

Sample(x) = (a*x<=t) is a distinguisher with probability 1/8

Mikkel Thorup
 University of Copenhagen
 Copenhagen, Denmark.
 Email: mikkel2thorup@gmail.com

Abstract

A random sampling function $Sample : U \rightarrow \{0, 1\}$ for a key universe U is a *distinguisher with probability α* if for any given assignment of values $v(x)$ to the keys $x \in U$, including at least one non-zero $v(x) \neq 0$, the sampled sum $\sum\{v(x) \mid x \in U \wedge Sample(x) = 1\}$ is non-zero with probability at least α . Here the key values may come from any commutative monoid (addition is commutative and associative and zero is neutral). Such distinguishers were introduced by Vazirani [PhD thesis 1986], and Naor and Naor used them for their small bias probability spaces [STOC'90]. Constant probability distinguishers are used for testing in contexts where the key values are not computed directly, yet where the sum is easily computed. A simple example is when we get a stream of key value pairs $(x_1, v_1), (x_2, v_2), \dots, (x_n, v_n)$ where the same key may appear many times. The accumulated value of key x is $v(x) = \sum\{v_i \mid x_i = x\}$. For space reasons, we may not be able to maintain $v(x)$ for every key x , but the sampled sum is easily maintained as the single value $\sum\{v_i \mid Sample(x_i) = 1\}$.

Here we show that when dealing with w -bit integers, if a is a uniform odd w -bit integer and t is a uniform w -bit integer, then $Sample(x) = [ax \bmod 2^w \leq t]$ is a distinguisher with probability $1/8$. Working with standard units, that is, $w = 8, 16, 32, 64$, we exploit that w -bit multiplication works modulo 2^w , discarding overflow automatically, and then the sampling decision is implemented by the C-code `a*x<=t`. Previous such samplers were much less computer friendly, e.g., the distinguisher of Naor and Naor [STOC'90] was more complicated and involved a 7-independent hash function.

Keywords

Hashing; Distinguishers; Small Bias Sampling Spaces;

I. INTRODUCTION

A random sampling function $Sample : U \rightarrow [2] = \{0, 1\}$ for a set U of keys is viewed as sampling the keys $x \in U$ with $Sample(x) = 1$. It is a *distinguisher with probability α* , or, for short, an α -*distinguisher*, if for any commutative monoid R (addition is commutative and associative and zero is neutral) and any given assignment of values $v(x) \in R$ to the keys $x \in U$, including at least one non-zero $v(x) \neq 0$, the *sampled sum*

$$\sum\{v(x) \mid x \in U \wedge Sample(x) = 1\}$$

is non-zero with probability at least α . Such distinguishers were introduced by Vazirani [1], and Naor and Naor used them for their small bias probability spaces [2]. As an example, with $R = \mathbb{F}_2$, the goal is to sample an odd number of elements from the set X of elements $x \in U$ with $v(x) = 1$. This is very different from the more typical goal in sampling (see, e.g., [3]) where we with some small sampling probability p hope to sample roughly $p|X|$ elements from X .

To simplify notation, we often identify the sampling function $Sample$ with the set $Sample^{-1}(1)$ of keys it samples from U , that is, $x \in Sample \iff Sample(x) = 1$. Thinking of $Sample$ as a random subset of U is often simpler, but our goal is to implement $Sample$ as a membership function, returning 1, or “true”, for the elements in the subset. Also, for any key set $X \subseteq U$, we define the multiset $v(X) = \{v(x) \mid x \in X\}$. Then the sampled sum is written as $\sum v(Sample)$. Finally, we define a *non-zero value function* $v : U \rightarrow R$ as one that has a non-zero value $v(x)$ for at least one $x \in U$.

It is easy to see that a fully random sampling function $Sample : U \rightarrow [2]$ is a distinguisher with probability $1/2$; for consider any non-zero value function $v : U \rightarrow R$. We can make the independent sampling decisions $Sample(x)$, $x \in U$, one by one. When we get to the last non-zero valued x in U , if the sum over the previous samples is non-zero, the final sampled sum is non-zero if x is not sampled, which happens with probability $1/2$. On the other hand, if the sum over the previous samples is zero, the final sampled sum is non-zero if x is included, which happens with probability $1/2$. Thus, if the sample of x is independent of the previous samples, then we do get a distinguisher with probability $1/2$. The best possible distinguisher samples a uniformly random non-empty set, that is, we re-run the fully random sampling if no samples are made. It is a distinguisher with probability $1/2 + 1/2^u$ where $u = |U|$. Fully random sampling functions are, however, not realistic for large sets of keys.

As a framework for more realistic hash functions, Wegman and Carter [4] define a random hash function $h : U \rightarrow [m] = \{0, \dots, m - 1\}$ to be k -independent if it is truly random when restricted to any k distinct keys x_1, \dots, x_k , that is, for any $t_1, \dots, t_k \in [m]$, $\Pr[\bigwedge_{i=1}^k h(x_i) = t_i] = 1/m^k$. Unfortunately, k -independent sampling does not in itself suffice for distinguishers unless we have full independence with $k = u$. To see this, consider the constant value function $v : U \rightarrow \mathbb{F}_2$ where $v(x) = 1$ for all $x \in U$. The sampled sum over \mathbb{F}_2 is zero exactly if we sample an even number of keys. Making a uniformly random such sample is $(u - 1)$ -independent; for it can be described by taking a fully random sample from any $u - 1$ given keys, and then make the sampling decision for the last key so that the total number of samples is even. This $(u - 1)$ -independent sampling is a distinguisher with probability 0.

Thus k -independent sampling functions cannot be used directly as distinguishers, but the concept is still useful in the design of distinguishers. The previous constant probability distinguisher defined by [2, §3.1.1, pp. 843–846] uses, as a subroutine, a 7-independent hash function, and it is a distinguisher with probability $1/8$.

In this paper we present some very simple and efficient constant probability distinguishers. First we have a generic construction based on 2-independent hashing. Let $h : U \rightarrow [m]$ be a 2-independent hash function. Moreover let the “threshold” t be uniformly distributed in $[m]$. Define $Sample(x) = [h(x) \leq t]$. Assuming that the value function $v : U \rightarrow R$ has $n > 0$ non-zeros, we will show that $\Pr[\sum v(Sample) \neq 0] \geq (1 - n^2/m^2)/8$, so for $n \ll m$, $Sample$ is essentially a $1/8$ -probability distinguisher.

For a canonical example of a 2-independent hash function h , we pick a prime p such that $U \subseteq \mathbb{Z}_p$, set $m = p$, and choose a and b uniformly in \mathbb{Z}_p . Then $h(x) = (ax + b) \bmod p$ is 2-independent, so $Sample(x) = [(ax + b \bmod p) \leq t]$ is a distinguisher with probability $(1 - n^2/p^2)/8$. Not surprisingly, the addition of b is superfluous. In fact, if a is uniform in $[m]_+ = \{1, \dots, m - 1\}$, we will prove that $Sample(x) = [ax \bmod p \leq t]$ is a strict $1/8$ -probability distinguisher. More efficient 2-independent schemes exist, see, e.g., [5], but we will show something even more efficient inspired by the extremely fast universal hashing scheme of Dietzfelbinger et al. [6]. The new scheme works for w -bit integers, that is, with key set $U = [2^w]$. Let a be a uniform odd w -bit integer and let t be a uniform w -bit integer. We will show that $Sample(x) = [ax \bmod 2^w \leq t]$ is a distinguisher with probability $1/8$. The analysis for this result is, technically, the hardest part of this paper. Working with standard units, that is, $w = 8, 16, 32, 64$, we can exploit that w -bit multiplication works modulo 2^w , discarding overflow automatically, and then $Sample(x)$ is implemented by the C-code `a*x<=t`. We could also use `a*x<t`, but this yields a less clean analysis and slightly weaker bounds.

