

On Parallel Scalable Uniform SAT Witness Generation * **

Supratik Chakraborty¹, Daniel J. Fremont², Kuldeep S. Meel³,
Sanjit A. Seshia², and Moshe Y. Vardi³

¹ Indian Institute of Technology, Bombay

² University of California, Berkeley

³ Department of Computer Science, Rice University

Abstract. Constrained-random verification (CRV) is widely used in industry for validating hardware designs. The effectiveness of CRV depends on the uniformity of test stimuli generated from a given set of constraints. Most existing techniques sacrifice either uniformity or scalability when generating stimuli. While recent work based on random hash functions has shown that it is possible to generate almost uniform stimuli from constraints with 100,000+ variables, the performance still falls short of today’s industrial requirements. In this paper, we focus on pushing the performance frontier of uniform stimulus generation further. We present a random hashing-based, easily parallelizable algorithm, **UniGen2**, for sampling solutions of propositional constraints. **UniGen2** provides strong and relevant theoretical guarantees in the context of CRV, while also offering significantly improved performance compared to existing almost-uniform generators. Experiments on a diverse set of benchmarks show that **UniGen2** achieves an average speedup of about 20× over a state-of-the-art sampling algorithm, even when running on a single core. Moreover, experiments with multiple cores show that **UniGen2** achieves a near-linear speedup in the number of cores, thereby boosting performance even further.

1 Introduction

Functional verification is concerned with the verification and validation of a *Design Under Verification* (DUV) with respect to design specifications. With the increasing complexity of DUVs, functional verification has become one of the most challenging and time-consuming steps in design validation [3]. In view of the

* The full version is available at <http://www.cs.rice.edu/CS/Verification/Projects/UniGen>.

** The authors would like to thank Suguman Bansal and Karthik Murthy for valuable comments on the earlier drafts, Armando Solar-Lezama for benchmarks, and Mate Soos for tweaking CMS to support UniGen2. This work was supported in part by NSF grants CNS 1049862, CCF-1139011, CCF-1139138, by NSF Expeditions in Computing project "ExCAPE: Expeditions in Computer Augmented Program Engineering", by BSF grant 9800096, by a gift from Intel, by a grant from Board of Research in Nuclear Sciences, India, by the Shared University Grid at Rice funded by NSF under Grant EIA-0216467 and a partnership between Rice University, Sun Microsystems, and Sigma Solutions, Inc., and by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

high computational cost of formal verification, simulation-based techniques have been extensively employed in industrial practice. The success of such techniques depends on the *quality* of input stimuli with which the design is simulated. The generation of high-quality stimuli that uncover hidden bugs continues to be a challenging problem even today [21].

The problem of high-quality stimulus generation has led to the emergence of *constrained-random simulation*, also known as *constrained-random verification* (CRV) [22]. In CRV, a verification engineer is tasked with the construction of verification scenarios, expressed as constraints over stimuli. Typically, constructing these scenarios involves applying past user experience, inputs from design engineers, and domain-specific knowledge. A constraint solver is then invoked to generate random stimuli satisfying the constraints. Since the distribution of errors in the design is not known *a priori*, each random stimulus is just as likely to produce an error as any other. Therefore, achieving a uniformly random distribution over stimuli satisfying the constraints is highly desirable.

While constraint-solving technologies have witnessed significant advancements over the last decade, methods of generating uniformly distributed solutions still face huge scalability hurdles. This has been observed repeatedly in the literature [6] and by industry practitioners⁴. In this paper, we take a step towards remedying the current situation by proposing an easily parallelizable sampling algorithm for Boolean constraints that provides strong theoretical guarantees (similar to those provided by an almost-uniform generator) in the context of CRV, and also runs significantly faster than current state-of-the-art techniques on a diverse set of benchmark problems.

Since constraints arising in CRV can often be encoded as propositional formulae in conjunctive normal form (CNF), we focus on almost-uniform sampling of satisfying assignments of CNF formulae (known as *SAT witnesses*). This problem has been extensively studied in both theoretical and practical contexts, and has many applications, including probabilistic reasoning, approximate model counting, and Markov logic networks [4]. Until recently, approaches to solving this problem belonged to one of two classes: those which provide strong guarantees of uniformity but scale poorly [2,24], and those which scale to large problem instances but rely on heuristics and hence offer very weak or no uniformity guarantees [11,16,14].

Recently, Chakraborty, Meel, and Vardi [4] proposed a new algorithmic approach to bridge the gap between these two extremes. The main idea behind their approach is to use universal hashing in order to partition the space of witnesses into roughly equal “cells”. Under an appropriate partitioning scheme, choosing a random witness from a randomly chosen cell provides strong uniformity guarantees. The most recent instance of this approach is called UniGen [6]. While UniGen scales to formulae much larger than those that can be handled by previous state-of-the-art techniques, the runtime performance of UniGen still falls short of industry requirements.

⁴ Private Communication: R. Kurshan

Since the end of Dennard scaling, there has been a strong revival of interest in parallelizing a wide variety of algorithms to achieve improved performance [10]. One of the main goals in parallel-algorithm design is to achieve a speedup nearly linear in the number of processors, which requires the avoidance of dependencies among different parts of the algorithm [8]. Most of the sampling algorithms used for uniform witness generation fail to meet this criterion, and are hence not easily parallelizable. In contrast, the algorithm proposed in this paper is inherently parallelizable, and achieves a near-linear speedup.

Our primary contribution is a new algorithm, UniGen2, that addresses key performance deficiencies of UniGen. Significantly, UniGen2 generates many more samples (witnesses) per iteration compared to UniGen, thereby reducing the number of SAT calls required per sample to a *constant*. While this weakens the guarantee of independence among samples, we show that this does not hurt the primary objective of CRV. Specifically, we prove that UniGen2 provides almost as strong guarantees as UniGen with respect to discovery of bugs in a CRV setting. On the practical front, we present an implementation of UniGen2, and show by means of extensive experiments that it significantly outperforms existing state-of-the-art algorithms, while generating sample distributions that are indistinguishable from those generated by an ideal uniform sampler. UniGen2 is also inherently parallelizable, and we have implemented a parallel version of it, called parallel UniGen2. Our experiments show that parallel UniGen2 achieves a near-linear speedup with the number of cores.

2 Notation and Preliminaries

Let F denote a Boolean formula in conjunctive normal form (CNF), and let X be the set of variables appearing in F . The set X is called the *support* of F . Given a set of variables $S \subseteq X$ and an assignment σ of truth values to the variables in X , we write $\sigma|_S$ for the projection of σ on S . A *satisfying assignment* or *witness* of F is an assignment that makes F evaluate to true. We denote the set of all witnesses of F by R_F and the projection of R_F on S by $R_{F|S}$. For the rest of the paper, we use S to denote the *sampling set*, the set of variables on which we desire assignments to be projected. Even when no projection is explicitly desired, S can often be restricted to a small subset of X , called an *independent support*, such that $|R_{F|S}| = R_F$ (see [6] for details). For notational simplicity, we omit mentioning F and S when they are clear from the context.

We use $\Pr[X]$ to denote the probability of event X . We say that a set of events $\{E_1, E_2, \dots, E_n\}$ are (l, u) *almost-independent almost-identically distributed* (henceforth, (l, u) -*a.a.d.*) if if $\forall i, l \leq \Pr[E_i] \leq u$ and $l \leq \Pr[E_i | (\{E_1, E_2, \dots, E_n\} \setminus E_i)] \leq u$. Note that this is similar to the notion of independently identically distributed (i.i.d.) events, but somewhat weaker.

Given a Boolean formula F and sampling set S , a *probabilistic generator* of witnesses of $R_{F|S}$ is a probabilistic algorithm that generates a random witness in $R_{F|S}$. A *uniform generator* $\mathcal{G}^u(\cdot, \cdot)$ is a probabilistic generator that guarantees $\Pr[\mathcal{G}^u(F, S) = y] = 1/|R_{F|S}|$, for every $y \in R_{F|S}$. An *almost-uniform generator*

$\mathcal{G}^{au}(\cdot, \cdot, \cdot)$ relaxes the above guarantees, ensuring only that $1/((1 + \varepsilon)|R_{F|S}|) \leq \Pr[\mathcal{G}^{au}(F, S, \varepsilon) = y] \leq (1 + \varepsilon)/|R_{F|S}|$ for every $y \in R_{F|S}$. Probabilistic generators are allowed to occasionally “fail” by returning no witness although $R_{F|S} \neq \emptyset$. The failure probability must be bounded by a constant strictly less than 1.

A special class of hash functions, called *r-wise independent* hash functions, play a crucial role in our work. Let n, m and r be positive integers, and let $H(n, m, r)$ denote a family of r -wise independent hash functions mapping $\{0, 1\}^n$ to $\{0, 1\}^m$. We use $h \xleftarrow{R} H(n, m, r)$ to denote the probability space obtained by choosing a hash function h uniformly at random from $H(n, m, r)$. The property of r -wise independence guarantees that for all $\alpha_1, \dots, \alpha_r \in \{0, 1\}^m$ and for all distinct $y_1, \dots, y_r \in \{0, 1\}^n$, $\Pr[\bigwedge_{i=1}^r h(y_i) = \alpha_i : h \xleftarrow{R} H(n, m, r)] = 2^{-mr}$. For every $\alpha \in \{0, 1\}^m$ and $h \in H(n, m, r)$, let $h^{-1}(\alpha)$ denote the set $\{y \in \{0, 1\}^n \mid h(y) = \alpha\}$. Given $R_{F|S} \subseteq \{0, 1\}^{|S|}$ and $h \in H(|S|, m, r)$, we use $R_{F|S, h, \alpha}$ to denote the set $R_{F|S} \cap h^{-1}(\alpha)$. If we keep h fixed and let α range over $\{0, 1\}^m$, the sets $R_{F|S, h, \alpha}$ form a partition of $R_{F|S}$.

We use a particular class of such hash functions, denoted by $H_{xor}(n, m)$, which is defined as follows. Let $h(y)[i]$ denote the i^{th} component of the vector $h(y)$. This family of hash functions is then defined as $\{h \mid h(y)[i] = a_{i,0} \oplus (\bigoplus_{k=1}^n a_{i,k} \cdot y[k]), a_{i,j} \in \{0, 1\}, 1 \leq i \leq m, 0 \leq j \leq n\}$, where \oplus denotes the XOR operation. By choosing values of $a_{i,j}$ randomly and independently, we can effectively choose a random hash function from $H_{xor}(n, m)$. It was shown in [12] that this family is 3-wise independent. For every $m \in \{1, \dots, |S| - 1\}$, the m^{th} prefix-slice of h , denoted $h^{(m)}$, is a map from $\{0, 1\}^{|S|}$ to $\{0, 1\}^m$, such that $h^{(m)}(y)[i] = h(y)[i]$, for all $y \in \{0, 1\}^{|S|}$ and for all $i \in \{1, \dots, m\}$. Similarly, the m^{th} prefix-slice of α , denoted $\alpha^{(m)}$, is an element of $\{0, 1\}^m$ such that $\alpha^{(m)}[i] = \alpha[i]$ for all $i \in \{1, \dots, m\}$.

3 Related Work

Uniform generation of SAT witnesses was first studied by Jerrum, Valiant, and Vazirani [15], who showed that the problem can be solved in probabilistic polynomial time given access to a Σ_2^P oracle. In addition, they showed that almost-uniform generation is polynomially inter-reducible with approximate model counting. Bellare, Goldreich, and Petrank [2] improved this result and provided an algorithm in BPP^{NP} . Unfortunately, their algorithm fails to scale beyond a few tens of variables in practice [4]. A completely different approach to uniform generation of SAT witnesses is due to Yuan et al. [24], wherein a sample is generated by performing a random walk over a weighted binary decision diagram (WBDD). The high space requirement of this technique limits its applicability in practice.

In several settings (some industrial), generation of stimuli for CRV is typically done via heuristic methods that provide very weak or no guarantees of uniformity. One of the earliest such methods was to randomly seed a SAT solver [19]. While this is simple in principle, the distributions generated by random seeding have been shown to be highly skewed in practice [17]. An alternative approach

focusing on the generation of “diverse” solutions was proposed by Nadel [20], but it also fails to provide theoretical guarantees of coverage.

Markov Chain Monte Carlo (MCMC) algorithms, such as those based on simulated annealing or the Metropolis-Hastings algorithm, have been studied extensively in the literature [18] in the context of generating samples from a probability space. The eventual convergence to the target distribution for MCMC methods is often impractically slow in practice under mild requirements. Most MCMC-based sampling tools therefore use heuristic adaptations [17,16] to improve performance and reduce correlation between samples. Unfortunately, these heuristics significantly weaken or even destroy the theoretical guarantees.

