains at least $d + T \cdot k$ bits of entropy. When d is large enough, it is possible that all the randomness of hashed sequence comes from the randomness of the hash family. Indeed, if H is T -wise independent (which is possible with $d \simeq T \cdot m$), then $(H(X_1), \dots, H(X_T))$ is uniform when X_1, \dots, X_T are all distinct. Therefore,

$$\Delta((H(X_1), \dots, H(X_T)), U_{[M]^T}) \leq \Pr[\text{not all } X_1, \dots, X_T \text{ are distinct}]$$

Thus, $K = \Omega(T^2)$ (independent of M) suffices to make the hashed value close to uniform.

Theorem 4.10 *Let N, M, T be positive integers, and d a positive real number such that $N \geq MT/d$. Let $\delta \in (0, 1)$, $\alpha > 1$ be real numbers such that $\alpha \cdot 2^d < \delta^3 \cdot e^{T/32}/128$. Let $H : [N] \rightarrow [M]$ be a random hash function from an hash family \mathcal{H} of size at most 2^d . There exists an integer $K = \Omega(\delta^2 MT/d \cdot \log(\alpha/\delta))$, and a block K -source $\mathbf{X} = (X_1, \dots, X_T)$ such that $\mathbf{Y} = (H(X_1), \dots, H(X_T))$ is $(1 - \delta)$ -far from any distribution $\mathbf{Z} = (Z_1, \dots, Z_T)$ with $\text{cp}(\mathbf{Z}) \leq \alpha/M^T$. In particular, \mathbf{Y} is $(1 - \delta)$ -far from uniform.*

Think of α and δ as constants. Then the theorem says that when the hash function contains $d \leq T/(32 \ln 2) - O(1)$ bits of randomness, $K = \Omega(MT/d)$ is necessary for the hashed sequence to be close to uniform. For example, in some typical hash applications, $N = \text{poly}(M)$ and the hash function is 2-universal or $O(1)$ -wise independent. In this case, $d = O(\log M)$ and we need $K = \Omega(MT/\log M)$. (Recall that our upper bound in Theorem 3.1 says that $K = O(MT)$ suffices.)

Proof. We will deduce the theorem from Theorem 4.5. Replacing the parameter α by $\alpha \cdot 2^d$ in Theorem 4.5, we know that there exists an integer $K = \Omega(\delta^2 MT/d \cdot \log(\alpha/\delta))$ and a block K -source $\mathbf{X} = (X_1, \dots, X_T)$ such that $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$ is $(1 - \delta)$ -far from any distribution (H', \mathbf{Z}) with $\text{cp}(H', \mathbf{Z}) \leq \alpha \cdot 2^d / (2^d \cdot M^T) = \alpha/M^T$. Now, suppose we are given a random variable \mathbf{Z} on $[M]^T$ with $\Delta(\mathbf{Y}, \mathbf{Z}) \leq 1 - \delta$. Then we can define an H' such that $\Delta((H, \mathbf{Y}), (H', \mathbf{Z})) = \Delta(\mathbf{Y}, \mathbf{Z})$ (Indeed, define the conditional distribution $H'|_{\mathbf{Z}=\mathbf{z}}$ to equal $H|_{\mathbf{Y}=\mathbf{z}}$ for every $\mathbf{z} \in [M]^T$.) Then we have

$$\text{cp}(\mathbf{Z}) \geq \text{cp}(H', \mathbf{Z}) > \frac{\alpha}{M^T}. \quad \blacksquare$$

One limitation of the above lower bound is that it only works when $d \leq T/(32 \ln 2) - O(1)$. For example, the lower bound cannot be applied when the hash function is T -wise independent.

Although $d = \Omega(T)$ may not be interesting in practice, for the sake of completeness, we provide another simple lower bound to cover this parameter region.