A. Applications

We will now discuss some applications of distinguishers. These applications are all known, and could be based on the distinguisher from [2]. Our contribution is a very efficient distinguisher that is easy to implement, thus making distinguishers more attractive to use in applications.

Constant probability distinguishers are used for testing in contexts where the key values are not computed directly, yet where the sum is easily computed. A simple example is when we get a stream of key value pairs $(x_1, v_1), (x_2, v_2), \dots, (x_n, v_n)$ where the same key may appear many times. The accumulated value of key x is

$v(x) = \sum\{v_i \mid x_i = x\}$. For space reasons, we may not be able to maintain $v(x)$ for every key x , but the sampled sum is easily maintained as the single value $\sum\{v_i \mid \text{Sample}(x_i) = 1\}$.

Another example suggested by [2] is to implement Freivald's technique [7] to verify a matrix product $AB = C$ for $n \times n$ matrices. We use a random sampling function $\text{Sample} : [n] \rightarrow [2]$, which we view as a 0/1-vector $\vec{s} = (\text{Sample}(0), \dots, \text{Sample}(n-1))^\top$. We now just test if $A(B\vec{s}) = C\vec{s}$. This is trivially done in $O(n^2)$ time, and if Sample is an α -distinguisher, then we discover $C \neq AB$ with probability α . In Freivald's original algorithm [7], the vector \vec{r} is fully random, but we generate it here using only $O(\log n)$ bits. We note that for this concrete application, we have good specialized techniques [8], but they do not allow us to compute $\text{Sample}(x)$ directly for a given key $x \in [n]$ in constant time.

In the streaming context above, testing is also a natural application: if someone claims to have the correct value function $v' : U \rightarrow R$, then we can test if $v' - v$ is non-zero by comparing the sampled sums.

A quite different application from [9], inspired by [10], [11], works with values from $R = \mathbb{F}_2$. We have a tree T in a graph $G = (V, E)$ and we want to test if some edge from E is leaving T . The value of an edge is the number of end-points it has in T , which in \mathbb{F}_2 is zero if the edge has both end-points in T , so we have a non-zero value exactly when there is an edge leaving T . Accumulating the sampled sum to test for a non-zero is comparatively easy in many contexts.

Following [2], to decrease the probability of missing a non-zero from α to ε , we can apply $d = \lceil \log_\alpha \varepsilon \rceil$ independent α -distinguishers $\text{Sample}_1, \dots, \text{Sample}_d : U \rightarrow [2]$, and ask if the vector of sampled sums $(\sum v(\text{Sample}_1), \dots, \sum v(\text{Sample}_d))$ is all zeros. In fact, these samplers do not really need to be independent. As described in [2, §3.2], based on a random walk in an expander, we can generate $d = O(\log(1/\varepsilon))$ random samplers using $O(\log u + \log(1/\varepsilon))$ random bits, so that the probability that all sampled sums are zero is ε . We shall generally refer to this as a *vector distinguisher*.

In [2] the main focus is the case where the values are from \mathbb{F}_2 . The goal is to create an ε -biased sampling space¹ defined by a sampling function $\text{SAMPLE} : U \rightarrow [2]$ that for any non-empty subset $I \subseteq U$, samples an odd number of elements from I with probability $(1 \pm \varepsilon)/2$. Translating to our framework, let the value function $v : U \rightarrow \mathbb{F}_2$ be the characteristic function for I , that is, $v(x) = 1 \iff x \in I$. Then we want $\Pr[\sum v(\text{SAMPLE}) = 1] = (1 \pm \varepsilon)/2$. To define such a function SAMPLE , we can use the above vector distinguisher with $d = \lceil \log_\alpha \varepsilon \rceil$ regular α -probability distinguishers $\text{Sample}_1, \dots, \text{Sample}_d : U \rightarrow [2]$, and take the inner product with a fully random bit vector $\vec{b} = (b_1, \dots, b_d) \in [2]^d$. Thus we define

$$\text{SAMPLE}(x) = \left\langle \vec{b}, (\text{Sample}_1(x), \dots, \text{Sample}_d(x)) \right\rangle = \sum_{i=1}^d (b_i \cdot \text{Sample}_i(x)). \quad (1)$$

Since $\Pr[(\sum v(\text{Sample}_1), \dots, \sum v(\text{Sample}_d)) = 0^d] \leq \varepsilon$, we get

$$(1 - \varepsilon)/2 \leq \Pr \left[\sum v(\text{SAMPLE}) = 1 \right] \leq 1/2.$$

Like the vector distinguisher, this construction can be implemented using $O(\log u + \log(1/\varepsilon))$ random bits.

Naor and Naor's paper [2] on small bias sampling spaces have been very influential with more than 500 citations, and we cannot cover all applications. The point in the current paper is to show that constant probability distinguishers, which often play a central role, now have a truly practical implementation. We are not claiming any new asymptotic bounds in terms of O -notation.

B. Comparison with previous implementations

When it comes to the implementation of constant probability distinguishers, we already mentioned the original construction of Naor and Naor [2, §3.1.1, pp. 843–846] involving, among other components, a 7-independent hash function, and coincidentally, it is a 1/8-probability distinguisher like our $\text{a} * \text{x} < = \text{t}$.

¹We follow the definition from http://en.wikipedia.org/wiki/Small-bias_sample_space, which is different but mathematically equivalent to the one in [2].

Small bias sampling spaces: Alon et al. [12] have presented some simple constructions of ε -biased sampling spaces over \mathbb{F}_2 . The goal is to make sampling decisions r_i for $i \in [n]$. Quoting [13]: *Our three constructions are so simple that they can be described in the three corresponding paragraphs below:*

- 1) A point in the first sample space is specified by two bit strings of length $m \stackrel{\text{def}}{=} \log(n/\varepsilon)$ each, denoted $f_0 \cdots f_{m-1}$ and $s_0 \cdots s_{m-1}$, where $f_0 = 1$ and $f(t) = t^m + \sum_{i=0}^{m-1} f_i \cdot t^i$ is an irreducible polynomial. The n -bit sample string, denoted $r_0 \cdots r_{n-1}$, is determined by $r_i = s_i$ for $i \leq m$ and $r_i = \sum_{j=0}^{m-1} f_j \cdot r_{i-m+j}$ for $i \geq m$.
- 2) A point in the second sample space is specified by residue x modulo a fixed prime $p \geq (n/\varepsilon)^2$. The n -bit sample string, denoted $r_0 \cdots r_{n-1}$, is determined by $r_i = 0$ if $x + i$ is a quadratic residue modulo p and $r_i = 1$ otherwise.
- 3) A point in the third sample space is specified by two bit strings of length $m \stackrel{\text{def}}{=} \log(n/\varepsilon)$ each, denoted x and y . The n -bit sample string, denoted $r_0 \cdots r_{n-1}$, is determined by letting r_i equal the inner product mod-2 of the binary vectors x^i and y where x^i is the i^{th} power of x when considered as an element of $GF(2^m)$.