Interval propagation [14] has been used extensively in industrial practice to achieve scalable stimulus generation. Techniques based on interval propagation, however, generate highly non-uniform distributions. Recent efforts via the conversion of constraints into belief networks [11,7] have also failed to achieve the desired balance between performance and guarantees of uniformity.

Recently, several random hashing-based techniques have been proposed to bridge the wide gap between scalable algorithms and those that give strong guarantees of uniformity when sampling witnesses of propositional constraints [4,6,9]. Hashing-based sampling techniques were originally pioneered by Sipser [23] and further used by Jerrum, Valiant, and Vazirani [15], and Bellare, Goldreich, and Petrank [2]. The key idea in hashing-based techniques is to first partition the space of satisfying assignments into small “cells” of roughly equal size using r -wise independent hash functions (for a suitable value of r), and then randomly choose a solution from a randomly picked cell. Bellare et al. showed that by choosing $r = n$ (where the propositional constraint has n variables), we can guarantee uniform generation. The resulting algorithm, however, does not scale in practice. Chakraborty, Meel, and Vardi [4] subsequently showed that with $r = 3$, a significantly more scalable near-uniform generator named UniWit can be designed. Building on the principle underlying UniWit, Ermon et al. [9] suggested further algorithmic improvements to uniform generation of witnesses.

Recently, Chakraborty et al. proposed a new algorithm named UniGen [5], which improves upon the ideas of UniWit. In particular, UniGen provides stronger guarantees of uniformity by exploiting a deep connection between approximate counting and almost-uniform sampling [15]. Furthermore, UniGen has been shown to scale to formulae with hundreds of thousands of variables. Even so, UniGen is typically 2-3 orders of magnitude slower than a single call to a SAT solver and therefore, its runtime performance falls short of the performance of heuristic methods commonly employed in industry to generate stimuli for CRV ⁵. In this paper, we offer several improvements to UniGen and obtain an algorithm with substantially improved performance that can be further scaled by parallelization to match the requirements of industry.

⁵ A random-constrained test case generator is typically allowed to be 10× slower than a constraint solver (private communication with industry expert W. Hung).

4 A Parallel SAT Sampler

In this section, we first motivate the need for sampling solutions of constraints in parallel, and then provide technical details of our algorithm, UniGen2.

Parallelization:

While simulation-based verification typically involves running in parallel many simulations with different input stimuli, the generation of these stimuli is often done sequentially. This is because existing approaches to stimulus generation are not efficiently parallelizable without degrading guarantees of uniformity. For example, approaches based on random seeding of a SAT solver maintain information about which regions of the solution space have already been explored, since the random seed often is not enough to steer the solver towards new regions of the solution space [17]. Different threads generating solutions must therefore communicate with each other, impeding efficient parallelization. In MCMC-based approaches, to generate independent samples in parallel each thread has to take a walk until a stationary distribution is reached. This often takes exponential time in the case of hard combinatorial spaces with complex internal structure [9]. Heuristics to speed up MCMC-based techniques destroy guarantees of uniformity even in the sequential case [17]. Methods based on random walks on WBDDs are amenable to parallelization, but they are known not to scale beyond a few hundred variables. The lack of techniques for sampling solutions of constraints in parallel while preserving guarantees of effectiveness in finding bugs is therefore a major impediment to high-performance CRV.

The algorithm UniGen2 presented in this section takes a step forward in addressing the above problem. It has an initial preprocessing step that is sequential but low-overhead, followed by inherently parallelizable sampling steps. It generates samples (stimuli) that are provably almost as effective as those generated by a uniform sampler for purposes of detecting a bug. Furthermore, our experiments demonstrate that a parallel implementation of UniGen2 achieves a near-linear speedup in the number of processor cores. Given that current practitioners are forced to trade guarantees of effectiveness in bug hunting for scalability, the above properties of UniGen2 are significant. Specifically, they enable a new paradigm of CRV wherein parallel stimulus generation and simulation can provide the required runtime performance while also providing theoretical guarantees.

Algorithm:

Our algorithm, named UniGen2, bears some structural similarities with the UniGen algorithm proposed earlier in [6]. Nevertheless, there are key differences that allow UniGen2 to outperform UniGen significantly. Like UniGen, UniGen2 takes a CNF formula F , a sampling set S and a tolerance ε (that is chosen to be at least 6.84 for technical reasons). Note that the formula F and set S uniquely define the solution set $R_{F|S}$.

Similarly to UniGen, UniGen2 works by partitioning $R_{F|S}$ into “cells” using random hash functions, then randomly selecting a cell by adding appropriate

constraints to F . If the chosen cell has the right size (where the acceptable size range depends on the desired tolerance ε), we can enumerate all the solutions in it and return a uniform random sample from among them. Unlike `UniGen`, however, `UniGen2` samples multiple times from the same cell. This decreases the generation time per sample by a large factor (about $10\times$ in our experiments), while preserving strong guarantees of effectiveness of the samples in finding bugs.

Algorithm 1 EstimateParameters(F, S, ε)

/* Returns (hashBits, loThresh, hiThresh) as required by `GenerateSamples` */

```

1: Find  $\kappa \in (0, 1)$  such that  $\varepsilon = (1 + \kappa)(7.44 + \frac{0.392}{(1-\kappa)^2}) - 1$ 
2: pivot  $\leftarrow \lceil 4.03 (1 + \frac{1}{\kappa})^2 \rceil$ 
3: hiThresh  $\leftarrow \lceil 1 + \sqrt{2}(1 + \kappa)\text{pivot} \rceil$ ; loThresh  $\leftarrow \lfloor \frac{1}{\sqrt{2}(1+\kappa)}\text{pivot} \rfloor$ 
4:  $i \leftarrow 0$ 
5: while  $i < n$  do
6:    $i \leftarrow i + 1$ 
7:   Choose  $h$  at random from  $H_{xor}(|S|, i)$ 
8:   Choose  $\alpha$  at random from  $\{0, 1\}^i$ 
9:    $Y \leftarrow \text{BSAT}(F \wedge (h(S) = \alpha), 61, S)$ 
10:  if  $1 \leq |Y| \leq 60$  then
11:    return (round( $\log |Y| + i + \log 1.8 - \log \text{pivot}$ ), loThresh, hiThresh)
12: return  $\perp$ 

```

Algorithm 2 GenerateSamples($F, S, \text{hashBits}, \text{loThresh}, \text{hiThresh}$)

```

1: Pick an order  $V$  of the values  $\{\text{hashBits} - 2, \text{hashBits} - 1, \text{hashBits}\}$ 
2: Choose  $h$  at random from  $H_{xor}(|S|, \text{hashBits})$ 
3: Choose  $\alpha$  at random from  $\{0, 1\}^{\text{hashBits}}$ 
4: for  $i \in V$  do
5:    $Y \leftarrow \text{BSAT}(F \wedge (h_i(S) = \alpha_i), \text{hiThresh}, S)$ 
6:   if ( $\text{loThresh} \leq |Y| < \text{hiThresh}$ ) then
7:     return loThresh distinct random elements of  $Y$ 
8: return  $\perp$ 

```

`UniGen2` is an algorithmic framework that operates in two stages: the first stage, `EstimateParameters` (Algorithm 1), performs low-overhead one-time pre-processing for a given F , S , and ε to compute numerical parameters ‘hashBits’, ‘loThresh’, and ‘hiThresh’. The quantity `hashBits` controls how many cells $R_{F|S}$ will be partitioned into, while `loThresh` and `hiThresh` delineate the range of acceptable sizes for a cell. In the second stage, `GenerateSamples` (Algorithm 2) uses these parameters to generate `loThresh` samples. If more samples are required, `GenerateSamples` is simply called again with the same parameters. Theorem 3 below shows that invoking `GenerateSamples` multiple times does not cause the loss of any theoretical guarantees. We now explain the operation of the two subroutines in detail.

Lines 1–3 of `EstimateParameters` compute numerical parameters based on the tolerance ε which are used by `GenerateSamples`. The variable ‘pivot’ can be thought of as the ideal cell size we are aiming for, while as mentioned above ‘loThresh’ and ‘hiThresh’ define the allowed size range around this ideal. For simplicity of exposition, we assume that $|R_{F|S}| > \max(60, \text{hiThresh})$. If not, there are very few solutions and we can do uniform sampling by enumerating all of them as in `UniGen` [6].

Lines 4–11 of `EstimateParameters` compute ‘hashBits’, an estimate of the number of hash functions required so that the corresponding partition of $R_{F|S}$ (into 2^{hashBits} cells) has cells of the desired size. This is done along the same lines as in `UniGen`, which used an approximate model counter such as `ApproxMC` [5]. The procedure invokes a SAT solver through the function `BSAT(ϕ, m, S)`. This returns a set, consisting of models of the formula ϕ which all differ on the set of variables S , that has size m . If there is no such set of size m , the function returns a maximal set. If the estimation procedure fails, `EstimateParameters` returns \perp on line 12. In practice, it would be called repeatedly until it succeeds. Theorem 8 below shows that on average few repetitions are needed for `EstimateParameters` to succeed, and this is borne out in practice.

The second stage of `UniGen2`, `GenerateSamples`, begins on lines 1–4 by picking a hash count i close to `hashBits`, then selecting a random hash function from the family $H_{xor}(|S|, i)$ on line 2. On line 3 we pick a random output value α , so that the constraint $h(S) = \alpha$ picks out a random cell. Then, on line 5 we invoke `BSAT` on F with this additional constraint, obtaining at most `hiThresh` elements Y of the cell. If $|Y| < \text{hiThresh}$ then we have enumerated every element of $R_{F|S}$ in the cell, and if $|Y| \geq \text{loThresh}$ the cell is large enough for us to get a good sample. So if $\text{loThresh} \leq |Y| < \text{hiThresh}$, we randomly select `loThresh` elements of Y and return them on line 7.

If the number of elements of $R_{F|S}$ in the chosen cell is too large or too small, we choose a new hash count on line 4. Note that line 1 can pick an arbitrary order for the three hash counts to be tried, since our analysis of `UniGen2` does not depend on the order. This allows us to use an optimization where if we run `GenerateSamples` multiple times, we choose an order which starts with the value of i that was successful in the previous invocation of `GenerateSamples`. Since `hashBits` is only an estimate of the correct value for i , in many benchmarks on which we experimented, `UniGen2` initially failed to generate a cell of the right size with $i = \text{hashBits} - 2$, but then succeeded with $i = \text{hashBits} - 1$. In such scenarios, beginning with $i = \text{hashBits} - 1$ in subsequent iterations saves considerable time. This heuristic is similar in spirit to “leapfrogging” in `ApproxMC` [5] and `UniWit` [4], but does not compromise the theoretical guarantees of `UniGen2` in any way.

If all three hash values tried on line 4 fail to generate a correctly-sized cell, `GenerateSamples` fails and returns \perp on line 8. Theorem 1 below shows that this happens with probability at most 0.38. Otherwise, `UniGen2` completes by returning `loThresh` samples.

Parallelization of UniGen2

As described above, UniGen2 operates in two stages: `EstimateParameters` is initially called to do one-time preprocessing, and then `GenerateSamples` is called to do the actual sampling. To generate N samples, we can invoke `EstimateParameters` once, and then `GenerateSamples` $N/\text{loThresh}$ times, since each of the latter calls generates `loThresh` samples (unless it fails). Furthermore, each invocation of `GenerateSamples` is completely independent of the others. Thus if we have k processor cores, we can just perform $N/(k \cdot \text{loThresh})$ invocations of `GenerateSamples` on each. There is no need for any inter-thread communication: the “leapfrogging” heuristic for choosing the order on line 1 can simply be done on a per-thread basis. This gives us a linear speedup in the number of cores k , since the per-thread work (excluding the initial preprocessing) is proportional to $1/k$. Furthermore, Theorem 3 below shows that assuming each thread has its own source of randomness, performing multiple invocations of `GenerateSamples` in parallel does not alter its guarantees of uniformity. This means that UniGen2 can scale to an arbitrary number of processor cores as more samples are desired, while not sacrificing any theoretical guarantees.

5 Analysis

In this section, we present a theoretical analysis of the uniformity, effectiveness in discovering bugs, and runtime performance of UniGen2. For lack of space, we defer all proofs to the full version. For technical reasons, we assume that $\varepsilon > 6.84$. Our first result bounds the failure probabilities of `EstimateParameters` and `GenerateSamples`.

Theorem 1. *EstimateParameters and GenerateSamples return \perp with probabilities at most 0.009 and 0.38 respectively.*

Next we show that a single invocation of `GenerateSamples` provides guarantees nearly as strong as those of an almost-uniform generator.