Theorem 4.11 *Let N, M, T be positive integers, and $\delta \in (0, 1)$, $\alpha > 1$, $d > 0$ real numbers. Let $H : [N] \rightarrow [M]$ be a random hash function from an hash family \mathcal{H} of size at most 2^d . Suppose $K \leq N$ be an integer such that $K \leq (\delta^2/4\alpha \cdot 2^d)^{1/T} \cdot M$. Then there exists a block K -source $\mathbf{X} = (X_1, \dots, X_T)$ such that $\mathbf{Y} = (H(X_1), \dots, H(X_T))$ is $(1 - \delta)$ -far from any distribution $\mathbf{Z} = (Z_1, \dots, Z_T)$ with $\text{cp}(\mathbf{Z}) \leq \alpha/M^T$. In particular, \mathbf{Y} is $(1 - \delta)$ -far from uniform.*

Again, think of α and δ as constants. The theorem says that $K = \Omega(M/2^{d/T})$ is necessary for the hashed sequence to be close to uniform. In particular, when $d = \Theta(T)$, $K = \Omega(M)$ is necessary. Theorem 4.10 gives the same conclusion, but only works for $d \leq T/(32 \ln 2) - O(1)$. On the other hand, when $d = o(T)$, Theorem 4.10 gives better lower bound $K = \Omega(MT/d)$.

Proof. Let X be any flat K -source, i.e., a uniform distribution over a set of size K . We simply take $\mathbf{X} = (X_1, \dots, X_T)$ to be T independent copies of X . Note that \mathbf{Y} has support at most as large as (H, \mathbf{X}) . Thus,

$$|\text{supp}(\mathbf{Y})| \leq |\text{supp}(H, \mathbf{X})| = 2^d \cdot K^T \leq \frac{\delta^2}{4\alpha} \cdot M^T.$$

Therefore, \mathbf{Y} is $(1 - \delta^2/4\alpha)$ -far from uniform. By Lemma 4.8, \mathbf{Y} is $(1 - \delta)$ -far from any distribution $\mathbf{Z} = (Z_1, \dots, Z_T)$ with $\text{cp}(\mathbf{Z}) \leq \alpha/M^T$. \blacksquare

4.4 Lower Bound for 2-universal Hash Functions

In this subsection, we show Theorem 3.1 is almost tight in the following sense. We show that there exists $K = \Omega(MT/\varepsilon \cdot \log(1/\varepsilon))$, a 2-universal hash family \mathcal{H} , and a block K -source \mathbf{X} such that (H, \mathbf{Y}) is ε -far from having collision probability $100/(|\mathcal{H}| \cdot M^T)$. The improvement over Theorem 4.5 is the almost tight dependency on ε . Recall that Theorem 3.1 says that for 2-universal hash family, $K = O(MT/\varepsilon)$ suffices. The upper and lower bound differs by a factor of $\log(1/\varepsilon)$. In particular, our result for 4-wise independent hash functions (Theorem 3.5) cannot be achieved with 2-universal hash functions. The lower bound can further be extended to the distribution of hashed sequence $\mathbf{Y} = (H(X_1), \dots, H(X_T))$ as in the previous subsection. Furthermore, since the 2-universal hash family we use has small size, we only pay a factor of $O(\log M)$ in the lower bound on K . Formally we prove the following results.

Theorem 4.12 *For every prime power M , real numbers $\varepsilon \in (0, 1/4)$ and $\alpha \geq 1$, the following holds. For all integers t and N such that $\varepsilon \cdot M^{t-1} \geq 1$ and $N \geq 6\varepsilon M^{2t}$, and for $T = \lceil \varepsilon^2 M^{2t-1} \log(\alpha/\varepsilon) \rceil$,⁴ there exists an integer $K = \Omega(MT/\varepsilon \cdot \log(\alpha/\varepsilon))$, and a 2-universal*

⁴For technical reason, our lower bound proof does not work for every sufficiently large T . However, note that the density of T such that the lower bound holds is $1/M^2$.

hash family \mathcal{H} from $[N]$ to $[M]$, and a block K -source $\mathbf{X} = (X_1, \dots, X_T)$ such that $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$ is ε -far from any distribution (H', \mathbf{Z}) with $\text{cp}(H', \mathbf{Z}) \leq \alpha/(|\mathcal{H}| \cdot M^T)$.