Mathematically, these descriptions are all very simple, and could all be made to work if we want to generate all n bits r_i . The setting of this paper, however, is that of hashing, where we want to compute r_i directly for any given key $i \in [n]$. Moreover, we typically think of n as much larger than the available space, e.g., $n = u = 2^{64}$ for 64-bit keys. In this hashing context, Construction 1 has problems because r_i is determined from r_{i-m}, \dots, r_{i-1} . Construction 2 and 3 both appear to require $\Omega(\log u)$ time for computing r_i directly, e.g., using Euler's criterion for the quadratic residues. Thus, with $u \gg 1/\varepsilon$, it is much better to use the function *SAMPLE* from (1) which makes $O(\log(1/\varepsilon))$ constant time calls to a constant probability distinguisher, e.g., either the one from [2], or, for even better constants, the new `a*x<=t` distinguisher proposed in here.

We note that there has been interesting recent work [14] on reducing the number of random bits needed to generate ε -biased sampling spaces. While the number of random bits is important in some contexts, it is not the concern in this paper.

AMS sketches: In the case where the values are integers or reals, we can also get a distinguisher from the AMS sketches [15]. These sketches are used to estimate $\sum_{x \in U} (v(x))^2$, which is non-zero if and only if v has a non-zero. Based on the analysis in [15], it is quite easy to see that if $v : U \rightarrow \mathbb{R}$ has a non-zero, then the AMS sketch returns non-zero with probability $1/3$.

The AMS sketch is based on 4-independent hashing, and using the same analysis as in [15], we can easily show that any 4-independent sampling function *Sample* : $U \rightarrow [2]$ is a distinguisher with probability $1/3$ when the values are reals. This is a nice contrast to the situation with for \mathbb{F}_2 where even $(|U| - 1)$ -independence does not suffice on its own.

We shall return to the AMS sketch and the 4-independent sampling of reals in Section V.

C. Theory and Practice

In theory of computing we often calculate running times with O -notation, ignoring constant factors since we try to identify the dominant term as the problem size n goes to infinity. Nevertheless, constants do matter, and they can be appealing to study when we know that we are getting close to the right answer, e.g., for the number of comparisons needed for sorting.

In this paper, we show that a constant probability distinguisher can be implemented as `a*x<=t` using just two standard instructions on integers. The constant two is clearly optimal as no single standard instruction will do on its own.

As with many theoretical measures, we do not expect such an instruction count to fully capture the practical performance on real computers, because for example the relative performance of different instructions vary from computer to computer. Nevertheless, with such a small instruction count, we would expect it to be very competitive in practice.

code/scheme	time (ns)
<code>a*x<=t</code>	1.02
<code>a*x>>63</code>	0.76
<code>if (a*x<=t) S+=x;</code>	1.15
<code>if (a*x>>63) S+=x;</code>	0.97
7-independent hashing	52.12

Table I

AVERAGE TIME PER COMPUTATION OVER 10 MILLION 64-BIT KEYS x ON A MACBOOK AIR WITH A 1.7 GHz INTEL CORE I5 PROCESSOR. THE TABLE COMPARES OUR NEW `a*x<=t` DISTINGUISHER WITH THE UNIVERSAL HASHING SCHEME `a*x>>63` WHICH IS NOT A DISTINGUISHER BUT A GENERAL BENCHMARK FOR FAST HASHING. THEY ARE TIMED BOTH ALONE, AND INSIDE A CONDITIONAL STATEMENT. THE TABLE ALSO INCLUDES THE AVERAGE TIME FOR 7-INDEPENDENT HASHING WHICH IS NEEDED AS A SUBROUTINE OF THE PREVIOUS DISTINGUISHER BY [2]

For 64-bit keys, we performed some simple experiments summarized in Table I. We compared our `a*x<=t` distinguisher with the fastest known universal hashing scheme `a*x>>s` from [6]. The time bounds reported are averaged over 10 million executions with different values of x . More precisely, we had a loop where we in each round incremented x with the same 64-bit random number. Our reported time bounds include the cost of this addition and the cost of the loop, but these costs are much smaller than that of a multiplication. With $s = 63$, both schemes can be viewed as returning a single bit, and both combine a multiplication of 64-bit integers with a single AC^0 operation, so a similar performance could be expected. Our `a*x<=t` scheme took 1.02 ns whereas `a*x>>s` took 0.76 ns, so we were 34% slower. A possible explanation is that while a comparison itself is very fast, most computers do not return the answer as a bit in a register. Rather a comparison sets some flags intended for later branching. To check this in the context of a sampled sum, we tested “`if (a*x<=t) S+=x;`” which took 1.15 ns against “`if (a*x>>63) S+=x;`” which took 0.97 ns. Now the difference is less than 20%, and it is also smaller in absolute terms with 0.18 ns instead of 0.24 ns. With this fast branching in mind, we could implement the small bias sampler *SAMPLE* from (1) using logical AND (`&&`) rather than bit-wise AND (`&`), so as to benefit from short-circuit evaluation where the evaluation of a Boolean expression terminates as soon as the value is known. Thus we could implement *SAMPLE* as

```
SAMPLE(x) {B=0; for (i=0;i++;i<d) if (b[i]&&(a[i]*x<=t[i])) B++; return B&1;};
```

If `b[i]` is false, then with short-circuit evaluation, we skip immediately to the next iteration of the loop. In Section VII, we show for completeness that `a*key>>63` is not a distinguisher. Rather we think of `a*key>>63` as a general benchmark for fast hashing, and our new `a*x<=t` distinguisher was less than 50% slower.

Recall now that the previous distinguisher from [2, §3.1.1, pp. 843–846] involves 7-independent hashing as a subroutine, so it is at least as slow as 7-independent hashing. To test 7-independent hashing, we used the implementation of k -independent hashing of 64-bit keys from [16, §A.14]. This is a tuned implementation in standard portable C-code. It uses a random degree $d = k - 1$ polynomial modulo a prime $p > 2^{64}$. As recommended in [17], it uses a Mersenne prime $p = 2^{89} - 1$ so that “mod p ” can be computed with shifts and bit-wise Boolean operations. Still we have the issue that standard 64-bit multiplication works modulo 2^{64} , so several of these are needed to implement a single multiplication modulo $2^{89} - 1$. As described in [16, §A.14], the implementation ends up using $59d + 7$ standard instructions, including $6d$ multiplications. We have $d = 6$ for 7-independence, so we expect 7-independence to be much slower than our `a*x<=t` scheme. In our experiments from Table I, we see that 7-independent hashing takes 52 ns on the average. This is 50 times slower than the 1.02 ns of our `a*x<=t` scheme.

The relative performance of different codes vary a lot from computer to computer (see, e.g., the experiments in [16, Table 1] on three different computers). There are also other ways of implementing 7-independent hashing [4], particularly if one starts using more specialized instructions for multiplication over certain fields [18]. One

could also have considered double tabulation [19], but it is only relevant for much higher independence, e.g., for 64-bit keys, it requires 24 lookups in tables of size 2^{22} , but then it gives 100-independent hashing.

The above mentioned experiments only serve to illustrate that our new scheme $\text{a*x}<=t$ is much faster than that in [2]. This may be important in time-critical situations like in the processing of a high volume data stream. Perhaps more important for the potential of practical impact, our new $\text{a*x}<=t$ scheme is trivial to implement for anyone considering the theory of distinguishers and small bias sampling schemes for use in practice.

D. Contents

In Sections II–IV we will present our general distinguishers, starting from the mathematically simplest based on 2-independent hashing, and ending with the most efficient one from the title. In the remaining sections, we will discuss some possible alternatives, and their limitations.