Theorem 2. *For given F , S , and ε , let L be the set of samples generated using UniGen2 with a single call to `GenerateSamples`. Then for each $y \in R_{F|S}$, we have*

$$\frac{\text{loThresh}}{(1 + \varepsilon)|R_{F|S}|} \leq \Pr[y \in L] \leq 1.02 \cdot (1 + \varepsilon) \frac{\text{loThresh}}{|R_{F|S}|}.$$

Now we demonstrate that these guarantees extend to the case when `GenerateSamples` is called multiple times, sequentially or in parallel.

Theorem 3. *For given F , S , and ε , and for hashBits, loThresh, and hiThresh as estimated by `EstimateParameters`, let `GenerateSamples` be called N times with these parameters in an arbitrary parallel or sequential interleaving. Let $E_{y,i}$ denote the event that $y \in R_{F|S}$ is generated in the i^{th} call to `GenerateSamples`. Then the events $E_{y,i}$ are (l, u) -a.a.d. with $l = \frac{\text{loThresh}}{(1 + \varepsilon)|R_{F|S}|}$ and $u = \frac{1.02 \cdot (1 + \varepsilon) \text{loThresh}}{|R_{F|S}|}$.*

Next we show that the above result establishes very strong guarantees on the effectiveness of UniGen2 in discovering bugs in the CRV context. In this context, the objective of uniform generation is to maximize the probability of discovering a bug by using a diverse set of samples. Let us denote the fraction of stimuli that trigger a bug by f , i.e. if B is the set of stimuli that trigger a bug, then $f = |B|/|R_{F|S}|$. Furthermore, if N is the desired number of stimuli we wish to generate, we want to minimize the failure probability, i.e. the probability that the N randomly generated stimuli fail to intersect the set B . If the stimuli are generated uniformly, the failure probability is $(1 - f)^N$. Using binomial expansion, the failure probability can be shown to decrease exponentially in N , with decay rate of f (henceforth denoted as *failure decay rate*). We can evaluate the effectiveness of a stimulus-generation method by comparing the failure decay rate it achieves to that of a uniform generator. Alternatively, given some $\delta > 0$, we can ask how many samples are needed to ensure that the failure probability is at most δ . Normalizing by the number of samples needed by an ideal uniform generator gives the *relative number of samples needed* to find a bug. Our next theorem shows that UniGen2 is as effective as an almost-uniform generator according to both of these metrics but needs many fewer SAT calls.

Theorem 4. *Given F , S , ε , and $B \subseteq R_{F|S}$, let $f = |B|/|R_{F|S}| < 0.8$, $\nu = \frac{1}{2}(1 + \varepsilon)f$, and $\hat{\nu} = 1.02 \cdot \text{loThresh} \cdot \nu < 1$. Then we have the following bounds:*

generator type	<i>uniform</i>	UniGen	UniGen2
failure decay rate	f	$\frac{f}{1+\varepsilon}$	$(1 - \hat{\nu}) \frac{f}{1+\varepsilon}$
relative # of samples needed	1	$(1 + \nu)(1 + \varepsilon)$	$\frac{1+\hat{\nu}}{1-\hat{\nu}}(1 + \varepsilon)$
relative expected # of SAT calls	1	$\frac{3 \cdot \text{hiThresh}(1+\nu)(1+\varepsilon)}{0.52}$	$\frac{3 \cdot \text{hiThresh}}{0.62 \cdot \text{loThresh}} \frac{1+\hat{\nu}}{1-\hat{\nu}}(1 + \varepsilon)$

If $8.07 \leq \varepsilon \leq 243$ and $f \leq 1/1000$, then UniGen2 uses fewer SAT calls than UniGen on average.

Thus under reasonable conditions such as occur in industrial applications, UniGen2 is more efficient than UniGen at finding bugs. We illustrate the significance of this improvement with an example. Suppose 1 in 10^4 inputs causes a bug. Then to find a bug with probability 1/2, we would need approximately $6.93 \cdot 10^3$ uniformly generated samples. To achieve the same target, we would need approximately $1.17 \cdot 10^5$ samples from an almost-uniform generator like UniGen, and approximately $1.20 \cdot 10^5$ samples from UniGen2, using a tolerance (ε) of 16 in both cases. However, since UniGen2 picks multiple samples from each cell, it needs fewer SAT calls. In fact, the expected number of calls made by UniGen2 is only $3.38 \cdot 10^6$, compared to $4.35 \cdot 10^7$ for UniGen – an order of magnitude difference! Therefore, UniGen2 provides as strong guarantees as UniGen in terms of its ability to discover bugs in CRV, while requiring far fewer SAT calls.

Finally, since the ratio of hiThresh to loThresh can be bounded above, the average number of SAT calls per generated sample in UniGen2 can be bounded by a constant.

Theorem 5. *There exists a fixed constant $\lambda = 40$ such that for every F , S , and ε , the expected number of SAT queries made by UniGen2 per generated sample is at most λ .*

In contrast, the number of SAT calls per generated sample in UniGen is proportional to hiThresh and thus to ε^{-2} . An upper bound on the expected number of SAT queries makes it possible for UniGen2 to approach the performance of heuristic methods like random seeding of SAT solvers, which make only one SAT query per generated sample (but fail to provide any theoretical guarantees).

6 Evaluation

To evaluate the performance of UniGen2, we built a prototype implementation in C++ that employs the solver CryptoMiniSAT [1] to handle CNF-SAT augmented with XORs efficiently⁶. We conducted an extensive set of experiments on diverse public domain benchmarks, seeking to answer the following questions:

1. How does UniGen2’s runtime performance compare to that of UniGen, a state-of-the-art almost-uniform SAT sampler?
2. How does the performance of parallel UniGen2 scale with the # of cores?
3. How does the distribution of samples generated by UniGen2 compare with the ideal distribution?
4. Does parallelization affect the uniformity of the distribution of the samples?

Our experiments showed that UniGen2 outperforms UniGen by a factor of about $20\times$ in terms of runtime. The distribution generated by UniGen2 is statistically indistinguishable from that generated by an ideal uniform sampler. Finally, the runtime performance of parallel UniGen2 scales linearly with the number of cores, while its output distribution continues to remain uniform.

6.1 Experimental Setup

We conducted experiments on a heterogeneous set of benchmarks used in earlier related work [6]. The benchmarks consisted of ISCAS89 circuits augmented with parity conditions on randomly chosen subsets of outputs and next-state variables, constraints arising in bounded model checking, bit-blasted versions of SMTLib benchmarks, and problems arising from automated program synthesis. For each benchmark, the sampling set S was either taken to be the independent support of the formula or was provided by the corresponding source. Experiments were conducted on a total of 200+ benchmarks. We present results for only a subset of representative benchmarks here. A detailed list of all the benchmarks is available in the Appendix.

For purposes of comparison, we also ran experiments with UniGen [6], a state-of-the-art almost-uniform SAT witness generator. We employed the Mersenne

⁶ The tool (with source code) is available at <http://www.cs.rice.edu/CS/Verification/Projects/UniGen>

Twister to generate pseudo-random numbers, and each thread was seeded independently using the C++ class `random_device`. Both tools used an overall timeout of 20 hours, and a BSAT timeout of 2500 seconds. All experiments used $\varepsilon = 16$, corresponding to `loThresh = 11` and `hiThresh = 64`. The experiments were conducted on a high-performance computer cluster, where each node had a 12-core, 2.83 GHz Intel Xeon processor, with 4GB of main memory per core.

6.2 Results

Runtime performance

We compared the runtime performance of UniGen2 with that of UniGen for all our benchmarks. For each benchmark, we generated between 1000 and 10000 samples (depending on the size of the benchmark) and computed the average time taken to generate a sample on a single core. The results of these experiments for a representative subset of benchmarks are shown in Table 1. The columns in this table give the benchmark name, the number of variables and clauses, the size of the sampling set, the success probability of UniGen2, and finally the average runtime per sample for both UniGen2 and UniGen in seconds. The success probability of UniGen2 was computed as the fraction of calls to `GenerateSamples` that successfully generated samples.

Benchmark	#vars	#clas	S	UniGen2		UniGen
				Succ. Prob	Runtime(s)	Runtime(s)
s1238a_3_2	686	1850	32	1.0	0.3	7.17
s1196a_3_2	690	1805	32	1.0	0.23	4.54
s832a_15_7	693	2017	23	1.0	0.04	0.51
case_1_b12_2	827	2725	45	1.0	0.24	6.77
squaring16	1627	5835	72	1.0	4.16	79.12
squaring7	1628	5837	72	1.0	0.79	21.98
doublyLinkedList	6890	26918	37	1.0	0.04	1.23
LoginService2	11511	41411	36	1.0	0.05	0.55
Sort	12125	49611	52	1.0	4.15	82.8
20	15475	60994	51	1.0	19.08	270.78
enqueue	16466	58515	42	1.0	0.87	14.67
Karatsuba	19594	82417	41	1.0	5.86	80.29
lltraversal	39912	167842	23	1.0	0.18	4.86
llreverse	63797	257657	25	1.0	0.73	7.59
diagStencil_new	94607	2838579	78	1.0	3.53	60.18
tutorial3	486193	2598178	31	1.0	58.41	805.33
demo2_new	777009	3649893	45	1.0	3.47	40.33

Table 1. Runtime performance comparison of UniGen2 and UniGen (on a single core).

Table 1 clearly shows that UniGen2 significantly outperforms UniGen on all types of benchmarks, even when run on a single core. Over the entire set of 200+ benchmarks, UniGen2’s runtime performance was about 20× better than that of UniGen on average (using the geometric mean). The observed performance gain can be attributed to two factors. First, UniGen2 generates `loThresh` (11 in our experiments) samples from every cell instead of just 1 in the case of UniGen. This provides a speedup of about 10×. Second, as explained in Section 4, UniGen2 uses “leapfrogging” to optimize the order in which the values of i in line 4 of

Algorithm 2 are chosen. In contrast, UniGen uses a fixed order. This provides an additional average speedup of $2\times$ in our experiments. Note also that the success probability of UniGen2 is consistently very close to 1 across the entire set of benchmarks.

Parallel speedup

To measure the effect of parallelization on runtime performance, we ran the parallel version of UniGen2 with 1 to 12 processor cores on our benchmarks. In each experiment with C cores, we generated 2500 samples per core, and computed the C -core resource usage as the ratio of the average individual core runtime to the total number of samples (i.e. $C \times 2500$). We averaged our computations over 7 identical runs. The speedup for C cores was then computed as the ratio of 1-core resource usage to C -core resource usage. Figure 1 shows how the speedup varies with the number of cores for a subset of our benchmarks. The figure illustrates that parallel UniGen2 generally scales almost linearly with the number of processor cores.

To obtain an estimate of how close UniGen2’s performance is to real-world requirements (roughly $10\times$ slowdown compared to a simple SAT call), we measured the slowdown of UniGen2 (and UniGen) running on a single core relative to a simple SAT call on the input formula. The (geometric) mean slowdown for UniGen2 turned out to be 21 compared to 470 for UniGen. This shows that UniGen2 running in parallel on 2–4 cores comes close to matching the requirements of CRV in industrial practice.

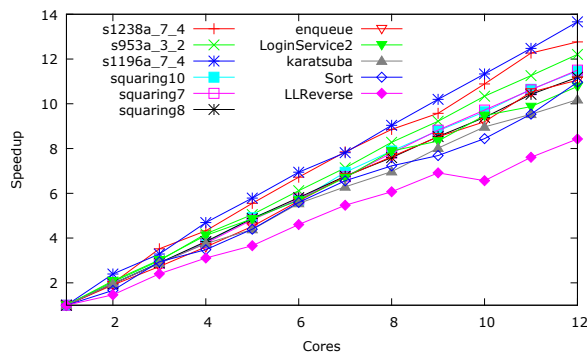


Fig. 1. Effect of parallelization on the runtime performance of UniGen2.