Theorem 4.13 *For every prime power M , real numbers $\varepsilon \in (0, 1/4)$ and $\alpha \geq 1$, the following holds. For all integers t and N such that $\varepsilon \cdot M^{t-1} \geq 1$ and $N \geq 6\varepsilon M^{2t}$, and for $T = \lceil \varepsilon^2 M^{2t-1} \log(\alpha M/\varepsilon) \rceil$, there exists an integer $K = \Omega(MT/\varepsilon \cdot \log(\alpha M/\varepsilon))$, and a 2-universal hash family \mathcal{H} from $[N]$ to $[M]$, and a block K -source $\mathbf{X} = (X_1, \dots, X_T)$ such that $\mathbf{Y} = (H(X_1), \dots, H(X_T))$ is ε -far from any distribution \mathbf{Z} with $\text{cp}(\mathbf{Z}) \leq \alpha/M^T$.*

Basically, the idea is to show that the Markov Inequality applied in the proof of Theorem 3.1 (see Inequality (1)) is tight for a single block. More precisely, we show that there exists a 2-universal hash family \mathcal{H} , and a K -source X such that with probability ε over $h \leftarrow \mathcal{H}$, $\text{cp}(h(X)) \geq 1/M + \Omega(1/K\varepsilon)$. Intuitively, if we take $T = \Theta(K\varepsilon \cdot \log(\alpha/\varepsilon)/M)$ independent copies of such X , then the collision probability will satisfy $\text{cp}(h(X_1), \dots, h(X_T)) \geq (1 + \Omega(M/K\varepsilon))^T/M^T \geq \alpha/(\varepsilon M^T)$, and so the overall collision probability is $\text{cp}(H, \mathbf{Y}) \geq \alpha/(|\mathcal{H}| \cdot M^T)$. Formally, we analyze our construction below using Hellinger distance, and show that the collision probability remains high even after modifying a $\Theta(\varepsilon)$ -fraction of distribution.

Proof of Theorem 4.12: Fix a prime power M , and $\varepsilon > 0$, we identify $[M]$ with the finite field \mathbb{F} of size M . Let t be an integer parameter such that $M^{t-1} > 1/\varepsilon$. Recall that the set \mathcal{H}_0 of linear functions $\{h_{\vec{a}} : \mathbb{F}^t \rightarrow \mathbb{F}\}_{\vec{a} \in \mathbb{F}^t}$ where $h_{\vec{a}}(\vec{x}) = \sum_i a_i x_i$ is 2-universal. Note that picking a random hash function $h \leftarrow \mathcal{H}_0$ is equivalent to picking a random vector $\vec{a} \leftarrow \mathbb{F}^t$. Two special properties of \mathcal{H}_0 are (i) when $\vec{a} = \vec{0}$, the whole domain \mathbb{F}^t is sent to $0 \in \mathbb{F}$, and (ii) the size of hash family $|\mathcal{H}_0|$ is the same as the size of the domain, namely $|\mathbb{F}^t|$. We will use \mathcal{H}_0 as a building block in our construction.

We proceed to construct the hash family \mathcal{H} . We partition the domain $[N]$ into several sub-domains, and apply different hash function to each sub-domain. Let s be an integer parameter to be determined later. We require $N \geq s \cdot M^t$, and partition $[N]$ into D_0, D_1, \dots, D_s , where each of D_1, \dots, D_s has size M^t and is identified with \mathbb{F}^t , and D_0 is the remaining part of $[N]$. In our construction, the data \mathbf{X} will never come from D_0 . Thus, wlog, we can assume D_0 is empty. For every $i = 1, \dots, s$, we use a linear hash function $h_{\vec{a}_i} \in \mathcal{H}_0$ to send D_i to \mathbb{F} . Thus, a hash function $h \in \mathcal{H}$ consists of s linear hash function $(h_{\vec{a}_1}, \dots, h_{\vec{a}_s})$, and can be described by s vectors $\vec{a}_1, \dots, \vec{a}_s \in \mathbb{F}^t$. Note that to make \mathcal{H} 2-universal, it suffices to pick $\vec{a}_1, \dots, \vec{a}_s$ pairwise independently. Specifically, we identify \mathbb{F}^t with the finite field $\hat{\mathbb{F}}$ of size M^t , and pick $(\vec{a}_1, \dots, \vec{a}_s)$ by picking $a, b \in \hat{\mathbb{F}}$, and output $(a + \alpha_1 \cdot b, a + \alpha_2 \cdot b, \dots, a + \alpha_s \cdot b)$ for some distinct constants $\alpha_1, \dots, \alpha_s \in \hat{\mathbb{F}}$. Formally, we define the hash family to be