II. USING 2-INDEPENDENT HASHING

First we present our simple generic distinguisher using an arbitrary 2-independent hash function $h : U \rightarrow [m]$ and a uniformly random $t \in [m]$, defining $\text{Sample} : U \rightarrow [2]$ as

$$\text{Sample}_t^h(x) = [h(x) \leq t].$$

We also have a value function $v : U \rightarrow R$ mapping keys $x \in U$ into elements of some commutative monoid R with plus as binary operator and zero the neutral element. We assume that v is non-zero for some key. Let $S \subseteq U$ be the set of keys with non-zero values, and let $n = |S| > 0$. We define the *sampled sum*

$$V^h(t) = \sum v(\text{Sample}_t^h) = \sum \{v(x) \mid x \in U \wedge \text{Sample}_t^h(x)\}.$$

To prove that Sample_t^h is a distinguisher with probability $(1 - n^2/m^2)/8$, we need to prove

$$\Pr[V^h(t) \neq 0] \geq (1 - n^2/m^2)/8. \quad (2)$$

In most of our analysis, we will think of the hash function h as fixed. Only at the very end will we use that h is a random 2-independent hash function. Let x_1, \dots, x_n denote the keys in S sorted according to their hash values, with ties broken arbitrarily. For $i = 1, \dots, n - 1$, let $I_i = [h(x_i), h(x_{i+1}))$. We view I_i as an interval of $h(x_{i+1}) - h(x_i)$ integers, which is empty if $h(x_i) = h(x_{i+1})$. Moreover, let $I_0 = [0, h(x_1))$ and $I_n = [h(x_n), m)$. Note for any I_i , that $V^h(t)$ is a constant for $t \in I_i$, and we denote this constant $V^h(I_i)$. Then $V^h(I_0) = 0$ and $V^h(I_i) = V^h(I_{i-1}) + v(x_i)$ for $i = 1, \dots, n$. We say an interval I_i is “good” if $V^h(I_i) \neq 0$; otherwise it is “bad”. Let \mathcal{GOOD}^h be the union of the good intervals. Then $V^h(t) \neq 0$ if and only if $t \in \mathcal{GOOD}^h$. When h is fixed while t is uniformly random in $[m]$, we have $\Pr[V^h(t) \neq 0] = |\mathcal{GOOD}^h|/m$, so when $h : U \rightarrow [m]$ is random,

$$\Pr_{h,t}[V^h(t) \neq 0] = E_h[|\mathcal{GOOD}^h|]/m.$$

Generally, for each key $x_i \in S$, we define the neighboring intervals $L^h(x_i) = I_{i-1}$ and $R^h(x_i) = I_i$. Since $V^h(I_i) = V^h(I_{i-1}) + v(x_i)$ for $i = 1, \dots, n$, we know that either $L^h(x_i)$ or $R^h(x_i)$ is good.

The analysis below is focused on an arbitrary $x \in S$. The length of the good interval around x is lower bounded by $\ell^h(x) = \min\{|L^h(x)|, |R^h(x)|\}$. Therefore $\sum_{x \in S} \ell^h(x)/2$ is a lower bound on $|\mathcal{GOOD}^h|$. The following observation follows directly from the definition of $\ell^h(x)$.

Observation 1. $\ell^h(x) \geq \delta$ if and only if the following conditions are satisfied:

- (a) for all $y \in S \setminus \{x\}$, $h(y) \notin (h(x) - \delta, h(x) + \delta)$, and
- (b) $\delta \leq h(x) \leq m - \delta$.

Lemma 2. If $h : U \rightarrow [m]$ is 2-independent and $|S| = n$, then $\Pr_h[\ell^h(x) \geq \delta] \geq 1 - n(2\delta - 1)/m$.

Proof: We are going to apply a union bound over the violating events from Observation 1. For each $y \in S \setminus \{x\}$ in (a), for any value of $h(x)$, there are $2\delta - 1$ violating values of $h(y)$, each happening with probability $1/m$. In (b) there are $2\delta - 1$ violating values of $h(x)$, each happening with probability $1/m$. By a union bound, the probability that at least one of all these violating events happen is bounded by $n(2\delta - 1)/m$. ■

Theorem 3. *If $h : U \rightarrow [m]$ is 2-independent and t is uniform in $[m]$, then $\text{Sample}(x) = [h(x) \leq t]$ is a $(1 - n^2/m^2)/8$ probability distinguisher.*

Proof: We have $\mathbb{E}_h[\ell^h(x)] = \sum_{\delta=1}^{\infty} \Pr[\ell^h(x) \geq \delta] \geq \sum_{\delta=1}^k \Pr[\ell^h(x) \geq \delta]$ for any integer k . Hence, by Lemma 2,

$$\begin{aligned} \mathbb{E}[\ell^h(x)] &\geq \sum_{\delta=1}^k \Pr[\ell^h(x) \geq \delta] \geq \sum_{\delta=1}^k (1 - (2\delta - 1)n/m) \\ &= k - (2k(k + 1)/2 - k)n/m = k - k^2n/m. \end{aligned}$$

The quadratic function $f : k \mapsto k - k^2n/m$ is concave and symmetric around its maximum in the reals at $k_0 = m/(2n)$. We pick the integer $k \in [k_0 - 1/2, k_0 + 1/2)$. Then

$$\begin{aligned} \mathbb{E}[\ell^h(x)] &\geq f(k) \geq f(k_0 - 1/2) \\ &= (m/(2n) - 1/2) - (m/(2n) - 1/2)^2n/m \\ &= (m/n - n/m)/4. \end{aligned}$$

Since $\sum_{x \in S} \ell^h(x)/2$ is a lower bound on $|\mathcal{GOOD}^h|$, we now get

$$\begin{aligned} \Pr_{h,t}[V^h(t) \neq 0] &= \mathbb{E}_h[|\mathcal{GOOD}^h|]/m \geq \mathbb{E}[\sum_{x \in S} \ell^h(x)]/(2m) \\ &\geq n/(2m) \cdot (m/n - n/m)/4 = (1 - n^2/m^2)/8. \end{aligned}$$

■

III. SIMPLE TUNING

So far we have considered h to be any 2-independent hash function, e.g., for some prime p , we could choose $a, b \in [p]$ uniformly at random, and then use the classic $h(x) = (ax + b) \bmod p$. We will argue that it suffices to use $h(x) = ax \bmod p$. Note that if $a \neq 0$, then h is a permutation. We will show that by picking a uniformly in $[p]_+ = [p] \setminus \{0\}$, we get a distinguisher with probability strictly bigger than $1/8$.

To analyze a hash function like $h(x) = ax \bmod p$, we need a better lower bound on the good interval around a key. Consider the case where $x = 0$ is the only key with non-zero value. Then $h(x) = 0$ so the smallest interval around x has length $\ell^h(x) = 0$. However, recall that $V^h(I_0) = 0$ while $V^h(I_1) = v(x_1) \neq 0$. To exploit this, we define a tighter lower bound $\ell_+^h(x)$ on the good interval around $h(x)$. We set $\ell_+^h(x_1) = |I_1|$ while for $i = 2, \dots, n$, $\ell_+^h(x_i) = \min\{|I_{i-1}|, |I_i|\}$. We can now tighten Observation 1 for the new measure:

Observation 4. $\ell_+^h(x) \geq \delta$ if and only if the following conditions are satisfied:

- (a) for all $y \in S \setminus \{x\}$, $|h(y) - h(x)| \geq \delta$, and
- (b) $h(x) \leq m - \delta$.

We note that Observation 4 (a) is equivalent to Observation 1 (a) while Observation 4 (b) has dropped the condition from Observation 1 (b) that $h(x) \geq \delta$.

The rest of the paper is devoted to hash functions of the form

$$h_a^m(x) = ax \bmod m, \tag{3}$$

both with m is a prime and with m a power of two. We shall use the notation

$$|x|_{\bmod m} = \min\{x \bmod m, -x \bmod m\}.$$

Also, for each $x \in S$, we define

$$S_x = \{y - x \mid y \in S, y \neq x\} \cup (\{x\} \setminus \{0\})$$

Note that $0 \notin S_x$.