Uniformity comparison

To measure the quality of the distribution generated by UniGen2 and parallel UniGen2 in practice, we implemented an *ideal sampler*, henceforth denoted as IS. Given a formula F , the sampler IS first enumerates all witnesses in $R_{F|S}$, and then picks an element of $R_{F|S}$ uniformly at random. We compared the distribution generated by IS with that generated by UniGen2 run sequentially, and with that generated by UniGen2 run in parallel on 12 cores. In the last case, the samples generated by all the cores were aggregated before comparing the

distributions. We had to restrict the experiments for comparing distributions to a small subset of our benchmarks, specifically those which had less than 100,000 solutions. We generated a large number N ($\geq 4 \times 10^6$) of samples for each benchmark using each of IS, sequential UniGen2, and parallel UniGen2. Since we chose N much larger than $|R_{F|S}|$, all witnesses occurred multiple times in the list of samples. We then computed the frequency of generation of individual witnesses, and grouped witnesses appearing the same number of times together. Plotting the distribution of frequencies — that is, plotting points (x, y) to indicate that each of x distinct witnesses were generated y times — gives a convenient way to visualize the distribution of the samples. Figure 2 depicts this for one representative benchmark (case110, with 16,384 solutions). It is clear from Figure 2 that

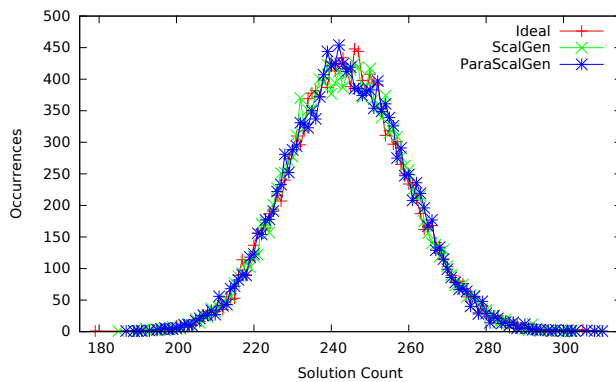


Fig. 2. Uniformity comparison between an ideal sampler (IS), UniGen2, and parallel UniGen2. Results from benchmark ‘case110’ with $N = 4 \cdot 10^6$.

the distribution generated by UniGen2 is practically indistinguishable from that of IS. Furthermore, the quality of the distribution is not affected by parallelization. Similar observations also hold for the other benchmarks for which we were able to enumerate all solutions. For the example shown in Fig. 2, the Jensen-Shannon distance between the distributions from sequential UniGen2 and IS is 0.049, while the corresponding figure for parallel UniGen2 and IS is 0.052. These small Jensen-Shannon distances make the distribution of UniGen2 (whether sequential or parallel) indistinguishable from that of IS (See Section IV(C) of [13]).

7 Conclusion

Constrained-random simulation has been the workhorse of functional verification for the past few decades. In this paper, we introduced a new algorithm, UniGen2, that outperforms state-of-the-art techniques by a factor of about $20\times$. UniGen2 trades off independence of samples for speed while still providing strong guarantees of discovering bugs with high probability. Furthermore, we showed that the parallel version of UniGen2 achieves a linear speedup with increasing number of cores. This suggests a new paradigm for constrained-random verification,

wherein we can obtain the required runtime performance through parallelization without losing guarantees of effectiveness in finding bugs.

References

1. CryptoMiniSAT. <http://www.msoos.org/cryptominisat2/>.
2. M. Bellare, O. Goldreich, and E. Petrank. Uniform generation of NP-witnesses using an NP-oracle. *Information and Computation*, 163(2):510–526, 2000.
3. L. Bening and H. Foster. *Principles of verifiable RTL design – a functional coding style supporting verification processes*. Springer, 2001.
4. S. Chakraborty, K. S. Meel, and M. Y. Vardi. A scalable and nearly uniform generator of SAT witnesses. In *Proc. of CAV*, pages 608–623, 2013.
5. S. Chakraborty, K. S. Meel, and M. Y. Vardi. A scalable approximate model counter. In *Proc. of CP*, pages 200–216, 2013.
6. S. Chakraborty, K. S. Meel, and M. Y. Vardi. Balancing scalability and uniformity in SAT-witness generator. In *Proc. of DAC*, pages 1–6, 2014.
7. R. Dechter, K. Kask, E. Bin, and R. Emek. Generating random solutions for constraint satisfaction problems. In *AAAI*, pages 15–21, 2002.
8. D. L. Eager, J. Zahorjan, and E. D. Lazowska. Speedup versus efficiency in parallel systems. *IEEE Trans. on Computers*, 38(3):408–423, 1989.
9. S. Ermon, C. P. Gomes, A. Sabharwal, and B. Selman. Embed and project: Discrete sampling with universal hashing. In *Proc. of NIPS*, 2013.
10. H. Esmailzadeh, E. Blem, R. St Amant, K. Sankaralingam, and D. Burger. Dark silicon and the end of multicore scaling. In *Proc. of ISCA*, pages 365–376. IEEE, 2011.
11. V. Gogate and R. Dechter. A new algorithm for sampling CSP solutions uniformly at random. In *Proc. of CP*, pages 711–715. Springer, 2006.
12. C. P. Gomes, A. Sabharwal, and B. Selman. Near uniform sampling of combinatorial spaces using XOR constraints. In *Proc. of NIPS*, pages 670–676, 2007.
13. I. Grosse, P. Bernaola-Galván, P. Carpena, R. Román-Roldán, J. Oliver, and H. E. Stanley. Analysis of symbolic sequences using the Jensen-Shannon divergence. *Physical Review E*, 65(4):041905, 2002.
14. M. A. Iyer. Race: A word-level ATPG-based constraints solver system for smart random simulation. In *Proc. of ITC*, pages 299–308. Citeseer, 2003.
15. M. Jerrum, L. Valiant, and V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *TCS*, 43(2-3):169–188, 1986.
16. N. Kitchen. *Markov Chain Monte Carlo Stimulus Generation for Constrained Random Simulation*. PhD thesis, University of California, Berkeley, 2010.
17. N. Kitchen and A. Kuehlmann. Stimulus generation for constrained random simulation. In *Proc. of ICCAD*, pages 258–265, 2007.
18. N. Madras. *Lectures on Monte Carlo Methods*, volume 16 of *Fields Institute Monographs*. AMS, 2002.
19. M. W. Moskewicz, C. F. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an efficient SAT solver. In *Proc. of DAC*, pages 530–535. ACM, 2001.
20. A. Nadel. Generating diverse solutions in SAT. In *Proc. of SAT*, pages 287–301. Springer, 2011.
21. R. Naveh and A. Metodi. Beyond feasibility: CP usage in constrained-random functional hardware verification. In *Proc. of CP*, pages 823–831. Springer, 2013.

22. Y. Naveh, M. Rimon, I. Jaeger, Y. Katz, M. Vinov, E. Marcus, and G. Shurek. Constraint-based random stimuli generation for hardware verification. In *Proc. of AAAI*, pages 1720–1727, 2006.
23. M. Sipser. A complexity theoretic approach to randomness. In *Proc. of STOC*, pages 330–335, 1983.
24. J. Yuan, A. Aziz, C. Pixley, and K. Albin. Simplifying Boolean constraint solving for random simulation vector generation. *TCAD*, 23(3):412–420, 2004.

APPENDIX

In this section, we first provide an extended version of Table 1 in Section A. Section B then provides detailed proofs for the theorems stated in Section 5.

A Detailed Experimental Results

Table 2: Extended Runtime performance comparison of UniGen2 and UniGen (on a single core)

Benchmark	#vars	#clas	S	UniGen2		UniGen
				Succ. Prob	Runtime(s)	Runtime(s)
109_new	60	55	36	1.0	0.14	19.35
32_new	60	49	38	1.0	0.12	22.66
70_new	62	49	40	1.0	0.13	16.34
29_new	69	55	45	1.0	0.12	6.05
case100	72	178	24	1.0	0.01	0.2
case101	72	178	24	1.0	0.01	0.2
10_new	103	135	46	1.0	0.14	2.18
case47	118	328	28	1.0	0.01	0.08
case124	133	386	31	1.0	0.12	3.43
case55	149	442	26	1.0	0.01	0.15
case8	160	525	26	1.0	0.04	0.96
lltraversal_new	163	359	41	1.0	0.19	9.73
case105	170	407	59	1.0	0.3	7.07
case5	176	518	36	1.0	0.65	5.09
treemin_new	177	451	29	1.0	0.12	2.55
s344_3_2	197	464	24	1.0	0.12	1.38
s349_3_2	198	469	24	1.0	0.12	1.46
case201	200	544	45	1.0	0.17	5.46
case202	200	544	45	1.0	0.18	5.44
case56	202	722	23	1.0	0.01	0.17
case54	203	725	23	1.0	0.01	0.17
case106	204	509	60	1.0	0.35	8.61
19_new	211	594	48	1.0	0.11	6.76
case133	211	615	42	0.98	136.04	1330.62
case136	211	615	42	0.98	128.91	1710.95
case203	214	580	49	1.0	0.13	5.22
case205	214	580	49	1.0	0.13	5.24
case204	214	580	49	1.0	0.13	5.23
tree_delete3_new	215	521	44	0.99	0.27	2.81
s344_7_4	215	540	24	1.0	0.2	1.66

Continued on next page

Benchmark	#vars	#clas	S	UniGen2		UniGen
				Succ. Prob	Runtime(s)	Runtime(s)
s349_7_4	216	545	24	1.0	0.2	1.58
case146	219	558	64	1.0	24.22	386.47
case145	219	558	64	1.0	17.49	478.73
case132	236	708	41	1.0	0.14	2.04
case135	236	708	41	1.0	0.15	2.02
case_1_b14.1	238	681	45	1.0	0.16	5.47
case_2_b14.1	238	681	45	1.0	0.15	5.28
case_3_b14.1	238	681	45	1.0	0.18	5.26
case109	241	915	31	1.0	0.03	0.39
case14	247	649	67	1.0	108.11	2675.16
s382_3_2	263	635	24	1.0	0.04	0.41
case123	267	980	34	1.0	0.4	5.22
case119	267	787	59	1.0	1.21	39.86
case_1_b14.2	270	805	43	1.0	0.17	5.38
case_2_b14.2	270	805	43	1.0	0.16	5.34
case_3_b14.2	270	805	43	1.0	0.18	5.59
case9	279	753	67	1.0	91.03	4632.83
s382_7_4	281	711	24	1.0	0.04	0.47
case61	282	753	66	1.0	103.77	1964.81
s344_15_7	284	824	24	1.0	0.08	1.27
case120	284	851	61	1.0	4.18	198.08
s349_15_7	285	829	24	1.0	0.09	1.23
case57	288	1158	32	1.0	0.85	4.88
s444_3_2	290	712	24	1.0	0.04	0.4
case121	291	975	48	1.0	0.17	5.5
case62	291	1165	33	1.0	0.21	5.59
s420_3_2	294	694	34	1.0	0.36	7.38
s420_new1_3_2	294	694	34	1.0	0.35	7.41
s420_new_3_2	294	694	34	1.0	0.24	6.2
case3	294	1110	26	1.0	0.04	0.84
case2	296	1116	26	1.0	0.05	0.83
s510_3_2	298	768	25	1.0	0.1	1.27
case126	302	1129	34	1.0	0.24	4.88
case_1_b14.3	304	941	40	1.0	0.18	5.61
case_2_b14.3	304	941	40	1.0	0.18	5.62
case_3_b14.3	304	941	40	1.0	0.18	5.76
s444_7_4	308	788	24	0.99	0.13	1.28
s420_new1_7_4	312	770	34	1.0	0.22	5.79
s420_new_7_4	312	770	34	1.0	0.17	5.81
s420_7_4	312	770	34	1.0	0.24	5.79

Continued on next page

Benchmark	#vars	#clas	S	UniGen2		UniGen
				Succ. Prob	Runtime(s)	Runtime(s)
case122	314	1258	27	1.0	0.07	1.46
s510.7_4	316	844	25	1.0	0.1	1.21
case6	329	996	52	1.0	0.2	6.96
case_0_b11.1	340	1026	48	1.0	0.21	6.06
case_1_b11.1	340	1026	48	1.0	0.23	6.16
s510.15_7	340	948	25	1.0	0.09	1.23
s382.15_7	350	995	24	1.0	0.14	1.41
s420_new.15_7	351	934	34	1.0	0.16	5.31
s526.3_2	365	943	24	1.0	0.04	0.87
s420.15_7	366	994	34	1.0	0.18	5.3
s420_new1.15_7	366	994	34	1.0	0.19	5.37
s526a.3_2	366	944	24	1.0	0.06	0.6
s444.15_7	377	1072	24	1.0	0.06	0.84
s526.7_4	383	1019	24	1.0	0.09	0.86
77_new	384	2171	44	1.0	0.12	26.92
s526a.7_4	384	1020	24	1.0	0.08	0.98
case125	393	1555	35	1.0	0.44	6.48
case35	400	1414	46	1.0	0.27	8.88
case34	409	1597	39	1.0	0.2	5.78
case143	427	1592	48	1.0	0.2	5.24
case_0_b12.1	427	1385	37	1.0	0.17	4.41
case_2_b12.1	427	1385	37	1.0	0.18	4.39
case_1_b12.1	427	1385	37	1.0	0.19	4.57
case115	428	1851	28	1.0	0.11	2.44
case114	428	1851	28	1.0	0.1	2.43
case131	432	1830	36	1.0	0.26	2.97
case116	438	1881	28	1.0	0.09	2.4
s526.15_7	452	1303	24	1.0	0.07	1.4
s526a.15_7	453	1304	24	1.0	0.07	1.37
isolateRightmost_new	483	1498	64	1.0	0.28	18.56
squaring51	496	1947	42	1.0	0.13	2.86
squaring50	500	1965	42	1.0	0.14	2.86
s953a.3_2	515	1297	45	1.0	0.67	11.6
s953a.7_4	533	1373	45	1.0	22.42	1303.7
s953a.15_7	602	1657	45	1.0	0.49	10.64
s820a.7_4	616	1703	23	1.0	0.01	0.15
s832a.7_4	624	1733	23	1.0	0.01	0.12
s820a.15_7	685	1987	23	1.0	0.01	0.11
s1238a.3_2	686	1850	32	1.0	0.3	7.17
s1196a.3_2	690	1805	32	1.0	0.23	4.54