$$\mathcal{H} = \{h^{a,b} : [N] \rightarrow [M]\}_{a,b \in \hat{\mathbb{F}}}, \text{ where } h^{a,b} = (h_{a+\alpha_1 b}, \dots, h_{a+\alpha_s b}) \stackrel{\text{def}}{=} (h_1^{a,b}, \dots, h_s^{a,b}).$$

It is easy to verify that \mathcal{H} is indeed 2-universal, and $|\mathcal{H}| = M^{2t}$.

We next define a single block K -source X that makes the Markov Inequality (1) tight. We simply take X to be a uniform distribution over $D_1 \cup \dots \cup D_s$, and so $K = s \cdot M^t$. Consider a hash function $h^{a,b} \in \mathcal{H}$. If all $h_i^{a,b}$ are non-zero and distinct, then $h^{a,b}(X)$ is the uniform distribution. If exactly one $h_i^{a,b} = 0$, then $h^{a,b}$ sends $M^t + (s-1)M^{t-1}$ elements in $[N]$ to 0, and $(s-1)M^{t-1}$ elements to each nonzero $y \in \mathbb{F}$. Let us call such $h^{a,b}$ *bad* hash functions. Thus, if $h^{a,b}$ is bad, then

$$\begin{aligned} \text{cp}(h^{a,b}(X)) &= \left(\frac{M^t + (s-1)M^{t-1}}{K} \right)^2 + (M-1) \cdot \left(\frac{(s-1)M^{t-1}}{K} \right)^2 \\ &= \frac{1}{M} + \frac{M-1}{s^2 M} \geq \frac{1}{M} + \frac{1}{2s^2}. \end{aligned}$$

Note that $h^{a,b}$ is bad with probability

$$\Pr[\text{exactly one } h_i^{a,b} = 0] = \Pr[b \neq 0 \wedge \exists i \quad (a + \alpha_i b = 0)] = \left(1 - \frac{1}{M^t}\right) \cdot \frac{s}{M^t} \geq \frac{s}{2M^t}.$$

We set $s = \lceil 4\varepsilon M^t \rceil \leq M^t$. It follows that with probability at least 2ε over $h \leftarrow \mathcal{H}$, the collision probability satisfies $\text{cp}(h(X)) \geq 1/M + 1/(4K\varepsilon)$, as we intuitively desired. However, instead of working with collision probability directly, we need to use Hellinger closeness to measure the growth of distance to uniform (see Definition 3.9.) The following claim upper bounds the Hellinger closeness of $h(X)$ for bad hash functions h . The proof of the claim is deferred to the end of this section.

Claim 4.14 *Suppose h is a bad hash function defined as above, then the Hellinger closeness of $h(X)$ satisfies $C(h(X)) \leq 1 - M/(64K\varepsilon)$.*

Finally, for every integer $T \in [\varepsilon^2 M^{2t-1} \log(\alpha/\varepsilon), c_0 \cdot \varepsilon^2 M^{2t-1} \log(\alpha/\varepsilon)]$, we can write $T = c \cdot (64K\varepsilon/M) \cdot \ln(800\alpha/\varepsilon)$ for some constant $c < c_0$. Let $\mathbf{X} = (X_1, \dots, X_T)$ be T independent copies of X . We now show that $K, \mathcal{H}, \mathbf{X}$ satisfy the conclusion of the theorem. That is, $K = \Omega(MT/(\varepsilon \log(\alpha/\varepsilon)))$ (as follows from the definition of T) and $(H, \mathbf{Y}) = (H, H(X_1), \dots, H(X_T))$ is ε -far from any distribution (H', \mathbf{Z}) with $\text{cp}(H', \mathbf{Z}) \leq \alpha/(|\mathcal{H}| \cdot M^T)$.

