

Balancing Scalability and Uniformity in SAT Witness Generator ^{*}

Supratik Chakraborty
Indian Institute of Technology, Bombay
supratik@cse.iitb.ac.in

Kuldeep S. Meel, Moshe Y. Vardi
Rice University
kuldeep@rice.edu, vardi@cs.rice.edu

ABSTRACT

Constrained-random simulation is the predominant approach used in the industry for functional verification of complex digital designs. The effectiveness of this approach depends on two key factors: the quality of constraints used to generate test vectors, and the randomness of solutions generated from a given set of constraints. In this paper, we focus on the second problem, and present an algorithm that significantly improves the state-of-the-art of (almost-)uniform generation of solutions of large Boolean constraints. Our algorithm provides strong theoretical guarantees on the uniformity of generated solutions and scales to problems involving hundreds of thousands of variables.

1 Introduction

Functional verification constitutes one of the most challenging and time-consuming steps in the design of modern digital systems. The primary objective of functional verification is to expose design bugs early in the design cycle. Among various techniques available for this purpose, those based on simulation overwhelmingly dominate industrial practice. The state of simulation technology today is mature enough to allow simulation of large designs within reasonable time using available computational resources. Generating input patterns that exercise diverse corners of the design’s behavior space, however, remains a challenging problem [4].

In recent years, constrained-random simulation (also called constrained-random verification, or CRV) [21] has emerged as a practical approach to address the problem of simulating designs with “random enough” input patterns. In CRV, the verification engineer declaratively specifies a set of con-

straints on the values of circuit inputs. Typically, these constraints are obtained from usage requirements, environmental constraints, constraints on operating conditions and the like. A constraint solver is then used to generate random values for the circuit inputs satisfying the constraints. Since the distribution of errors in the design’s behavior space is not known *a priori*, every solution to the set of constraints is as likely to discover a bug as any other solution. It is therefore important to sample the space of all solutions uniformly or almost-uniformly (defined formally below) at random. Unfortunately, guaranteeing uniformity poses significant technical challenges when scaling to large problem sizes. This has been repeatedly noted in the literature (see, for example, [9, 22, 17]) and also confirmed by industry practitioners¹. The difficulties of generating solutions with guarantees of uniformity have even prompted researchers to propose alternative techniques for generating input patterns [9, 22]. This paper takes a step towards remedying this situation. Specifically, we describe an algorithm for generating solutions to a set of Boolean constraints, with stronger guarantees on uniformity and with higher scalability in practice than that achieved earlier.

Since constraints that arise in CRV of digital circuits are encodable as Boolean formulae, we focus on uniform generation of solutions of Boolean formulae. Henceforth, we call such solutions SAT *witnesses*. Besides its usefulness in CRV and in other applications [2, 23], uniform generation of SAT witnesses has had strong theoretical interest as well [15]. Most prior approaches to solving this problem belong to one of two categories: those that focus on strong guarantees of uniformity but scale poorly in practice (examples being [25, 3, 15]), and those that provide practical heuristics to scale to large problem instances with weak or no guarantees of uniformity (examples being [8, 17]). In [5], Chakraborty, Meel and Vardi attempted to bridge these extremes through an algorithm called UniWit. More recently, Ermon, Gomes, Sabharwal and Selman [10] proposed another sampling algorithm called PAWS. Unfortunately, both algorithms suffer from inherent limitations that make it difficult to scale them to Boolean constraints with tens of thousands of variables and beyond. In addition, the guarantees provided by these algorithms (in the context of uniform generation of SAT witnesses) are weaker than what one would desire in practice.

In this paper, we propose an algorithm called UniGen that addresses some of the deficiencies of UniWit and PAWS. This enables us to improve both the theoretical guarantees *and* practical performance vis-a-vis earlier algorithms in the con-

^{*}The full version is available as [6]

[†] This work was supported in part by NSF grants CNS 1049862 and CCF-1139011, by NSF Expeditions in Computing project “ExCAPE: Expeditions in Computer Augmented Program Engineering”, by BSF grant 9800096, by gift from Intel, by a grant from Board of Research in Nuclear Sciences, India, and 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.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

DAC ’14, June 01-05 2014, San Francisco, CA, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2370-5/14/06\$15.00.

http://dx.doi.org/10.1145/2593069.2593097.

¹Private communication: R. Kurshan

text of uniform generation of SAT witnesses. UniGen is the first algorithm to provide strong two-sided guarantees of almost-uniformity, while scaling to problems involving hundreds of thousands of variables. We also improve upon the success probability of the earlier algorithms significantly, both in theory and as evidenced by our experiments.