Continued on next page

Benchmark	#vars	#clas	S	UniGen2		UniGen
				Succ. Prob	Runtime(s)	Runtime(s)
s832a_15_7	693	2017	23	1.0	0.04	0.51
squaring24	695	2193	61	1.0	0.29	6.98
squaring22	695	2193	61	1.0	0.28	6.89
squaring20	696	2198	61	1.0	0.29	6.96
squaring21	697	2203	61	1.0	0.27	6.78
s1238a_7_4	704	1926	32	1.0	0.23	3.11
s1196a_7_4	708	1881	32	1.0	0.29	3.3
squaring23	710	2268	61	1.0	0.28	7.05
GuidanceService2	715	2181	27	1.0	0.02	0.34
s1238a_15_7	773	2210	32	1.0	0.21	3.55
s1196a_15_7	777	2165	32	1.0	0.18	4.87
tree_delete3	795	2734	32	1.0	0.2	3.64
case_0_b12_2	827	2725	45	1.0	0.25	6.74
case_2_b12_2	827	2725	45	1.0	0.24	6.72
case_1_b12_2	827	2725	45	1.0	0.24	6.77
squaring27	837	2901	61	1.0	0.36	6.39
squaring25	846	2947	61	1.0	0.35	6.66
squaring3	885	2809	72	1.0	0.58	15.94
squaring2	885	2809	72	1.0	0.6	17.15
squaring6	885	2809	72	1.0	0.76	15.81
squaring5	885	2809	72	1.0	0.58	15.49
squaring1	891	2839	72	1.0	0.69	16.0
squaring4	891	2839	72	1.0	0.66	15.49
squaring26	894	3187	61	1.0	0.4	6.92
squaring11	966	3213	72	1.0	0.85	18.53
GuidanceService	988	3088	27	1.0	0.02	0.23
squaring30	1031	3693	61	1.0	0.55	15.3
squaring28	1060	3839	61	1.0	0.46	15.67
llreverse_new	1096	4217	47	1.0	0.18	10.1
squaring10	1099	3632	72	1.0	0.73	20.65
squaring8	1101	3642	72	1.0	0.76	19.64
squaring29	1141	4248	61	1.0	0.65	19.42
79_new	1217	4034	40	1.0	2.93	21.24
IssueServiceImpl	1393	4319	30	1.0	0.01	0.1
squaring9	1434	5028	72	1.0	1.03	20.35
squaring14	1458	5009	72	1.0	2.62	48.73
10	1494	2215	46	1.0	0.33	85.45
squaring12	1507	5210	72	1.0	3.25	62.44
27	1509	2707	32	1.0	0.22	6.37
squaring16	1627	5835	72	1.0	4.16	79.12

Continued on next page

Benchmark	#vars	#clas	S	UniGen2		UniGen
				Succ. Prob	Runtime(s)	Runtime(s)
squaring7	1628	5837	72	1.0	0.79	21.98
PhaseService	1686	5655	27	1.0	0.01	0.18
27_new	1792	6717	32	1.0	0.38	13.45
ActivityService	1837	5968	27	1.0	0.01	0.17
55_new	1874	8384	46	1.0	3.05	146.83
IterationService	1896	6732	27	1.0	0.01	0.23
ActivityService2	1952	6867	27	1.0	0.01	0.19
aig_insertion1	2296	9326	60	1.0	0.18	3.57
111	2348	5479	36	1.0	0.48	15.79
ConcreteActivityService	2481	9011	28	1.0	0.02	0.36
53	2586	10747	32	1.0	0.26	6.96
aig_insertion2	2592	10156	60	1.0	0.18	3.54
55	3128	12145	46	1.0	31.11	178.17
ProjectService3	3175	11019	55	1.0	0.68	17.32
NotificationServiceImpl2	3540	13425	36	1.0	0.12	1.34
109	3565	14012	36	1.0	0.88	12.99
51	3708	14594	38	1.0	0.52	18.77
32	3834	13594	38	1.0	0.47	19.39
70	4670	15864	40	1.0	0.78	24.58
ProcessBean	4768	14458	64	1.0	0.8	32.2
56	4842	17828	38	1.0	0.61	15.98
35	4915	10547	52	1.0	1.33	65.12
80	4969	17060	48	1.0	0.98	181.87
tree_delete	5758	22105	30	1.0	0.02	0.35
7	6683	24816	50	1.0	1.69	160.65
doublyLinkedList	6890	26918	37	1.0	0.04	1.23
19	6993	23867	48	1.0	3.34	52.28
LoginService	8200	26689	34	1.0	0.08	0.9
29	8866	31557	45	1.0	8.19	100.46
17	10090	27056	45	1.0	35.0	526.58
parity_new	10137	44830	50	1.0	4.09	41.08
81	10775	38006	51	1.0	15.19	285.7
LoginService2	11511	41411	36	1.0	0.05	0.55
Sort	12125	49611	52	1.0	4.15	82.8
77	14535	27573	44	1.0	11.33	38.54
20	15475	60994	51	1.0	19.08	270.78
enqueue	16466	58515	42	1.0	0.87	14.67
Karatsuba	19594	82417	41	1.0	5.86	80.29
lltraversal	39912	167842	23	1.0	0.18	4.86
LLReverse	63797	257657	25	1.0	0.73	7.59

Continued on next page

Benchmark	#vars	#clas	S	UniGen2		UniGen
				Succ. Prob	Runtime(s)	Runtime(s)
diagStencil_new	94607	2838579	78	1.0	3.53	60.18
demo4_new	381129	1801463	45	1.0	4.01	74.68
tutorial3	486193	2598178	31	1.0	58.41	805.33
demo2_new	777009	3649893	45	1.0	3.47	40.33
demo3_new	865935	3509158	45	1.0	6.36	87.12

B Detailed Proofs

In this section, we provide proofs of the various theorems stated previously. Section B.1 presents proofs of our main results, Theorems 1–5. We also include a result (Theorem 6) bounding the probability of a particular witness being generated by UniGen2, which we omitted above for lack of space. These proofs depend on a number of lemmas about `GenerateSamples` and `EstimateParameters`, presented in Sections B.2 and B.3 respectively.

B.1 Analysis of UniGen2

Theorem 1. *EstimateParameters and GenerateSamples return \perp with probabilities at most 0.009 and 0.38 respectively.*

Proof. By Lemmas 12 and 9 below respectively.

Theorem 2. *For given F , S , and ε , let L be the set of samples generated using UniGen2 with a single call to `GenerateSamples`. Then for each $y \in R_{F|S}$, we have*

$$\frac{\text{loThresh}}{(1+\varepsilon)|R_{F|S}|} \leq \Pr[y \in L] \leq 1.02 \cdot (1+\varepsilon) \frac{\text{loThresh}}{|R_{F|S}|}.$$

Proof. By Lemma 8 below.

Theorem 3. *For given F , S , and ε , and for hashBits, loThresh, and hiThresh as estimated by `EstimateParameters`, let `GenerateSamples` be called N times with these parameters in an arbitrary parallel or sequential interleaving. Let $E_{y,i}$ denote the event that $y \in R_{F|S}$ is generated in the i^{th} call to `GenerateSamples`. Then the events $E_{y,i}$ are (l, u) -a.a.d. with $l = \frac{\text{loThresh}}{(1+\varepsilon)|R_{F|S}|}$ and $u = \frac{1.02 \cdot (1+\varepsilon) \text{loThresh}}{|R_{F|S}|}$.*

Proof. Different invocations of `GenerateSamples` use independent randomness for the choices on lines 2, 3, and 7. Therefore the only part of `GenerateSamples` which can be affected by earlier invocations is the ordering heuristic used on line 1. But Lemma 8 shows that the probability that `GenerateSamples` returns a particular witness is between l and u regardless of the order used. Therefore $l \leq \Pr[E_{y,i}] \leq u$ even if conditioned on the results of previous invocations, and so the events $E_{y,i}$ are (l, u) -a.a.d..

In the following results, given F , S , ε , and $B \subseteq R_{F|S}$, we define $f = |B|/|R_{F|S}|$, $\nu = \frac{1}{2}(1+\varepsilon)f$, and $\hat{\nu} = 1.02 \cdot \text{loThresh} \cdot \nu$. We also denote the number of independent uniform samples from $R_{F|S}$ needed to hit B with probability at least δ by $N_{\text{ideal}}(|B|, \delta)$.

Theorem 4. *Given F , S , ε , B , f , ν , and $\hat{\nu}$ as above, assume $f < 0.8$ and $\hat{\nu} < 1$. Then we have the following bounds:*

generator type	<i>uniform</i>	UniGen	UniGen2
failure decay rate	f	$\frac{f}{1+\varepsilon}$	$(1-\hat{\nu})\frac{f}{1+\varepsilon}$
relative # of samples needed	1	$(1+\nu)(1+\varepsilon)$	$\frac{1+\hat{\nu}}{1-\hat{\nu}}(1+\varepsilon)$
relative expected # of SAT calls	1	$\frac{3 \cdot \text{hiThresh}(1+\nu)(1+\varepsilon)}{0.52}$	$\frac{3 \cdot \text{hiThresh}}{0.62 \cdot \text{loThresh}} \frac{1+\hat{\nu}}{1-\hat{\nu}}(1+\varepsilon)$

If $8.07 \leq \varepsilon \leq 243$ and $f \leq 1/1000$, then UniGen2 uses fewer SAT calls than UniGen on average.

Proof. The failure decay rates and relative numbers of samples needed for UniGen and UniGen2 are from Lemmas 2 and 3 below respectively.

Both UniGen and GenerateSamples make at most hiThresh SAT calls per iteration of their loops, and iterate at most 3 times (see [6] for UniGen). Since UniGen succeeds with probability at least 0.52 (by Theorem 1 of [6]), the expected number of SAT calls it makes per generated sample is at most $\frac{3 \cdot \text{hiThresh}}{0.52}$. Since GenerateSamples succeeds with probability at least 0.62 (by Theorem 1), and generates loThresh samples when it does, its expected number of SAT calls per sample is at most $\frac{3 \cdot \text{hiThresh}}{0.62 \cdot \text{loThresh}}$.

Finally, the last statement can be checked by numerical computation, noting that under the given conditions we have $\text{loThresh} \geq 6$.

Lemma 1. *Given F , S , B , $f < 0.8$, and ν as above, and $0 < \delta \leq 1$, we have*

$$N_{\text{ideal}}(|B|, \delta) \geq \frac{\ln(1/\delta)}{f(1+\nu)}.$$

Proof. First observe that for any $x \in [0, 1)$,

$$\ln(1-x) = -x - \sum_{i=2}^{\infty} \frac{x^i}{i} \geq -x - \frac{x^2}{2} \sum_{i=0}^{\infty} x^i \geq -x - \frac{x^2}{2(1-x)} = -x \left(1 + \frac{x}{2(1-x)} \right).$$

Now since $(1-f)^{N_{\text{ideal}}(|B|, \delta)} = \delta$, we have

$$N_{\text{ideal}}(|B|, \delta) = \frac{\ln \delta}{\ln(1-f)} \geq \frac{\ln(1/\delta)}{f \left(1 + \frac{f}{2(1-f)} \right)}.$$

Therefore we obtain the desired result by noting that since $\varepsilon > 6.84$, we have

$$\nu = \frac{1}{2}(1+\varepsilon)f > 3.42f \geq \frac{f}{2(1-0.8)} \geq \frac{f}{2(1-f)}.$$

Lemma 2. Given $F, S, \varepsilon, B, f < 0.8$, and ν as above, let an almost-uniform generator $\mathcal{G}^{au}(F, \varepsilon, S)$ generate N independent samples in a list L . Then

1. $\Pr[B \cap L \neq \emptyset] \geq 1 - \exp\left[-N \left(\frac{f}{1+\varepsilon}\right)\right]$.
2. For $0 < \delta \leq 1$, if $N \geq (1+\nu)(1+\varepsilon)N_{\text{ideal}}(|B|, \delta)$, then $\Pr[B \cap L \neq \emptyset] \geq 1 - \delta$.