Consider the distribution $(h(X_1), \dots, h(X_T))$ for a bad hash function $h \in \mathcal{H}$. From the above claim, the Hellinger closeness satisfies

$$C(h(X_1), \dots, h(X_T)) = C(h(X))^T \leq (1 - M/64K\varepsilon)^T \leq e^{MT/64K\varepsilon} \leq \frac{800\alpha}{\varepsilon}.$$

By Lemma 3.10 and the definition of Hellinger closeness, we have

$$\Delta((h(X_1), \dots, h(X_T)), U_{[M]^T}) \geq 1 - C(h(X_1), \dots, h(X_T)) \geq 1 - \frac{800\alpha}{\varepsilon}.$$

By Lemma 4.8, $(h(X_1), \dots, h(X_T))$ is $(0.1, 2\alpha/\varepsilon)$ -nonuniform for at least 2ε -fraction of bad hash functions h . By the second statement of Lemma 4.9, this implies (H, \mathbf{Y}) is ε -far from any distribution (H', \mathbf{Z}) with $\text{cp}(H', \mathbf{Z}) \leq \alpha/(|\mathcal{H}| \cdot M^T)$.

In sum, given $M, \varepsilon, \alpha, t$ that satisfies the premise of the theorem, we set $K = \lceil 4\varepsilon M^t \rceil \cdot M^t$, and proved that for every $N \geq K$, and $T = \Theta((K\varepsilon/M) \cdot \ln(\alpha/\varepsilon))$, the conclusion of the theorem is true. It remains to prove Claim 4.14.

Proof of claim: Let $p(x) = M \cdot \Pr[h(X) = x]$ for every $x \in \mathbb{F}$. For a bad hash function h , we have $p(0) = (1 + (M-1)/s)$, and $p(x) = (1 - 1/s)$ for every $x \neq 0$. We will upper bound $C(h(X)) = (1/M) \cdot \sum_x \sqrt{p(x)}$ using Taylor series. Recall that for $z \geq 0$, there exists some $z', z'' \in [0, z]$ such that

$$\begin{aligned} \sqrt{1+z} &= 1 + \frac{z}{2} + \frac{z^2}{2} \cdot \left(-\frac{1}{4(1+z')^{3/2}} \right) \leq 1 + \frac{z}{2} - \frac{z^2}{8(1+z)^{3/2}}, \text{ and} \\ \sqrt{1-z} &= 1 - z \cdot \frac{1}{2\sqrt{1-z''}} \leq 1 - \frac{z}{2}. \end{aligned}$$

We have

$$\begin{aligned} C(h(X)) &= \frac{1}{M} \sum_x \sqrt{p(x)} \\ &\leq \frac{1}{M} \left(1 + \frac{M-1}{2s} - \frac{(M-1)^2}{8s^2 \cdot (1 + (M-1)/s)^{3/2}} + (M-1) \cdot \left(1 - \frac{1}{2s} \right) \right) \\ &= 1 - \frac{(M-1)^2}{8Ms^2(1 + (M-1)/s)^{3/2}} \end{aligned}$$

Recall that $M \geq 2$, $s = \varepsilon M^t \geq M$, and $s^2 = K\varepsilon$, we have

$$C(h(X)) \leq 1 - \frac{M^2}{64K\varepsilon}.$$

□

■

Recall that $|\mathcal{H}| = M^{2t}$. Theorem 4.13 follows from Theorem 4.12 by exactly the same argument as in the proof of Theorem 4.10.

Acknowledgments

We thank Wei-Chun Kao for helpful discussions in the early stages of this work, David Zuckerman for telling us about Hellinger distance, and Michael Mitzenmacher for suggesting parameter settings useful in practice.