Lemma 5. $\ell_+^{h_a^m}(x) \geq \delta$ if for all $z \in S_x$, $|az|_{\bmod m} \geq \delta$.

Proof: The lemma follows easily from Observation 4. First we note that $|h_a^m(y) - h_a^m(x)| < \delta$ implies $|h_a^m(y) - h_a^m(x)|_{\bmod m} < \delta$. Moreover, $|h_a^m(y) - h_a^m(x)|_{\bmod m} = |a(y - x)|_{\bmod m}$, so with $z = y - x$, our condition implies Observation 4 (a). Now, if $x = 0$, we always have $h_a^m(x) = 0$, and then Observation 4 (b) is satisfied. Otherwise $x \in S_x$, but then $|az|_{\bmod m} \geq \delta$ implies that $h_a^m(x) \leq m - \delta$, which again implies Observation 4 (b). ■

In the rest of this section, $m = p$ is a prime and a is uniform in $[p]_+$.

Lemma 6. If p is prime and a is uniformly distributed in $[p]_+ = \{1, \dots, p-1\}$, then for any $x \in S$, $\Pr_h[\ell_+^{h_a^p}(x) \geq \delta] \geq 1 - 2n(\delta - 1)/(p - 1)$.

Proof: The proof is similar to that of Lemma 2. From Lemma 5 we know that $\ell_+^{h_a^p}(x) \geq \delta$ unless for some $z \in S_x$, we have $|az|_{\bmod p} < \delta$. Each such z is non-zero and a uniform in $[p]_+$, so $az \bmod p$ is uniformly distributed in $[p]_+$. We therefore get $|az|_{\bmod p} < \delta$ only if $az \bmod p \in \{m - \delta + 1, \dots, m - 1, 1, \dots, \delta - 1\}$. This happens with probability $2(\delta - 1)/(p - 1)$. By a union bound, the probability that this happens for any of the at most n elements in S_x is at most $2n(\delta - 1)/(p - 1)$. ■

Theorem 7. If p is prime, a is uniform in $[p]_+$, and t is uniform in $[p]$, then $\text{Sample}(x) = [ax \bmod p \leq t]$ is a $1/8$ probability distinguisher.

Proof: The proof is similar to that of Theorem 3, except that we now use $\ell_+^{h_a^p}$ and Lemma 6. For any given integer k , we get

$$\begin{aligned} \mathbb{E}[\ell_+^{h_a^p}(x)] &\geq \sum_{\delta=1}^k \Pr[\ell_+^{h_a^p}(x) \geq \delta] \geq \sum_{\delta=1}^k (1 - 2n(\delta - 1)/(p - 1)) \\ &= k - k(k - 1)n/(p - 1) = k(1 + n/(p - 1)) - k^2n/(p - 1). \end{aligned}$$

The quadratic function $f : k \mapsto k(1 + n/(p - 1)) - k^2n/(p - 1)$ is concave and symmetric around its maximum in the reals at $k_0 = (p - 1)/(2n) + 1/2$. We pick the integer $k \in [k_0 - 1/2, k_0 + 1/2]$. Then

$$\begin{aligned} \mathbb{E}[\ell_+^{h_a^p}(x)] &\geq f(k) \geq f(k_0 - 1/2) \\ &= (p - 1)/(2n) \cdot (1 + n/(p - 1)) - ((p - 1)/(2n))^2n/(p - 1) \\ &= (p - 1)/(2n) + 1/2 - (p - 1)/(4n) = p/(4n) - 1/(4n) + 1/2 > p/(4n). \end{aligned}$$

The probability that we get distinction with a non-zero sampled sum $V^{h_a^p}(t)$ is thus lower bounded by

$$\Pr_{h,t}[V^{h_a^p}(t) \neq 0] = \mathbb{E}_h[|\mathcal{G}\mathcal{O}\mathcal{O}\mathcal{D}^{h_a^p}|]/p \geq \mathbb{E}[\sum_{x \in S} \ell_+^{h_a^p}(x)]/(2p) > n/(2p) \cdot p/(4n) = 1/8. \quad \blacksquare$$

IV. SAMPLE (X) = (A * X <= T)

We will now study a very computer friendly implementation of $h_a^m(x) = ax \bmod m$ from equation (3). For some integer parameter w , we set $m = 2^w$ and pick a as a uniformly distributed odd number in $[m]$. Since a is odd, it is relatively prime to m which is a power of two, so $h_a^{2^w}$ is a permutation on $[m]$. If w represents the bit length of a standard unit, that is, if $w \in \{8, 16, 32, 64\}$, then we get a particularly efficient portable code in a programming language such as C. When t is also a uniform w -bit number, we can compute the sampling decision for x with the full code $a * x <= t$. An important point here is that w -bit multiplication automatically discards overflow, so $a * x$ computes $ax \bmod 2^w$. We will show that this power-of-two scheme yields a distinguisher with probability bigger than $1/8$.

We note that the above power-of-two scheme with odd a has been studied before by Dietzfelbinger et al. [6]. They proved that the scheme $h_a^{m,s}(x) = \lfloor (ax \bmod 2^w) / 2^s \rfloor$ mapping $[2^w]$ into $[2^{w-s}]$ is universal in the sense that for any distinct $x, y \in [2^w]$, $\Pr[h(x) = h(y)] \leq 2^{s-w}$. This scheme is implemented efficiently as $(a * x) >> s$. Our analysis here is, however, quite different from that of Dietzfelbinger et al. [6].

By Lemma 5, for any k we choose, we get

$$\begin{aligned} \mathbb{E}[\ell_+^{h_a^{2^w}}(x)] &\geq \sum_{\delta=1}^k \Pr[\ell_+^{h_a^{2^w}}(x) \geq \delta] \geq \sum_{\delta=1}^k (1 - \sum_{z \in S_x} \Pr[|az| \bmod 2^w < \delta]) \\ &= k - \sum_{z \in S_x} \sum_{\delta=1}^k \Pr[|az| \bmod 2^w < \delta]. \end{aligned} \tag{4}$$

In the previous section, with m prime and random $a \neq [m]_+$, we proved that $\Pr[|az| \bmod m < \delta] \leq 2(\delta-1)/(p-1)$ for any $z \in [m]_+$ (cf. the proof of Lemma 6). However, for our power-of-two scheme with $m = 2^w$, we do not have a good general bound on $\Pr[|az| \bmod m < \delta]$ since the probability may change a lot with z . It turns out, however, that by exchanging summation order, we can get a good bound on the sum $\sum_{\delta=1}^k \Pr[|az| \bmod 2^w < \delta]$.

Lemma 8. *For any given $k, z \in [2^w]_+$, and uniformly random odd $a \in [2^w]$,*

$$\sum_{\delta=1}^k \Pr[|az| \bmod 2^w < \delta] \leq 2^{2-w} \lfloor k/2 \rfloor \lceil k/2 \rceil.$$

If k is even, the bound evaluates to $k^2/2^w$, and if k is odd, it is $(k-1)(k+1)/2^w = (k^2-1)/2^w$.