2 Notation and Preliminaries

Let F be 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 . A *satisfying assignment* or *witness* of F is an assignment of truth values to variables in its support such that F evaluates to true. We denote the set of all witnesses of F as R_F . Let $\mathcal{D} \subseteq X$ be a subset of the support such that there are no two satisfying assignments of F that differ only in the truth values of variables in \mathcal{D} . In other words, in every satisfying assignment of F , the truth values of variables in $X \setminus \mathcal{D}$ uniquely determine the truth value of every variable in \mathcal{D} . The set \mathcal{D} is called a *dependent* support of F , and $X \setminus \mathcal{D}$ is called an *independent* support of F . Note that there may be more than one independent supports of F . For example, $(a \vee \neg b) \wedge (\neg a \vee b)$ has three independent supports: $\{a\}$, $\{b\}$ and $\{a, b\}$. Clearly, if \mathcal{I} is an independent support of F , so is every superset of \mathcal{I} . For notational convenience, whenever the formula F is clear from the context, we omit mentioning it.

We use $\Pr[X]$ to denote the probability of event X . Given a Boolean formula F , a *probabilistic generator* of witnesses of F is a probabilistic algorithm that generates a random witness in R_F . A *uniform generator* $\mathcal{G}^u(\cdot)$ is a probabilistic generator that guarantees $\Pr[\mathcal{G}^u(F) = y] = 1/|R_F|$, for every $y \in R_F$. An *almost-uniform generator* $\mathcal{G}^{au}(\cdot, \cdot)$ ensures that for every $y \in R_F$, we have $\frac{1}{(1+\varepsilon)|R_F|} \leq \Pr[\mathcal{G}^{au}(F, \varepsilon) = y] \leq \frac{1+\varepsilon}{|R_F|}$, where $\varepsilon > 0$ is the specified *tolerance*. A *near-uniform generator* $\mathcal{G}^{nu}(\cdot)$ further relaxes the guarantee of uniformity, and ensures that $\Pr[\mathcal{G}^{nu}(F) = y] \geq c/|R_F|$ for a constant c , where $0 < c \leq 1$. Probabilistic generators are allowed to occasionally “fail” in the sense that no witness may be returned even if R_F is non-empty. The failure probability for such generators must be bounded by a constant strictly less than 1. The algorithm presented in this paper falls in the category of almost-uniform generators. An idea closely related to almost-uniform generation, and used in a key manner in our algorithm, is *approximate model counting*. Given a CNF formula F , a tolerance $\varepsilon > 0$ and a confidence $1 - \delta \in (0, 1]$, and *approximate model counter* $\text{ApproxModelCounter}(\cdot, \cdot, \cdot)$ ensures that $\Pr[\frac{|R_F|}{1+\varepsilon} \leq \text{ApproxModelCounter}(F, \varepsilon, 1 - \delta) \leq (1 + \varepsilon)|R_F|] \geq 1 - \delta$.

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 $h \stackrel{R}{\leftarrow} H(n, m, r)$ 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 $y_1, \dots, y_r \in \{0, 1\}^n$, $\Pr[\bigwedge_{i=1}^r h(y_i) = \alpha_i : h \stackrel{R}{\leftarrow} H(n, m, r)] = 2^{-mr}$. 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 $R_F \subseteq \{0, 1\}^n$ and $h \in H(n, m, r)$, we use $R_{F,h,\alpha}$ to denote the set $R_F \cap h^{-1}(\alpha)$. If we keep h fixed and

let α range over $\{0, 1\}^m$, the sets $R_{F,h,\alpha}$ form a partition of R_F .

3 Related Work

Marrying scalability with strong guarantees of uniformity has been the holy grail of algorithms that sample from solutions of constraint systems. The literature bears testimony to the significant tension between these objectives when designing random generators of SAT witnesses. Earlier work in this area either provide strong theoretical guarantees at the cost of scalability, or remedy the scalability problem at the cost of guarantees of uniformity. More recently, however, there have been efforts to bridge these two extremes.

Bellare, Goldreich and Petrank [3] showed that a provably uniform generator of SAT witnesses can be designed in theory to run in probabilistic polynomial time relative to an NP oracle. Unfortunately, it was shown in [5] that this algorithm does not scale beyond formulae with few tens of variables in practice. Weighted binary decision diagrams (BDD) have been used in [25] to sample uniformly from SAT witnesses. BDD-based techniques are, however, known to suffer from scalability problems [17]. Adapted BDD-based techniques with improved performance were proposed in [19] but the scalability was achieved at the cost of guarantees of uniformity. Random seeding of DPLL SAT solvers [20] has been shown to offer performance, although the generated distributions of witnesses can be highly skewed [17].