References

- [BBR85] Charles H. Bennett, Gilles Brassard, and Jean-Marc Robert. How to reduce your enemy's information (extended abstract). In Hugh C. Williams, editor, *Advances in Cryptology—CRYPTO '85*, volume 218 of *Lecture Notes in Computer Science*, pages 468–476. Springer-Verlag, 1986, 18–22 August 1985.
- [BM03] Andrei Z. Broder and Michael Mitzenmacher. Survey: Network applications of bloom filters: A survey. *Internet Mathematics*, 1(4), 2003.
- [CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.*, 17(2):230–261, 1988.
- [CV08] Kai-Min Chung and Salil Vadhan. Tight bounds for hashing block sources. In *APPROX-RANDOM*, 2008.
- [Dur04] Richard Durrett. *Probability: Theory and Examples. Third Edition*. Duxbury, 2004.
- [GS02] Alison L. Gibbs and Francis Edward Su. On choosing and bounding probability metrics. *International Statistical Review*, 70:419, 2002.
- [ILL89] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstracts). In *Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing*, pages 12–24, Seattle, Washington, 15–17 May 1989.
- [Knu98] Donald E. Knuth. *The art of computer programming, Volume 3: Sorting and Searching*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1998.
- [MV08] Michael Mitzenmacher and Salil Vadhan. Why simple hash functions work: Exploiting the entropy in a data stream. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '08)*, pages 746–755, 20–22 January 2008.
- [Mut03] S. Muthukrishnan. Data streams: algorithms and applications. In *SODA*, pages 413–413, 2003.
- [NT99] Noam Nisan and Amnon Ta-Shma. Extracting randomness: A survey and new constructions. *J. Comput. Syst. Sci.*, 58(1):148–173, 1999.
- [NZ96] Noam Nisan and David Zuckerman. Randomness is linear in space. *Journal of Computer and System Sciences*, 52(1):43–52, February 1996.
- [RT00] Jaikumar Radhakrishnan and Amnon Ta-Shma. Bounds for dispersers, extractors, and depth-two superconcentrators. *SIAM Journal on Discrete Mathematics*, 13(1):2–24 (electronic), 2000.

- [Rey04] Leonid Reyzin. A note on the statistical difference of small direct products. Technical Report BUCS-TR-2004-032, Boston University Computer Science Department, 2004.
- [SV99] Amit Sahai and Salil Vadhan. Manipulating statistical difference. In Panos Pardalos, Sanguthevar Rajasekaran, and José Rolim, editors, *Randomization Methods in Algorithm Design (DIMACS Workshop, December 1997)*, volume 43 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 251–270. American Mathematical Society, 1999.
- [Sha04] Ronen Shaltiel. Recent developments in extractors. In G. Paun, G. Rozenberg, and A. Salomaa, editors, *Current Trends in Theoretical Computer Science*, volume 1: Algorithms and Complexity. World Scientific, 2004.
- [Vad07] Salil Vadhan. The unified theory of pseudorandomness. *SIGACT News*, 38(3), September 2007.
- [Zuc96] David Zuckerman. Simulating BPP using a general weak random source. *Algorithmica*, 16(4/5):367–391, October/November 1996.

A Technical Lemma on Binomial Coefficients

Lemma A.1 (Claim 4.3, restated) *Let $N, K > 1$ be integers such that $N > 2K$, and $L \in [0, K/2]$, $\beta \in (0, \min\{1, \sqrt{L}\})$ real numbers. Let $S \subset [N]$ be a random subset of size K , and $T \subset [N]$ be a fixed subset of $[N]$ of arbitrary size. We have*

$$\Pr_S \left[\left| |S \cap T| - L \right| \leq \beta \sqrt{L} \right] \leq O(\beta).$$

Proof. By an abuse of notation, we use T to denote the size of set T . The probability can be expressed as a sum of binomial coefficients as follows.