Proof: Let $k, z \in [2^w]_+$ be fixed, and let $i = \text{lsb}(z)$ denote the position of the least significant set bit in the binary representation of z , that is, $z = z'2^i$ for some odd z' . The multiplier a is a uniformly distributed odd w -bit number a , and this means that $az \bmod 2^w$ is a uniformly distributed odd multiple of 2^i , that is, a number of the form $2^i + j2^{i+1}$ for $j \in [2^{w-i-1}]$. There are the same number of such odd multiples in $(2^w - \delta)$ and in $(0, \delta)$, and none in 0, so we conclude that

$$\Pr[|az| \bmod 2^w < \delta] = 2 \cdot \sum_{j: 0 < 2^i + j2^{i+1} < \delta} 1/2^{w-i-1} = 2^{2-w} \cdot \sum_{j: 0 < 2^i + j2^{i+1} < \delta} 2^i.$$

Hence

$$\begin{aligned}
2^{w-2} \cdot \sum_{\delta=1}^k \Pr[|az| \bmod 2^w < \delta] &= \sum_{1 \leq \delta \leq k} \sum_{j: 0 < 2^i + j2^{i+1} < \delta} 2^i \\
&= \sum_{0 \leq j < (k-2^i)/2^{i+1}} \sum_{\delta=2^i + j2^{i+1} + 1}^k 2^i \\
&= \sum_{0 \leq j < (k-2^i)/2^{i+1}} (k - 2^i - j2^{i+1})2^i \\
&= \sum_{0 \leq j < J_k} (k - 2^i - j2^{i+1})2^i,
\end{aligned}$$

where $J_k = \lceil (k - 2^i)/2^{i+1} \rceil$. Hence

$$2^{w-2} \cdot \sum_{\delta=1}^k \Pr[|az| \bmod 2^w < \delta] = J_k 2^i (k - 2^i) - 2^{2i+1} J_k (J_k - 1)/2 = J_k 2^i (k - J_k 2^i).$$

The quadratic mapping $\alpha \mapsto \alpha(k - \alpha)$ is concave and symmetric around its maximum in the reals at $\alpha_0 = k/2$. The maximum in the integers is attained at the points $\lfloor k/2 \rfloor$ and $\lceil k/2 \rceil$ where it evaluates to $\lfloor k/2 \rfloor (k - \lfloor k/2 \rfloor) = \lfloor k/2 \rfloor \lceil k/2 \rceil$. We have $\alpha = J_k 2^i = \lceil (k - 2^i)/2^{i+1} \rceil 2^i$, which is indeed an integer, and equal to $\lceil k/2 \rceil$ when $i = 0$. We conclude that

$$2^{w-2} \cdot \sum_{\delta=1}^k \Pr[|az| \bmod 2^w < \delta] = J_k 2^i (k - J_k 2^i) \leq \lfloor k/2 \rfloor \lceil k/2 \rceil.$$

This completes the proof of Lemma 8. ■

Lemma 9. For any $x \in S$, when a is a uniform odd number in $[2^w]$, we have

$$\mathbb{E}[\ell_+^{h_a^{2^w}}(x)] \geq 2^w / (4n).$$

Proof: Inserting the bound of Lemma 8 in (4), for any odd positive integer k ,

$$\mathbb{E}[\ell_+^{h_a^{2^w}}(x)] = k - \sum_{z \in S_x} \sum_{\delta=1}^k \Pr[|az| \bmod 2^w < \delta] \geq k - n \cdot (k^2 - 1)/2^w.$$

The quadratic function $f : k \mapsto k - n \cdot (k^2 - 1)/2^w$ is concave and symmetric around its maximum in the reals at $k_0 = 2^{w-1}/n$. We pick k as the odd integer in $[k_0 - 1, k_0 + 1)$. Then

$$\begin{aligned}
\mathbb{E}[\ell_+^{h_a^{2^w}}(x)] &= f(k) \geq f(k_0 - 1) = (k_0 - 1) - n \cdot k_0(k_0 - 2)/2^w \\
&= \frac{2^{w-1}}{n} - 1 - \frac{n}{2^w} \cdot \frac{2^{2(w-1)}}{n^2} + \frac{n}{2^w} \cdot \frac{2^w}{n} \\
&= \frac{2^w}{2n} - 1 - \frac{2^w}{4n} + 1 = \frac{2^w}{4n}.
\end{aligned}$$

Theorem 10. If w is a natural number, a is a uniform odd number in $[2^w]$, and t is uniform in $[2^w]$, then $\text{Sample}(x) = [ax \bmod 2^w \leq t]$ is a distinguisher with probability $1/8$. ■

Proof: Recall that $\ell_+^{h_a^{2^w}}(x)$ is a lower bound on the size of the good interval neighboring $h_a^{2^w}(x)$, so $\sum_{x \in S} \ell_+^{h_a^{2^w}}(x)/2$ is a lower bound of the total length $|\mathcal{GOOD}^{h_a^{2^w}}|$ of the good intervals. Hence, by Lemma 9,

$$E_a[|\mathcal{GOOD}^{h_a^{2^w}}|] \geq \sum_{x \in S} E[\ell_+^{h_a^{2^w}}(x)]/2 > n(2^w/(4n))/2 = 2^w/8.$$

We get a distinguisher when the sampled sum $V^{h_a^{2^w}}(t)$ is non-zero. By definition, this is exactly when $t \in \mathcal{GOOD}^{h_a^{2^w}}$, which for a given $\mathcal{GOOD}^{h_a^{2^w}}$ happens with probability $|\mathcal{GOOD}^{h_a^{2^w}}|/2^w$. Therefore

$$\Pr_{a,t}[V^{h_a^{2^w}}(t) \neq 0] = E_h[|\mathcal{GOOD}^{h_a^{2^w}}|]/2^w \geq (2^w/8)/2^w = 1/8. \quad \blacksquare$$

Recall that when $w \in \{8, 16, 32, 64\}$, we implement $[ax \bmod 2^w \leq t]$ in C as $(a*x <= t)$. Theorem 10 thus formalizes the title of this paper.

V. RELATION TO AMS SKETCHES AND 4-INDEPENDENCE WITH REAL VALUES

When we have a value function $v : U \rightarrow \mathbb{R}$, we can use the classic AMS sketches [15] as a kind of distinguisher. These sketches are used to estimate the second moment $Q = \sum_{x \in U} (v(x))^2$, which is non-zero if and only if v has a non-zero.

The AMS sketch takes a random sign function $Sign : U \rightarrow \{-1, +1\}$ and computes the sketch

$$X = \sum_{x \in U} v(x) Sign(x).$$

Then $Y = X^2$ is used as an estimator for Q . In [15] it is proved that $E[X^2] = Q$ if $Sign$ is 2-independent and that $\text{Var}[X^2] < 2Q$ when $Q > 0$ and $Sign$ is 4-independent. From the simple observation below it follows that Y and hence X are non-zero with probability bigger than $1/3$.

Observation 11. *If $Y \geq 0$ is a random variable and $\Pr[Y = 0] = p$, then $\text{Var}[Y] \geq \frac{p}{1-p} E[Y]^2$. In particular, if $p \geq 2/3$ then $\text{Var}[Y] \geq 2E[Y]^2$.*

Proof: The observation follows from the more general Paley-Zygmund inequality, but for this simple case, we provide a direct proof. Given that $\Pr[Y = 0] = p$, we get the smallest variance for a given mean γ if Y is a constant y when $Y \neq 0$. Then $\Pr[Y = y] = (1 - p)$ and $\mu = (1 - p)y$, so $\text{Var}[Y] = E[Y^2] - \gamma^2 = (1 - p)(\gamma/(1 - p))^2 - \gamma^2 = \frac{p}{1-p} \gamma^2$. \blacksquare

We now observe that we get the same result if we replace the 4-independent sign-function $Sign : U \rightarrow \{-1, +1\}$ with a 4-independent sample function $Sample : U \rightarrow [2]$, that is,

Proposition 12. *A 4-independent Sample : $U \rightarrow [2]$ is a distinguisher with probability $1/3$ for real valued functions.*

Proof: Consider a non-zero value function $v : U \rightarrow \mathbb{R}$. Let $x_1, \dots, x_n \in \mathbb{R} \setminus \{0\}$ be the non-zero values in the multiset $v(U)$. Each x_i gets sampled with probability $1/2$, yielding a random variable X_i with $E[X_i] = \mu_i = x_i/2$. Our sampled sum is $X = \sum_i X_i$ which has mean $\mu = \sum_i \mu_i$. We want to argue that $X \neq 0$ with constant probability using that the X_i are 4-independent.