Proof. Index the elements of L as $(y_i)_{1 \leq i \leq |L|}$. For any $b \in B$, define $p_b = \Pr[y_i = b]$, where the value is the same for all $i \in \{1, \dots, |L|\}$, and put $P = \sum_{b \in B} p_b$. By the definition of an almost-uniform generator, we have $1/((1+\varepsilon)|R_{F|S}|) \leq p_b \leq (1+\varepsilon)/|R_{F|S}|$. Therefore, $|B|/(|R_{F|S}|(1+\varepsilon)) \leq P \leq |B|(1+\varepsilon)/|R_{F|S}|$. So

$$\begin{aligned} \Pr[L \cap B \neq \emptyset] &\geq 1 - (1 - P)^N \\ &\geq 1 - \exp(-NP) \\ &\geq 1 - \exp\left(-\frac{N|B|}{(1+\varepsilon)|R_{F|S}|}\right). \end{aligned}$$

This proves the first part of the theorem.

Now if $N \geq (1+\varepsilon)(1+\nu)N_{\text{ideal}}$, using the bound on N_{ideal} from Lemma 1 and the fact that $P \geq \frac{f}{1+\varepsilon}$ gives

$$\begin{aligned} \Pr[B \cap L = \emptyset] &\leq (1 - P)^N \\ &\leq \exp[\ln(1 - P) \cdot N] \\ &\leq \exp[\ln(1 - P) \cdot (1 + \nu)(1 + \varepsilon)N_{\text{ideal}}] \\ &\leq \exp\left[\ln(1 - P) \cdot \frac{(1 + \nu)(1 + \varepsilon)}{f(1 + \nu)} \cdot \ln(1/\delta)\right] \\ &\leq \exp\left[\frac{\ln(1 - P)}{P} \cdot \ln(1/\delta)\right] \\ &\leq \exp[-1 \cdot \ln(1/\delta)] \\ &\leq \delta \end{aligned}$$

where we have used the fact that $\ln(1 - x) \leq -x$ for $x \in [0, 1)$.

Lemma 3. Given $F, S, \varepsilon, B, f < 0.8$, and $\hat{\nu} < 1$ as above, let UniGen2 generate N samples in a list L (by running `GenerateSamples` $N/\text{loThresh}$ times). Then

1. $\Pr[B \cap L \neq \emptyset] \geq 1 - \exp\left[-N(1 - \hat{\nu})\left(\frac{f}{1+\varepsilon}\right)\right]$.
2. For $0 < \delta \leq 1$, if $N \geq \frac{1+\hat{\nu}}{1-\hat{\nu}}(1+\varepsilon)N_{\text{ideal}}(|B|, \delta)$ then $\Pr[B \cap L \neq \emptyset] \geq 1 - \delta$.

Proof. Let R be the set returned by an invocation of `GenerateSamples`, and $r_1, \dots, r_{\text{loThresh}}$ be the elements of R . For any $b \in B$, define $p_b = \Pr[r_i = b]$, where the value is the same for all $i \in \{1, \dots, \text{loThresh}\}$, and $P = \sum_{b \in B} p_b$.

Now we have

$$\begin{aligned}
\Pr[R \cap B \neq \emptyset] &\geq \Pr[\{r_1, \dots, r_{\text{loThresh}}\} \cap B \neq \emptyset] \\
&\geq \text{loThresh} \cdot \Pr[r_1 \in B] - \binom{\text{loThresh}}{2} \cdot \Pr[r_1, r_2 \in B] \\
&= \text{loThresh} \cdot P - \binom{\text{loThresh}}{2} \sum_{b \in B} \sum_{b' \in B \setminus \{b\}} p_b p_{b'} \quad (\text{by pairwise independence}) \\
&= \text{loThresh} \cdot P - \binom{\text{loThresh}}{2} \sum_{b \in B} p_b (P - p_b) \\
&\geq \text{loThresh} \cdot P - \binom{\text{loThresh}}{2} \sum_{b \in B} p_b P \\
&\geq \text{loThresh} \cdot P \cdot \left(1 - \frac{\text{loThresh} \cdot P}{2}\right).
\end{aligned}$$

Now since $1/((1 + \varepsilon)|R_{F|S}|) \leq p_b \leq 1.02(1 + \varepsilon)/|R_{F|S}|$ by Theorem 3, we have $f/(1 + \varepsilon) \leq P \leq 1.02f(1 + \varepsilon)$. So

$$\Pr[R \cap B \neq \emptyset] \geq \text{loThresh} \cdot P \cdot \left(1 - \frac{\text{loThresh} \cdot 1.02f(1 + \varepsilon)}{2}\right) = (1 - \hat{\nu}) \cdot \text{loThresh} \cdot P.$$

Now by Theorem 3 this holds even when conditioned on the results of prior calls to `GenerateSamples`, so

$$\begin{aligned}
\Pr[L \cap B \neq \emptyset] &\geq 1 - (1 - (1 - \hat{\nu}) \cdot \text{loThresh} \cdot P)^{N/\text{loThresh}} \\
&\geq 1 - \exp(-(1 - \hat{\nu}) \cdot NP) \\
&\geq 1 - \exp\left[-N(1 - \hat{\nu}) \left(\frac{f}{1 + \varepsilon}\right)\right].
\end{aligned}$$

This proves the first part of the theorem. Now since $\hat{\nu} \geq \nu$, Lemma 1 shows that

$$N_{\text{ideal}}(|B|, \delta) \geq \frac{\ln(1/\delta)}{f(1 + \hat{\nu})}.$$

Therefore if $N \geq \frac{1 + \hat{\nu}}{1 - \hat{\nu}}(1 + \varepsilon)N_{\text{ideal}}$, from above we have

$$\begin{aligned}
\Pr[L \cap B = \emptyset] &\leq \exp\left[-N(1 - \hat{\nu}) \left(\frac{f}{1 + \varepsilon}\right)\right] \\
&\leq \exp\left[-\left(\frac{1 + \hat{\nu}}{1 - \hat{\nu}}\right)(1 + \varepsilon)N_{\text{ideal}} \cdot (1 - \hat{\nu}) \left(\frac{f}{1 + \varepsilon}\right)\right] \\
&= \exp(-f(1 + \hat{\nu}) \cdot N_{\text{ideal}}) \\
&\leq \exp[-1 \cdot \ln(1/\delta)] \\
&\leq \delta.
\end{aligned}$$

Theorem 5. *There exists a fixed constant $\lambda = 40$ such that for every F , S , and ε , the expected number of SAT queries made by UniGen2 per generated sample is at most λ .*

Proof. A successful invocation of `GenerateSamples` produces `loThresh` samples and makes at most $3 \cdot \text{hiThresh}$ SAT queries (at most `hiThresh` for each call to `BSAT`). Since by Lemma 1 `GenerateSamples` succeeds with probability at least 0.62, the expected number of SAT queries per generated sample is at most $(3 \cdot \text{hiThresh}) / (0.62 \cdot \text{loThresh})$. Optimization shows that $\text{hiThresh} / \text{loThresh} < 8.2$, so the expected number of queries per sample is less than 40.

Finally, we give a theorem left out of Section 5 for lack of space, bounding the probability of generating a given witness with multiple calls to `GenerateSamples`.

Theorem 6. *Given F , S , and ε as above, let UniGen2 generate N samples in a list L (by running `GenerateSamples` $N / \text{loThresh}$ times). Then for each $y \in R_{F|S}$,*

$$\frac{0.93 \cdot N}{(1 + \varepsilon)|R_{F|S}|} \leq \Pr[y \in L] \leq 1.02(1 + \varepsilon) \frac{N}{|R_{F|S}|}.$$

Proof. By Theorem 3, if R is the set returned by a single invocation of `GenerateSamples` we have

$$\frac{\text{loThresh}}{(1 + \varepsilon)|R_{F|S}|} \leq \Pr[y \in R] \leq \frac{1.02 \cdot \text{loThresh}(1 + \varepsilon)}{|R_{F|S}|}$$

regardless of the results of any prior invocations. Therefore

$$\Pr[y \in L] = 1 - \Pr[y \notin L] \geq 1 - \left(1 - \frac{\text{loThresh}}{(1 + \varepsilon)|R_{F|S}|}\right)^{N/\text{loThresh}}.$$

Now noting that

$$\frac{\text{loThresh}}{(1 + \varepsilon)|R_{F|S}|} \cdot \frac{N}{\text{loThresh}} = \frac{N}{(1 + \varepsilon)|R_{F|S}|} \leq \frac{1}{7.84},$$

applying the binomial theorem and observing that the sum of the cubic and higher order terms is positive, we have

$$\Pr[y \in L] \geq \frac{N}{(1 + \varepsilon)|R_{F|S}|} \left(1 - \frac{1}{2! \cdot 7.84}\right) = \frac{0.93 \cdot N}{(1 + \varepsilon)|R_{F|S}|}.$$

For the upper bound, a similar argument shows that

$$\Pr[y \in L] \leq 1 - \left(1 - \frac{1.02(1 + \varepsilon)\text{loThresh}}{|R_{F|S}|}\right)^{N/\text{loThresh}} \leq \frac{1.02(1 + \varepsilon)N}{|R_{F|S}|}.$$

B.2 Analysis of GenerateSamples

Throughout this section, we use the notations $R_{F|S}$ and $R_{F|S,h,\alpha}$ introduced in Section 2. We denote by U_y the event that witness $y \in R_{F|S}$ is output by GenerateSamples when called with the parameters calculated by EstimateParameters on inputs F , S , and ε . We are interested in providing lower and upper bounds for $\Pr[U_y]$. The proofs presented here follow the structure of the proofs in [6].

Let us denote $\text{round}(\log(|R_{F|S}| - 1) - \log \text{pivot})$ by m , where ‘pivot’ is the quantity computed on line 2 of EstimateParameters. The expression used for computing pivot ensures that $\text{pivot} \geq 17$. Also, as mentioned in Section 4, for simplicity we assume that $|R_{F|S}| > \max(60, \text{hiThresh})$ (in practice this can be checked by simply enumerating up to $\max(60, \text{hiThresh})$ witnesses). Finally, note that the expression for computing κ on line 1 of EstimateParameters requires $\varepsilon \geq 6.84$ in order to ensure that $\kappa \in [0, 1)$ can always be found.

The next lemma provides a lower bound on the probability of generation of a witness. Let $w_{i,y,\alpha}$ denote the probability $\Pr\left[\frac{\text{pivot}}{\sqrt{2}(1+\kappa)} \leq |R_{F|S,h,\alpha}| \leq 1 + \sqrt{2}(1+\kappa)\text{pivot} \text{ and } h(y) = \alpha : h \stackrel{R}{\leftarrow} H_{xor}(n, i)\right]$. The proof of the lemma also provides a lower bound on $w_{m,y,\alpha}$. Let $p_{i,y}$ denote the probability that GenerateSamples returns on line 7 with a particular value of i and with y in $R_{F|S,h,\alpha}$, where $\alpha \in \{0, 1\}^i$ is the value chosen on line 3. Also let $f_m = \Pr[q-2 \leq m \leq q]$, where q is shorthand for the quantity hashBits computed by EstimateParameters.

Lemma 4. *Regardless of the order chosen on line 1 of GenerateSamples, we have $\frac{\text{loThresh}}{\text{hiThresh}} \cdot f_m \cdot p_{m,y} \leq \Pr[U_y] \leq \frac{\text{loThresh}}{|Y|} \sum_{i=q-2}^q p_{i,y}$ for each $y \in R_{F|S}$.*

Proof. $\Pr[U_y] \geq \frac{\text{loThresh}}{\text{hiThresh}} \cdot \frac{1}{2^{q-m}} \cdot p_{m,y} \cdot f_m \geq \frac{\text{loThresh}}{\text{hiThresh}} \cdot \frac{1}{4} \cdot p_{m,y} \cdot f_m$. The upper bound follows from the definition of U_y .

All subsequent results in this section will bound $\Pr[U_y]$ using Lemma 4, so they also hold regardless of the order of hash counts. For notational simplicity we do not always mention this fact in the lemma statements.