$$\Pr_S \left[\left| |S \cap T| - L \right| \leq \beta \sqrt{L} \right] = \sum_{R=\lceil L-\beta\sqrt{L} \rceil}^{\lfloor L+\beta\sqrt{L} \rfloor} \frac{\binom{T}{R} \binom{N-T}{K-R}}{\binom{N}{K}}.$$

Note that there are at most $\lfloor 2\beta\sqrt{L} \rfloor + 1$ terms, it suffices to show that for every $R \in [L - \beta\sqrt{L}, L + \beta\sqrt{L}]$,

$$f(T) \stackrel{\text{def}}{=} \frac{\binom{T}{R} \binom{N-T}{K-R}}{\binom{N}{K}} \leq O\left(\sqrt{\frac{1}{L}}\right).$$

We use the following bound on binomial coefficients, which can be derived from Stirling’s formula.

Claim A.2 For integers $0 < i < a$, $0 < j < b$, we have

$$\frac{\binom{a}{i} \binom{b}{j}}{\binom{a+b}{i+j}} \leq O \left(\sqrt{\frac{a \cdot b \cdot (i+j) \cdot (a+b-i-j)}{i \cdot (a-i) \cdot j \cdot (b-j) \cdot (a+b)}} \right).$$

Note that $L \in [0, K/2]$ implies $K - R = \Omega(K)$. When $2R \leq T \leq N - 2K + 2R$, we have

$$\begin{aligned} f(T) &= \frac{\binom{T}{R} \binom{N-T}{K-R}}{\binom{N}{K}} \\ &= O \left(\sqrt{\frac{T(N-T)K(N-K)}{R(T-R)(K-R)(N-T-K+R)N}} \right) \\ &= O \left(\sqrt{\frac{1}{R} \cdot \frac{K}{K-R} \cdot \frac{N-K}{N} \cdot \frac{T(N-T)}{(T-R)(N-T-K+R)}} \right) \\ &= O \left(\sqrt{\frac{1}{R}} \right) = O \left(\sqrt{\frac{1}{L}} \right), \end{aligned}$$

as desired. Note that when $N > 2K$, such T exists. Finally, observe that $\beta^2 < L$ implies $R \geq 1$, and

$$\frac{f(T)}{f(T+1)} = \frac{(T-R+1)(N-T)}{(T+1)(N-T-K+R)}.$$

It follows that $f(T)$ is increasing when $T \leq 2R$, and $f(T)$ is decreasing when $T \geq N - 2K + 2R$. Therefore, $f(T) \leq f(2R) = O(\sqrt{1/L})$ for $T \leq 2R$, and $f(T) \leq f(N - 2K + 2R) = O(\sqrt{1/L})$ for $T \geq N - 2K + 2R$, which complete the proof. \blacksquare

B Proof of Lemma 4.4

Lemma B.1 (Lemma 4.4, restated) Let X and Y be random variables over $[M]$ such that $\Delta(X, Y) \geq \varepsilon$. Let $\mathbf{X} = (X_1, \dots, X_T)$ be T i.i.d. copies of X , and let $\mathbf{Y} = (Y_1, \dots, Y_T)$ be T i.i.d. copies of Y . We have

$$\Delta(\mathbf{X}, \mathbf{Y}) \geq \min\{\varepsilon_0, c\sqrt{T} \cdot \varepsilon\},$$

where ε_0, c are absolute constants.

Proof. We prove the lemma by the following two claims. The first claim reduces the lemma to the special case that X is a Bernoulli random variable with bias $\Omega(\varepsilon)$, and Y is a uniform coin. The second claim proves the special case.

Claim B.2 Let X and Y be random variables over $[M]$ such that $\Delta(X, Y) = \varepsilon$. Then there exists a randomized function $f : [M] \rightarrow \{0, 1\}$ such that $f(Y) = U_{\{0,1\}}$, and $\Delta(f(X), f(Y)) \geq \varepsilon/2$.