Markov Chain Monte Carlo methods (also called MCMC methods) [16] are widely considered to be a practical way to sample from a distribution of solutions. While MCMC methods guarantee eventual convergence to a target distribution under mild requirements, convergence is often impractically slow in practice. Heuristic adaptations were proposed in the works of [17, 18] for MCMC-based sampling in the context of constrained-random verification. Unfortunately, most of these adaptations are heuristic in nature, and do not preserve theoretical guarantees of uniformity. Sampling techniques based on interval-propagation and belief networks have been proposed in [8, 11, 14]. The simplicity of these approaches lend scalability to the techniques, but the generated distributions can deviate significantly from the uniform distribution, as shown in [18].

Sampling techniques based on hashing were originally pioneered by Sipser [24], and have been used subsequently by several researchers [3, 12, 5]. The core idea in hashing-based sampling is to use r -wise independent hash functions (for a suitable value of r) to randomly partition the space of witnesses into “small cells” of roughly equal size, and then randomly pick a solution from a randomly chosen cell. The algorithm of Bellare et al. referred to above uses this idea with n -wise independent algebraic hash functions (where n denotes the size of the support of F). Gomes, Sabharwal and Selman used 3-wise independent linear hash functions in [12] to design $\text{XORSample}'$, a near-uniform generator of SAT witnesses. Nevertheless, to realize the guarantee of near-uniformity, their algorithm requires the user to provide difficult-to-estimate input parameters. Although $\text{XORSample}'$ has been shown to scale to constraints involving a few thousand variables, Gomes et al. acknowledge the difficulty of scaling their algorithm to much larger problem sizes without sacrificing theoretical guarantees [12].

Recently, Chakraborty, Meel and Vardi [5] proposed a new hashing-based SAT witness generator, called UniWit, that

represents a small but significant step towards marrying the conflicting goals of scalability and guarantees of uniformity. Like `XORSample'`, the `UniWit` algorithm uses 3-wise independent linear hashing functions. Unlike `XORSample'`, however, the guarantee of near-uniformity of witnesses generated by `UniWit` does not depend on difficult-to-estimate input parameters. In [5], `UniWit` has been shown to scale to formulas with several thousand variables. In addition, Chakraborty et al proposed a heuristic called “leap-frogging” that allows `UniWit` to scale even further – to tens of thousands of variables [5]. Unfortunately, the guarantees of near-uniformity can no longer be established for `UniWit` with “leap-frogging”. More recently, Ermon et al. [10] proposed a hashing-based algorithm called `PAWS` for sampling from a distribution defined over a discrete set using a graphical model. While the algorithm presented in this paper has some similarities with `PAWS`, there are significant differences as well. Specifically, our algorithm provides much stronger theoretical guarantees vis-a-vis those offered by `PAWS` in the context of uniform generation of SAT witness. In addition, our algorithm scales to hundreds of thousands of variables while preserving the theoretical guarantees. `PAWS` faces the same scalability hurdles as `UniWit`, and is unlikely to scale beyond a few thousand variables without heuristic adaptations that compromise its guarantees.

4 The UniGen Algorithm

The new algorithm, called `UniGen`, falls in the category of hashing-based almost-uniform generators. `UniGen` shares some features with earlier hashing-based algorithms such as `XORSample'` [12], `UniWit` [5] and `PAWS` [10], but there are key differences that allow `UniGen` to significantly outperform these earlier algorithms, both in terms of theoretical guarantees and measured performance.

Given a CNF formula F , we use a family of 3-wise independent hash functions to randomly partition the set, R_F , of witnesses of F . Let $h : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a hash function in the family, and let y be a vector in $\{0, 1\}^n$. Let $h(y)[i]$ denote the i^{th} component of the vector obtained by applying h to y . The family of hash functions of interest is defined as $\{h(y) \mid h(y)[i] = a_{i,0} \oplus (\bigoplus_{k=1}^n a_{i,k} \cdot y[k]), a_{i,j} \in \{0, 1\}, 1 \leq i \leq m, 0 \leq j \leq n\}$, 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 the family. It has been shown in [12] that this family of hash functions is 3-wise independent. Following notation introduced in Section 2, we call this family $H_{xor}(n, m, 3)$.