We are going to study the equivalent event that $X^2 \neq 0$. We will argue that

$$\text{Var}[X^2] < 2E[X^2]^2 \quad (5)$$

By Observation 11 this implies that $\Pr[X^2 = 0] < 2/3$, hence that $\Pr[X \neq 0] > 1/3$, as desired.

Since $\text{Var}[X^2] = E[X^4] - E[X^2]^2$, we have the following equivalent form of (5):

$$E[X^4] < 3E[X^2]^2. \quad (6)$$

To prove (6), we will compute $E[(X - \mu)^k]$ and $E[X^k]$ for $k = 2, 3, 4$. The calculations are straightforward because of the symmetry $\Pr[(X_i - \mu_i) = \mu_i] = \Pr[(X_i - \mu_i) = -\mu_i] = 1/2$. Therefore $E[(X_i - \mu_i)^k] = 0$ for k odd while $E[(X_i - \mu_i)^k] = \mu_i^k$ for k even. Thus

$$\begin{aligned} E[(X - \mu)^2] &= E[X^2] - \mu^2 \\ E[(X - \mu)^2] &= E[(\sum_i X_i - \mu_i)^2] = \sum_i E[(X_i - \mu_i)^2] = \sum_i \mu_i^2 \\ E[X^2] &= \sum_i \mu_i^2 + \mu^2. \end{aligned}$$

$$\begin{aligned} E[(X - \mu)^3] &= E[X^3] - 3E[X^2]\mu + 2\mu^3 \\ E[(X - \mu)^3] &= E[(\sum_i X_i - \mu_i)^3] = 0 \\ E[X^3] &= 3E[X^2]\mu - 2\mu^3 = 3(\sum_i \mu_i^2 + \mu^2)\mu - 2\mu^3 \\ &= 3\mu \sum_i \mu_i^2 + \mu^3 \end{aligned}$$

$$\begin{aligned} E[(X - \mu)^4] &= E[X^4] - 4E[X^3]\mu + 6E[X^2]\mu^2 - 3\mu^4 \\ E[(X - \mu)^4] &= E[(\sum_i X_i - \mu_i)^4] = E[\sum_i (X_i - \mu_i)^4] + 3E[\sum_{i \neq j} (X_i - \mu_i)^2 (X_j - \mu_j)^2] \\ &= \sum_i \mu_i^4 + 3 \sum_{i \neq j} \mu_i^2 \mu_j^2 = 3(\sum_i \mu_i^2)^2 - 2 \sum_i \mu_i^4. \end{aligned}$$

$$\begin{aligned} E[X^4] &= E[(X - \mu)^4] + 4E[X^3]\mu - 6E[X^2]\mu^2 + 3\mu^4 \\ &= 3(\sum_i \mu_i^2)^2 - 2 \sum_i \mu_i^4 + 4(3\mu \sum_i \mu_i^2 + \mu^3)\mu - 6(\sum_i \mu_i^2 + \mu^2)\mu^2 + 3\mu^4 \\ &= 3(\sum_i \mu_i^2)^2 - 2 \sum_i \mu_i^4 + 12\mu^2 \sum_i \mu_i^2 + 4\mu^4 - 6\mu^2 \sum_i \mu_i^2 - 6\mu^4 + 3\mu^4 \\ &= 3(\sum_i \mu_i^2)^2 - 2 \sum_i \mu_i^4 + 6\mu^2 \sum_i \mu_i^2 + \mu^4 \\ &= 2E[X^2]^2 - 2 \sum_i \mu_i^4 - 2\mu^4 \\ &< 2E[X^2]^2. \end{aligned}$$

The last strict inequality follows because v was non-zero, hence that we have some $\mu_i \neq 0$. This completes the proof of (6), hence of (5). ■

Proposition 12 forms an interesting contrast to the fact that less than full independence does not suffice for a distinguisher for values in \mathbb{F}_2 .

Contrasting the above result, we will now argue that 2-independence does not suffice for a distinguisher for the reals.

Proposition 13. *There is a 2-independent sampling function $\text{Sample} : [u] \rightarrow [2]$ and a non-zero value function $v : [u] \rightarrow \mathbb{R}$ such that $\Pr[\sum v(\text{Sample}) \neq 1] \leq 4/u$.*

Proof: First we describe the value function. We assume $u = 4n$ for some positive integer n . We have $2n$ values that are $+1$ and $2n$ values that are -1 . The positive and negative values will be sampled independently, but using the same scheme. In both cases, for some parameter ε , the distribution samples nothing with probability ε and all with probability ε . Otherwise make a “balanced” sample of exactly n random values. If we get the balanced case for both negative and positive values, then the sampled sum is zero, and this happens with probability at least $1 - 4\varepsilon$.

We will now fix ε so as to get a 2-independent distribution. Regardless of ε , in the marginal distribution, each item is sampled with probability $1/2$. Because everything is symmetric, it now suffices to set ε such that any 2 given items are sampled with probability $1/4$. Thus we want

$$1/4 = \varepsilon + (1 - 2\varepsilon)(1/2)(n - 1)/(2n - 1) = \frac{\varepsilon(4n - 2) + n - 1 - 2\varepsilon(n - 1)}{4n - 2} = \frac{n - 1 + 2\varepsilon n}{4n - 2}$$

$$\iff \varepsilon = 1/(4n) = 1/u.$$

The probability of getting a non-zero sum is thus at most $4/u$. ■

VI. SIMPLE TABULATION

We will now briefly discuss what can be achieved using simple tabulation for constructing distinguishers. Simple tabulation dates back to Zobrist [20]. We view a key $x \in [u] = \{0, \dots, u - 1\}$ as a vector of $c > 1$ characters $x_0, \dots, x_{c-1} \in \Sigma = [u^{1/c}]$. For each character position, we initialize a fully random table H_i of size $|\Sigma|$ with values from $R = [2^r]$. The hash value of a key x is calculated as

$$h(x) = H_0[x_0] \oplus \dots \oplus H_{c-1}[x_{c-1}] .$$

Pătraşcu and Thorup [21] analyzed simple tabulation, showing that it works well in many algorithmic contexts. It is now tempting to set $r = 1$ and use $Sample = h : [u] \rightarrow [2]$ as a sampler.

We will now argue that the above sampler is not a distinguisher for values in \mathbb{F}_2 . Suppose $c = 2$ and consider the 4 keys aa , ab , ba , and bb . Then over \mathbb{F}_2 , we have $Sample(aa) + Sample(ab) + Sample(ba) + Sample(bb) = 0$. This is indeed the standard proof that simple tabulation is not 4-independent. It implies that if our four keys are the only ones with value 1, then the sampled sum is always 0.

On the other hand, we do claim that our simple tabulation sampler is a constant probability distinguisher over the reals. In Proposition 12 we proved this to be the case when $Sample$ is 4-independent, but simple tabulation is not 4-independent. However, in [21] it is proved that the fourth moment that we studied in the proof of Proposition 12 deviates only by a constant factor from that with full randomness, and all lower moments are the same since simple tabulation is 3-independent. We can therefore perform the same analysis as that in the proof of Proposition 12, and conclude that simple tabulation sampling is a constant probability distinguisher for real values.