Proof of claim: By the definition, there exists a set $T \subset [M]$ such that

$$|\Pr[X \in T] - \Pr[Y \in T]| = \varepsilon.$$

With out loss of generality, we can assume that $\Pr[Y \in T] = p \leq 1/2$ (because we can take the complement of T .) Let $g : [M] \rightarrow \{0, 1\}$ be the indicator function of T , so we have $\Pr_Y[g(Y) = 1] = p$. For every $x \in [M]$, we define $f(x) = g(x) \vee b$, where b is a biased coin with $\Pr[b = 0] = 1/(2(1-p))$. The claim follows by observing that

$$\Pr[f(Y) = 0] = \Pr[g(Y) = 0 \wedge b = 0] = (1-p) \cdot 1/(2(1-p)) = 1/2,$$

and

$$\Delta(f(X), f(Y)) \geq \Delta(X, Y) \cdot \Pr[b = 0] \geq \varepsilon/2.$$

□

Claim B.3 Let X be a Bernoulli random variable over $\{0, 1\}$ such that $\Pr[X = 0] = 1/2 - \varepsilon$. Let $\mathbf{X} = (X_1, \dots, X_T)$ be T independent copies of X . Then

$$\Delta(\mathbf{X}, U_{\{0,1\}^T}) \geq \min\{\varepsilon_0, c\sqrt{T}\varepsilon\},$$

where ε_0, c are absolute constants independent of ε and T .

Proof of claim: For $\mathbf{x} \in \{0, 1\}^T$, let the *weight* $\text{wt}(\mathbf{x})$ of \mathbf{x} to be the number of 1's in \mathbf{x} . Let

$$S = \left\{ x \in \{0, 1\}^T : \text{wt}(x) \leq \frac{T}{2} - \sqrt{T} \right\}$$

be the subset of $\{0, 1\}^T$ with small weight (This choice of S is the main source of improvement in our proof compared to that of Reyzin [Rey04], who instead considers the set of all x with weight at most $T/2$.) For every $\mathbf{x} \in S$, we have

$$\Pr[\mathbf{X} = \mathbf{x}] \leq \frac{1}{2^T} \cdot (1-\varepsilon)^{T/2+\sqrt{T}} \cdot (1+\varepsilon)^{T/2-\sqrt{T}} \leq \left(1 - \min \left\{ \frac{\sqrt{T} \cdot \varepsilon}{2}, \frac{1}{2} \right\} \right) \cdot \Pr[U_{\{0,1\}^T} = \mathbf{x}].$$

By standard results on large deviation, we have

$$\Pr[U_{\{0,1\}^T} \in S] \geq \Omega(1).$$

Combining the above two inequalities, we get

$$\begin{aligned}
\Delta(\mathbf{X}, U_{\{0,1\}^T}) &\geq \Pr[U_{\{0,1\}^T} \in S] - \Pr[\mathbf{X} \in S] \\
&\geq \left(1 - \left(1 - \min\left\{\frac{\sqrt{T} \cdot \varepsilon}{2}, \frac{1}{2}\right\}\right)\right) \cdot \Pr[U_{\{0,1\}^T} \in S] \\
&\geq \min\left\{\frac{\sqrt{T} \cdot \varepsilon}{2}, \frac{1}{2}\right\} \cdot \Omega(1) \\
&= \min\{c\sqrt{T}\varepsilon, \varepsilon_0\}
\end{aligned}$$

for some absolute constants c, ε_0 , which completes the proof. \square

Note that applying the same randomized function f on two random variables X and Y cannot increase the statistical distance. I.e., $\Delta(f(X), f(Y)) \leq \Delta(X, Y)$. The lemma following immediately by the above two claims:

$$\begin{aligned}
\Delta(\mathbf{X}, \mathbf{Y}) &\geq \Delta(((f_1(X_1), \dots, f_T(X_T)), ((f_1(Y_1), \dots, f_T(Y_T)) \\
&\geq \min\{\varepsilon_0, c\sqrt{T}\varepsilon\}
\end{aligned}$$

where f_1, \dots, f_T are independent copies of randomized function defined in Claim B.2, and ε_0, c are absolute constants from Claim B.3. \blacksquare