Lemma 5. *For every $y \in R_{F|S}$, $\Pr[U_y] \geq \frac{0.8(1-e^{-3/2})}{4(1.05+\kappa)(|R_{F|S}|-1)}$*

Proof. From Lemma 4, we have $\Pr[U_y] \geq \frac{\text{loThresh}}{\text{hiThresh}} \cdot f_m \cdot \frac{1}{4} \cdot p_{m,y}$. Therefore, $\Pr[U_y] \geq \frac{\text{loThresh}}{1+\sqrt{2}(1+\kappa)\text{pivot}} \cdot p_{m,y} \cdot f_m$. By Lemma 14, $f_m > 0.8$. The proof is now completed by showing $p_{m,y} \geq 1/2^m(1 - \frac{1}{(\kappa/(1+\kappa))^2(|R_{F|S}|-1)})$. This gives $\Pr[U_y] \geq \frac{0.8(1-e^{-3/2})\text{loThresh}}{4(1+\sqrt{2}(1+\kappa)\text{pivot})2^m} \geq \frac{0.8(1-e^{-3/2})\text{loThresh}}{4(1.05+\kappa)(|R_{F|S}|-1)}$. The last inequality uses the observation that $1/(\sqrt{2} \cdot \text{pivot}) \leq 0.05$.

To calculate $p_{m,y}$, we first note that since $y \in R_{F|S}$, the requirement “ $y \in R_{F|S,h,\alpha}$ ” reduces to “ $y \in h^{-1}(\alpha)$ ”. For $\alpha \in \{0, 1\}^n$, we define $w_{m,y,\alpha}$ as $\Pr\left[\frac{\text{pivot}}{\sqrt{2}(1+\kappa)} \leq |R_{F|S,h,\alpha}| \leq 1 + \sqrt{2}(1+\kappa) \text{ pivot} \text{ and } h(y) = \alpha : h \stackrel{R}{\leftarrow} H_{xor}(n, m)\right]$. Therefore, $p_{m,y} = \sum_{\alpha \in \{0,1\}^m} (w_{m,y,\alpha} \cdot 2^{-m})$. The proof is now completed by showing that $w_{m,y,\alpha} \geq (1 - e^{-3/2})/2^m$ for every $\alpha \in \{0, 1\}^m$ and $y \in \{0, 1\}^n$.

Towards this end, let us first fix a random y . Now we define an indicator variable $\gamma_{z,\alpha}$ for every $z \in R_F \setminus \{y\}$ such that $\gamma_{z,\alpha} = 1$ if $h(z) = \alpha$, and $\gamma_{z,\alpha} = 0$ otherwise. Let us fix α and choose h uniformly at random from $H_{xor}(n, m)$. The random choice of h induces a probability distribution on $\gamma_{z,\alpha}$ such that $E[\gamma_{z,\alpha}] = \Pr[\gamma_{z,\alpha} = 1] = 2^{-m}$. Since we have fixed y , and since hash functions chosen from $H_{xor}(n, m, 3)$ are 3-wise independent, it follows that for every distinct $z_a, z_b \in R_F \setminus \{y\}$, the random variables $\gamma_{z_a,\alpha}, \gamma_{z_b,\alpha}$ are 2-wise independent. Let $\Gamma_\alpha = \sum_{z \in R_F \setminus \{y\}} \gamma_{z,\alpha}$ and $\mu_\alpha = E[\Gamma_\alpha]$. Clearly, $\Gamma_\alpha = |R_{F|S,h,\alpha}| - 1$ and $\mu_\alpha = \sum_{z \in R_F \setminus \{y\}} E[\gamma_{z,\alpha}] = \frac{|R_F| - 1}{2^m}$. Also, $\Pr[\frac{\text{pivot}}{\sqrt{2(1+\kappa)}} \leq |R_{F|S,h,\alpha}| \leq 1 + \sqrt{2}(1 + \kappa)\text{pivot}] = \Pr[\frac{\text{pivot}}{\sqrt{2(1+\kappa)}} - 1 \leq |R_{F|S,h,\alpha}| - 1 \leq \sqrt{2}(1 + \kappa)\text{pivot}] \geq \Pr[\frac{\text{pivot}}{\sqrt{2(1+\kappa)}} \leq |R_{F|S,h,\alpha}| - 1 \leq \sqrt{2}(1 + \kappa)\text{pivot}]$. Using the expression for pivot, we get $2 \leq \lfloor e^{-1/2}(1 + 1/\epsilon)^2 \cdot \frac{|R_F| - 1}{2^m} \rfloor$. Therefore using Chebyshev Inequality and substituting $\text{pivot} = (|R_F| - 1)/2^m$, we get $\Pr[\frac{\text{pivot}}{\sqrt{2(1+\kappa)}} \leq |R_{F|S,h,\alpha}| - 1 \leq \sqrt{2}(1 + \kappa)\text{pivot}] \geq 1 - \frac{1}{(\kappa/(1+\kappa))^2 (|R_F| - 1)/2^m}$. Therefore, $\Pr[\frac{\text{pivot}}{\sqrt{2(1+\kappa)}} \leq |R_{F|S,h,\alpha}| \leq 1 + \sqrt{2}(1 + \kappa)\text{pivot}] \geq 1 - \frac{1}{(\kappa/(1+\kappa))^2 (|R_F| - 1)/2^m}$. Since h is chosen at random from $H_{xor}(n, m)$, we also have $\Pr[h(y) = \alpha] = 1/2^m$. It follows that $w_{m,y,\alpha} \geq (1 - \frac{1}{(\kappa/(1+\kappa))^2 (|R_F| - 1)})/2^m$

The next lemma provides an upper bound on $w_{i,y,\alpha}$ and $p_{i,y}$.

Lemma 6. For $i < m - 1$, both $w_{i,y,\alpha}$ and $p_{i,y}$ are bounded above by $\frac{1}{|R_{F|S}| - 1} \frac{1}{(1 - \frac{2(1+\kappa)}{2^{m-i}})^2}$.

Proof. We will use the terminology introduced in the proof of Lemma 5. Clearly, $\mu_\alpha = \frac{|R_{F|S}| - 1}{2^i}$. Since each $\gamma_{z,\alpha}$ is a 0-1 variable, $V[\gamma_{z,\alpha}] \leq E[\gamma_{z,\alpha}]$. Therefore, $\sigma_{z,\alpha}^2 \leq \sum_{z \neq y, z \in R_{F|S}} E[\gamma_{z,\alpha}] \leq \sum_{z \in R_{F|S}} E[\gamma_{z,\alpha}] = E[\Gamma_\alpha] = 2^{-i} (|R_{F|S}| - 1)$. So $\Pr[\frac{\text{pivot}}{\sqrt{2(1+\kappa)}} \leq |R_{F|S,h,\alpha}| \leq 1 + (1 + \kappa)\sqrt{2}\text{pivot}] \leq \Pr[|R_{F|S,h,\alpha}| - 1 \leq (1 + \kappa)\sqrt{2}\text{pivot}] \leq \Pr[|R_{F|S,h,\alpha}| - 1 \leq 2(1 + \kappa)\frac{|R_{F|S}| - 1}{2^m}]$. From Chebyshev's inequality, we know that $\Pr[|\Gamma_\alpha - \mu_\alpha| \geq \lambda \sigma_{z,\alpha}] \leq 1/\lambda^2$ for every $\kappa > 0$. By choosing $\lambda = (1 - \frac{2(1+\kappa)}{2^{m-i}}) \frac{\mu_{z,\alpha}}{\sigma_{z,\alpha}}$ (Note that $\lambda > 0$ for $i < m - 1$), we have $\Pr[|R_{F|S,h,\alpha}| - 1 \leq (1 + \kappa)2\frac{|R_{F|S}| - 1}{2^m}] \leq \Pr[|(R_{F|S,h,\alpha}| - 1) - \frac{|R_{F|S}| - 1}{2^i}| \geq (1 - \frac{2(1+\kappa)}{2^{m-i}}) \frac{|R_{F|S}| - 1}{2^i}] \leq \frac{1}{(1 - \frac{2(1+\kappa)}{2^{m-i}})^2} \cdot \frac{2^i}{|R_{F|S}| - 1}$. Since h is chosen at random from $H_{xor}(n, m)$, we also have $\Pr[h(y) = \alpha] = 1/2^i$. It follows that $w_{i,y,\alpha} \leq \frac{1}{|R_{F|S}| - 1} \frac{1}{(1 - \frac{2(1+\kappa)}{2^{m-i}})^2}$. The bound for $p_{i,y}$ is easily obtained by noting that $p_{i,y} = \sum_{\alpha \in \{0,1\}^i} (w_{i,y,\alpha} \cdot 2^{-i})$.

This allows us to give an upper bound for $\Pr[U_y]$.

Lemma 7. For every $y \in R_{F|S}$, $\Pr[U_y] \leq \frac{1+\kappa}{|R_{F|S}| - 1} (7.55 + \frac{0.29}{(1-\kappa)^2})$.

Proof. We will use the terminology introduced in the proof of Lemma 5. The proof below uses the inequality $2^m \times \text{pivot} \leq \frac{|R_{F|S}| - 1}{\sqrt{2}}$ at several points. Also, From Lemma 4, we have $\Pr[U_y] \leq \sum_{i=q-2}^q \frac{\text{loThresh}}{|Y|} p_{i,y} \leq \frac{\sqrt{2}(1+\kappa)\text{loThresh}}{\text{pivot}} \sum_{i=q-2}^q p_{i,y}$.

We can sub-divide the calculation of $\Pr[U_y]$ into three cases based on the range of the values m can take.

Case 1 : $q - 2 \leq m \leq q$.

Now there are three values that m can take.

1. $m = q - 2$. We know that $p_{i,y} \leq \Pr[h(y) = \alpha] = \frac{1}{2^i}$. Therefore, $\Pr[U_y|m = q - 2] \leq \frac{\sqrt{2}(1+\kappa)\text{loThresh}}{\text{pivot}} \cdot \frac{1}{2^{q-2}} \frac{7}{4}$. Substituting the value of pivot and m , we get $\Pr[U_y|m = q - 2] \leq \frac{7(1+\kappa)\text{loThresh}}{2(|R_{F|S}|-1)}$.
2. $m = q - 1$. For $i \in [q - 2, q]$ $p_{i,y} \leq \Pr[h(y) = \alpha] = \frac{1}{2^i}$. $\Pr[U_y|m = q - 1] \leq \frac{\sqrt{2}(1+\kappa)\text{loThresh}}{\text{pivot}} \cdot \frac{1}{2^{q-2}} \frac{7}{2}$. Substituting the value of pivot and m , we get $\Pr[U_y|m = q - 1] \leq \frac{7(1+\kappa)\text{loThresh}}{|R_{F|S}|-1}$.
3. $m = q$. For $i \in [q - 1, q]$ $p_{i,y} \leq \Pr[h(y) = \alpha] = \frac{1}{2^i}$. Using Lemma 6, we get $p_{q-2,y} \leq \frac{1}{|R_{F|S}|-1} \left(\frac{1}{(1-\frac{1+\kappa}{2})^2} \right)$. Therefore, $\Pr[U_y|m = q] \leq \frac{\sqrt{2}(1+\kappa)\text{loThresh}}{\text{pivot}} \left(\frac{1}{|R_{F|S}|-1} \left(\frac{1}{(1-\frac{1+\kappa}{2})^2} + \frac{3}{2^q} \right) \right)$. Noting that pivot ≥ 17 and $\kappa \leq 1$, $\Pr[U_y|m = q] \leq \frac{(1+\kappa)\text{loThresh}}{|R_{F|S}|-1} \left(6 + \frac{0.333}{(1-\kappa)^2} \right)$.

$\Pr[U_y|q - 2 \leq m \leq q] \leq \max_i(\Pr[U_y|m = i])$. Therefore, $\Pr[U_y|q - 2 \leq m \leq q] \leq \Pr[U_y|m = q] \leq \frac{(1+\kappa)\text{loThresh}}{|R_{F|S}|-1} \left(6.667 + \frac{0.333}{(1-\kappa)^2} \right)$.

Case 2 : $m < q - 2$. $\Pr[U_y|m < q - 3] \leq \frac{\sqrt{2}(1+\kappa)}{\text{pivot}} \cdot \frac{1}{2^{q-3}} \frac{7}{4}$. Substituting the value of pivot and maximizing $m = q + 3$, we get $\Pr[U_y|m < q - 2] \leq \frac{7(1+\kappa)\text{loThresh}}{4(|R_{F|S}|-1)}$.

Case 3 : $m > q$. $\Pr[U_y|m > q] \leq \Pr[U_y|m = q + 1] = \frac{\sqrt{2}(1+\kappa)\text{loThresh}}{\text{pivot}} \left(\frac{2}{2^m} + \frac{1}{|R_{F|S}|-1} \left(\sum_{i=q-2}^{q-1} \frac{1}{1-\frac{2(1+\kappa)}{2^{m-i}}} \right) \right)$. Noting that pivot ≥ 17 and expanding the summation, $\Pr[U_y|m > q] \leq \frac{(1+\kappa)\text{loThresh}}{|R_{F|S}|-1} \left(4 + \frac{\sqrt{2}}{17} \left(\frac{1}{(1-\frac{2(1+\kappa)}{2^3})^2} + \frac{1}{(1-\frac{2(1+\kappa)}{2^2})^2} \right) \right)$.