In the experiments with a 64-bit computer from [21, Table 1], the universal hashing `a*x>>s` takes 2.33 ns whereas simple tabulation takes 11.40 ns. In our own experiments, our distinguisher `a*x<=t` was only 34% slower than `a*x>>s`, so it should be 2-3 times faster than simple tabulation (the relative difference may vary a lot from computer to computer since `a*x<=t` depends on arithmetic operations while tabulation depends on the speed of cache). On top of that, simple tabulation uses much more space and it is not a general distinguisher. However, simple tabulation does have one interesting advantage as a distinguisher for reals; for if we use $r > 1$ output bits in h and the tables H_i , then they are completely independent, yielding r independent distinguishers that can be used for the vector distinguisher mentioned in Section I-A. Unfortunately, simple tabulation is not a general distinguisher, and we do not know of a better way of generating bits for a general vector distinguisher than to generate each one individually using our new `a*x<=t` scheme. This is left as an interesting open problem.

VII. OTHER NON-DISTINGUISHERS

We will now prove that some other tempting sampling schemes are not distinguishers.

In the introduction we considered the fastest universal hash function $a \cdot x \gg 63$ mapping 64-bit keys to [2]. Consider applying it to a pair of keys (x_1, x_2) differ in only in their most significant bit. If the multiplier is odd, then $ax_1 \bmod 2^{64}$ and $ax_2 \bmod 2^{64}$ will also only differ in the most significant bit, so $a \cdot x \gg 63$ will sample exactly one of them. On the other hand, if a is even, we will either sample both x_1 and x_2 , or none of them. Suppose we have one more pair (y_1, y_2) , also differing in only the most significant bit. Then if a is odd, we sample one key from each pair, and if a is even, we sample either both or none from each pair. In all cases we end up with an even number of keys. Thus, if x_1, x_2, y_1, y_2 are the keys with value 1, then the sampled sum over \mathbb{F}_2 is always 0.

Now note that $a \cdot x \gg 63$ is equivalent to $a \cdot x > \tau$ when τ is the constant 2^{63} , and that $a \cdot x < \tau$ would sample the complement set that also has an even number of elements. The problem persists for schemes like $a \cdot x + b < \tau$ as long as $\tau = 2^{63}$. It is thus no coincidence that our $a \cdot x < \tau$ distinguisher relies on τ being random.

ACKNOWLEDGMENTS

I would like to thank Martin Dietzfelbinger for suggesting some very nice simplifications of the analysis including the change of summation order used in the proof of Lemma 8. Also, I would like to thank a very thorough reviewer from FOCS'15 for a 6 page report with many insightful comments. Finally, I thank Noga Alon for pointing me to the Paley-Zygmund inequality.

Mikkel Thorup's research is partly supported by Advanced Grant DFF-0602-02499B from the Danish Council for Independent Research under the Sapere Aude research career programme.

REFERENCES

- [1] U. Vazirani, "Randomness, adversaries, and computation," Ph.D. dissertation, University of California, Berkeley, 1986.
- [2] J. Naor and M. Naor, "Small-bias probability spaces: Efficient constructions and applications," *SIAM J. Comput.*, vol. 22, no. 4, pp. 838–856, 1993, announced at STOC'90.
- [3] B. Chor and O. Goldreich, "On the power of two-point based sampling," *J. Complexity*, vol. 5, no. 1, pp. 96–106, 1989. [Online]. Available: [http://dx.doi.org/10.1016/0885-064X\(89\)90015-0](http://dx.doi.org/10.1016/0885-064X(89)90015-0)
- [4] M. N. Wegman and L. Carter, "New classes and applications of hash functions," *Journal of Computer and System Sciences*, vol. 22, no. 3, pp. 265–279, 1981, see also FOCS'79.
- [5] M. Dietzfelbinger, "Universal hashing and k -wise independent random variables via integer arithmetic without primes," in *Proc. 13th Symposium on Theoretical Aspects of Computer Science (STACS)*, 1996, pp. 569–580.
- [6] M. Dietzfelbinger, T. Hagerup, J. Katajainen, and M. Penttonen, "A reliable randomized algorithm for the closest-pair problem," *Journal of Algorithms*, vol. 25, no. 1, pp. 19–51, 1997.
- [7] R. Freivalds, "Probabilistic machines can use less running time," in *IFIP Congress*, 1977, pp. 839–842.
- [8] T. Kimbrel and R. K. Sinha, "A probabilistic algorithm for verifying matrix products using $O(n^2)$ time and $\log_2 n + O(1)$ random bits," *Inf. Process. Lett.*, vol. 45, no. 2, pp. 107–110, 1993. [Online]. Available: [http://dx.doi.org/10.1016/0020-0190\(93\)90224-W](http://dx.doi.org/10.1016/0020-0190(93)90224-W)
- [9] V. King, S. Kutten, B. Mountjoy, and M. Thorup, "Construction and impromptu repair of an MST in a distributed network with $o(m)$ communication," 2015.
- [10] K. J. Ahn, S. Guha, and A. McGregor, "Analyzing graph structure via linear measurements," in *Proc. 25th ACM/SIAM Symposium on Discrete Algorithms (SODA)*, 2012, pp. 459–467.

- [11] B. M. Kapron, V. King, and B. Mountjoy, “Dynamic graph connectivity in polylogarithmic worst case time,” in *Proc. 25th ACM/SIAM Symposium on Discrete Algorithms (SODA)*, 2013, pp. 1131–1142.
- [12] N. Alon, O. Goldreich, J. Håstad, and R. Peralta, “Simple construction of almost k-wise independent random variables,” *Random Struct. Algorithms*, vol. 3, no. 3, pp. 289–304, 1992.
- [13] ———, “Addendum to ”simple construction of almost k-wise independent random variables”,” *Random Struct. Algorithms*, vol. 4, no. 1, pp. 119–120, 1993.
- [14] A. Ben-Aroya and A. Ta-Shma, “Constructing small-bias sets from algebraic-geometric codes,” *Theory of Computing*, vol. 9, pp. 253–272, 2013, announced at FOCS’09.
- [15] N. Alon, Y. Matias, and M. Szegedy, “The space complexity of approximating the frequency moments,” *Journal of Computer and System Sciences*, vol. 58, no. 1, pp. 209–223, 1999, see also STOC’96.
- [16] M. Thorup and Y. Zhang, “Tabulation-based 5-independent hashing with applications to linear probing and second moment estimation,” *SIAM Journal on Computing*, vol. 41, no. 2, pp. 293–331, 2012, announced at SODA’04 and ALENEX’10.
- [17] L. Carter and M. N. Wegman, “Universal classes of hash functions,” *Journal of Computer and System Sciences*, vol. 18, no. 2, pp. 143–154, 1979, see also STOC’77.
- [18] S. Gueron and M. E. Kounavis, “Intel[®] carry-less multiplication instruction and its usage for computing the CGM mode. Intel Corporation White Paper 323640-001, Rev. 2.0, 2010,” 2010.
- [19] M. Thorup, “Simple tabulation, fast expanders, double tabulation, and high independence,” in *Proc. 25th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2013, pp. 90–99.
- [20] A. L. Zobrist, “A new hashing method with application for game playing,” Computer Sciences Department, University of Wisconsin, Madison, Wisconsin, Tech. Rep. 88, 1970.
- [21] M. Pătraşcu and M. Thorup, “The power of simple tabulation-based hashing,” *Journal of the ACM*, vol. 59, no. 3, p. Article 14, 2012, announced at STOC’11.