Using $\kappa < 1$ for the first term, $\Pr[U_y|m > q] \leq \frac{(1+\kappa)\text{loThresh}}{|R_{F|S}|-1} \left(4.333 + \frac{0.333}{(1-\kappa)^2} \right)$

Summing up all the above cases, $\Pr[U_y] = \Pr[U_y|m < q - 2] \times \Pr[m < q - 2] + \Pr[U_y|q - 2 \leq m \leq q] \times \Pr[q - 2 \leq m \leq q] + \Pr[U_y|m > q] \times \Pr[m > q]$. From Lemma 13, we have $\Pr[m < q - 1] + \Pr[m > q] \leq 0.177$ and $\Pr[q - 3 \leq m \leq q] \leq 1$. Also, $\Pr[U_y|m < q - 2] \leq \Pr[U_y|m > q]$. Therefore, $\Pr[U_y|m < q - 2] \times \Pr[m < q - 2] + \Pr[U_y|m > q] \times \Pr[m > q] \leq 0.177 \times \Pr[U_y|m > q]$. Therefore, $\Pr[U_y] \leq \frac{(1+\kappa)\text{loThresh}}{|R_{F|S}|-1} \left(7.44 + \frac{0.392}{(1-\kappa)^2} \right)$.

Combining Lemmas 5 and 7, the following lemma is obtained.

Lemma 8. *Regardless of the order chosen on line 1 of GenerateSamples, for every $y \in R_{F|S}$ and $\varepsilon > 6.84$ we have*

$$\frac{\text{loThresh}}{(1 + \varepsilon)|R_{F|S}|} \leq \Pr[U_y] \leq 1.02(1 + \varepsilon) \frac{\text{loThresh}}{|R_{F|S}|}.$$

Proof. The proof is completed by using Lemmas 5 and 7 and substituting $(1 + \varepsilon) = (1 + \kappa)(7.44 + \frac{0.392}{(1-\kappa)^2})$. To arrive at the results, we use the inequality $\frac{4(1.05+\kappa)}{0.8(1-e^{-3/2})} \leq (1+\kappa)(7.44 + \frac{0.392}{(1-\kappa)^2})$. Furthermore, we use $\frac{\text{loThreshold}}{(1+\varepsilon)|R_{F|S}|} < \frac{\text{loThreshold}}{(1+\varepsilon)(|R_{F|S}|-1)}$. Also, since we assume $|R_{F|S}| - 1 \geq 60$, we have $\frac{(1+\varepsilon)\text{loThreshold}}{|R_{F|S}|-1} < \frac{1.02(1+\varepsilon)\text{loThreshold}}{|R_{F|S}|}$.

Lemma 9. *GenerateSamples succeeds (i.e. does not return \perp) with probability at least 0.62.*

Proof. As mentioned above, we are assuming $|R_{F|S}| > 1 + \sqrt{2}(1 + \kappa)\text{pivot}$. Let P_{succ} denote the probability that `GenerateSamples` succeeds. Let p_i with $q - 2 \leq i \leq q$ denote the conditional probability that the condition on line 6 of `GenerateSamples` evaluates to true with $\frac{\text{pivot}}{\sqrt{2}(1+\kappa)} \leq |R_{F|S,h,\alpha}| \leq 1 + \sqrt{2}(1 + \kappa)\text{pivot}$, given that $|R_{F|S}| > 1 + \sqrt{2}(1 + \kappa)\text{pivot}$. Let $f_m = \Pr[q - 2 \leq m \leq q]$. Therefore as shown in Lemma 4, $P_{\text{succ}} \geq p_m f_m \geq 0.8p_m$. The theorem is now proved by using Lemma ?? to show that $p_m \geq 1 - e^{-3/2} \geq 0.77$.

For every $y \in \{0, 1\}^n$ and for every $\alpha \in \{0, 1\}^m$, define an indicator variable $\nu_{y,\alpha}$ as follows: $\nu_{y,\alpha} = 1$ if $h(y) = \alpha$, and $\nu_{y,\alpha} = 0$ otherwise. Let us fix α and y and choose h uniformly at random from $H_{xor}(n, m)$. The random choice of h induces a probability distribution on $\nu_{y,\alpha}$, such that $\Pr[\nu_{y,\alpha} = 1] = \Pr[h(y) = \alpha] = 2^{-m}$ and $\mathbb{E}[\nu_{y,\alpha}] = \Pr[\nu_{y,\alpha} = 1] = 2^{-m}$. In addition 3-wise independence of hash functions chosen from $H_{xor}(n, m)$ implies that for every distinct $y_a, y_b, y_c \in R_{F|S}$, the random variables $\nu_{y_a,\alpha}, \nu_{y_b,\alpha}$ and $\nu_{y_c,\alpha}$ are 3-wise independent.

Let $\Gamma_\alpha = \sum_{y \in R_{F|S}} \nu_{y,\alpha}$ and $\mu_\alpha = \mathbb{E}[\Gamma_\alpha]$. Clearly, $\Gamma_\alpha = |R_{F|S,h,\alpha}|$ and $\mu_\alpha = \sum_{y \in R_{F|S}} \mathbb{E}[\nu_{y,\alpha}] = 2^{-m}|R_{F|S}|$. Since $|R_{F|S}| > \text{pivot}$ and $i - l > 0$, using the expression for pivot we get $3 \leq \left[e^{-1/2}(1 + \frac{1}{\kappa})^{-2} \cdot \frac{|R_{F|S}|}{2^m} \right]$. Therefore, by Lemma ??, $\Pr \left[\frac{|R_{F|S}|}{2^m} \cdot \left(1 - \frac{\kappa}{1+\kappa}\right) \leq |R_{F|S,h,\alpha}| \leq (1 + \kappa) \frac{|R_{F|S}|}{2^m} \right] > 1 - e^{-3/2}$. Simplifying and noting that $\frac{\kappa}{1+\kappa} < \kappa$ for all $\kappa > 0$, we obtain $\Pr \left[(1 + \kappa)^{-1} \cdot \frac{|R_{F|S}|}{2^m} \leq |R_{F|S,h,\alpha}| \leq (1 + \kappa) \cdot \frac{|R_{F|S}|}{2^m} \right] > 1 - e^{-3/2}$. Also, $\frac{\text{pivot}}{\sqrt{2}(1+\kappa)} \leq \frac{1}{1+\kappa} \frac{|R_{F|S}|-1}{2^m} \leq \frac{|R_{F|S}|}{(1+\kappa)2^m}$ and $1 + \sqrt{2}(1+\kappa)\text{pivot} \geq 1 + \frac{(1+\kappa)(|R_{F|S}|-1)}{2^m} \geq \frac{(1+\kappa)|R_{F|S}|}{2^m}$. Therefore, $p_m = \Pr \left[\frac{\text{pivot}}{\sqrt{2}(1+\kappa)} \leq |R_{F|S,h,\alpha}| \leq 1 + \sqrt{2}(1+\kappa)\text{pivot} \right] \geq \Pr \left[(1 + \kappa)^{-1} \cdot \frac{|R_{F|S}|}{2^m} \leq |R_{F|S,h,\alpha}| \leq (1 + \kappa) \cdot \frac{|R_{F|S}|}{2^m} \right] \geq 1 - e^{-3/2}$.

B.3 Analysis of EstimateParameters

In this section we define $\ell = \log(60) - 1$ and $\mu = \mathbb{E}[|R_{F|S,h,\alpha}|] = 2^{-i}|R_{F|S}|$. Putting $HC(x) = \text{round}(\log x + \log 1.8 - \log \text{pivot})$, we show that the value `hashBits` computed by `EstimateParameters` is a good estimate of $HC(|R_{F|S}|)$ with high probability.

The following property of pairwise independent hash functions is the main tool in our analysis.

Lemma 10. *With h and α chosen as in EstimateParameters, for each $\gamma > 0$ we have*

$$\Pr[(1 - \gamma)\mu \leq |R_{F|S,h,\alpha}| \leq (1 + \gamma)\mu] \geq 1 - \frac{1}{\gamma^2\mu}.$$

Proof. By pairwise independence, the variance of $|R_{F|S,h,\alpha}|$ is at most μ . The result then follows from Chebyshev's inequality.

Lemma 11. *Let EstimateParameters return a hashBits value of c , with i being the final value of its loop counter. Then*

$$\Pr \left[HC((1.8)^{-1} \cdot |R_{F|S}|) \leq c \leq HC(1.8 \cdot |R_{F|S}|) \mid c \neq \perp \text{ and } i + \ell \leq \log_2 |R_{F|S}| \right] \geq 0.831.$$

Proof. Since $c \neq \perp$, by line 11 of the pseudocode we have $c = HC(2^i \cdot |R_{F|S,h,\alpha}|)$, where α, i and h denote (with abuse of notation) the values of the corresponding variables in the final iteration of the loop. As mentioned above, we are assuming that $|R_{F|S}| > 60$. Since $i + \ell \leq \log_2 |R_{F|S}|$, we have $\mu = 2^{-i}|R_{F|S}| \geq 2^\ell = 30$. Applying Lemma 10 with $\gamma = 0.8/(1 + 0.8) < 0.8$, we obtain

$$\Pr[(1.8)^{-1} \cdot 2^{-i}|R_{F|S}| \leq |R_{F|S,h,\alpha}| \leq (1.8) \cdot 2^{-i}|R_{F|S}|] \geq 1 - \frac{5.0625}{\mu} \geq 0.831.$$

Lemma 12. *Given $|R_{F|S}| > 60$, the probability that EstimateParameters returns non- \perp with $i + \ell \leq \log_2 |R_{F|S}|$, is at least 0.991.*

Proof. Let us denote $\log_2 |R_{F|S}| - \ell = \log_2 |R_{F|S}| - (\lfloor \log_2(60) \rfloor - 1)$ by m . Since $|R_{F|S}| > 60$ as noted above and $|R_{F|S}| \leq 2^n$, we have $\ell < m + \ell \leq n$. Let p_i ($\ell \leq i \leq n$) denote the conditional probability that EstimateParameters terminates in iteration i of its loop with $1 \leq |R_{F|S,h,\alpha}| \leq 60$, given $|R_{F|S}| > 60$. Since the choice of h and α in each iteration of the loop are independent of those in previous iterations, the conditional probability that EstimateParameters returns non- \perp with $i \leq \log_2 |R_{F|S}| = m + \ell$, given $|R_{F|S}| > 60$, is $p_\ell + (1 - p_\ell)p_{\ell+1} + \dots + (1 - p_\ell)(1 - p_{\ell+1}) \dots (1 - p_{m+\ell-1})p_{m+\ell}$. Let us denote this sum by P . Thus, $P = p_\ell + \sum_{i=\ell+1}^{m+\ell} \prod_{k=\ell}^{i-1} (1 - p_k)p_i \geq \left(p_\ell + \sum_{i=\ell+1}^{m+\ell-1} \prod_{k=\ell}^{i-1} (1 - p_k)p_i \right) p_{m+\ell} + \prod_{s=\ell}^{m+\ell-1} (1 - p_s)p_{m+\ell} = p_{m+\ell}$. The lemma is now proved by showing that $p_{m+\ell} \geq 0.991$. Applying Lemma 10 with $\gamma = 1 - 1/30$ and $i = m = \log_2 |R_{F|S}| - \ell$, and noting that $\mu = 2^{-i}|R_{F|S}| = 2^\ell = 30$, we have $\Pr[1 \leq |R_{F|S,h,\alpha}| \leq 59] \geq 0.991$.

Now we can establish that EstimateParameters provides a good estimate of $HC(|R_{F|S}|)$.

Lemma 13. *With hashBits computed by EstimateParameters, we have*

$$\Pr [c \neq \perp \text{ and } HC((1.8)^{-1} \cdot |R_{F|S}|) \leq \text{hashBits} \leq HC((1.8) \cdot |R_{F|S}|)] > 0.823.$$

Proof. Combine Lemmas 11 and 12, getting an overall success probability of at least $0.831 \cdot 0.991 > 0.823$.

This in turn means that `hashBits` is a good estimate of the quantity m used in the analysis of `GenerateSamples`.

Lemma 14. *Let $m = \text{round}(\log(|R_{F|S}| - 1) - \log \text{pivot})$ be defined as in Section B.2. For the value `hashBits` computed by `EstimateParameters`, we have*

$$\Pr[\text{hashBits} - 2 \leq m \leq \text{hashBits}] > 0.823.$$

Proof. Straightforward computation from Lemma 13, noting that $|R_{F|S}| > 60$.