While $H_{xor}(n, m, 3)$ was used earlier in `XORSample'`, `PAWS`, and (in a variant of) `UniWit`, there is a fundamental difference in the way we use it in `UniGen`. Let $X = \{x_1, x_2, \dots, x_{|X|}\}$ be the set of variables of F . Given $m > 0$, the algorithms `XORSample'`, `PAWS` and `UniWit` partition R_F by randomly choosing $h \in H_{xor}(|X|, m, 3)$ and $\alpha \in \{0, 1\}^m$, and by seeking witnesses of F conjoined with $\bigwedge_{i=1}^m (h(x_1, \dots, x_{|X|})[i] \leftrightarrow \alpha[i])$. By choosing a random $h(x_1, \dots, x_{|X|}) \in H_{xor}(|X|, m, 3)$, the set of *all* assignments to variables in X (regardless of whether they are witnesses of F) is partitioned randomly. This, in turn, ensures that the set of satisfying assignments of F is also partitioned randomly. Each conjunctive constraint of the form $(h(x_1 \dots x_{|X|})[i] \leftrightarrow \alpha[i])$ is an xor of a subset of variables of X and $\alpha[i]$, and is called an *xor-clause*. Observe that the expected number of variables in each such xor-clause is approximately $|X|/2$. It is well-known (see, for

example [13]) that the difficulty of checking satisfiability of a CNF formula with xor-clauses grows significantly with the number of variables per xor-clause. It is therefore extremely difficult to scale `XORSample'`, `PAWS` or `UniWit` to problems involving hundreds of thousands of variables. In [5], an alternative family of linear hash functions is proposed to be used with `UniWit`. Unfortunately, this also uses $|X|/2$ variables per xor-clause on average, and suffers from the same problem. In [13], a variant of $H_{xor}(|X|, m, 3)$ is used, wherein each variable in X is chosen to be in an xor-clause with a small probability $q (< 0.5)$. This mitigates the performance bottleneck significantly, but theoretical guarantees of (near-) uniformity are lost.

We address the above problem in `UniGen` by making two important observations: (i) an independent support \mathcal{I} of F is often far smaller (sometimes by a few orders of magnitude) than X , and (ii) since the value of every variable in $X \setminus \mathcal{I}$ in a satisfying assignment of F is uniquely determined by the values of variables in \mathcal{I} , the set R_F can be randomly partitioned by randomly partitioning its projection on \mathcal{I} . This motivates us to design an almost-uniform generator that accepts a subset S of the support of F as an additional input. We call S the set of *sampling variables* of F , and intend to use an independent support of F (not necessarily a minimal one) as the value of S in any invocation of the generator. Without loss of generality, let $S = \{x_1, \dots, x_{|S|}\}$, where $|S| \leq |X|$. The set R_F can now be partitioned by randomly choosing $h \in H_{xor}(|S|, m, 3)$ and $\alpha \in \{0, 1\}^m$, and by seeking solutions of $F \wedge \bigwedge_{i=1}^m (h(x_1, \dots, x_{|S|})[i] \leftrightarrow \alpha[i])$. If $|S| \ll |X|$ (as is often the case in our experience), the expected number of variables per xor-clause is significantly reduced. This makes satisfiability checking easier, and allows scaling to much larger problem sizes than otherwise possible. It is natural to ask if finding an independent support of a CNF formula F is computationally easy. While an algorithmic solution to this problem is beyond the scope of this paper, our experience indicates that a small, not necessarily minimal, independent support can often be easily determined from the source domain from which the CNF formula F is derived. For example, when a non-CNF formula G is converted to an equisatisfiable CNF formula F using Tseitin encoding, the variables introduced by the encoding form a dependent support of F .

The effectiveness of a hashing-based probabilistic generator depends on its ability to quickly partition the set R_F into “small” and “roughly equal” sized random cells. This, in turn, depends on the parameter m used in the choice of the hash function family $H(n, m, r)$. A high value of m leads to skewed distributions of sizes of cells, while a low value of m leads to cells that are not small enough. The best choice of m depends on $|R_F|$, which is not known *a priori*. Different algorithms therefore use different techniques to estimate a value of m . In `XORSample'`, this is achieved by requiring the user to provide some difficult-to-estimate input parameters. In `UniWit`, the algorithm sequentially iterates over values of m until a good enough value is found. The approach of `PAWS` comes closest to our, although there are crucial differences. In both `PAWS` and `UniGen`, an approximate model counter is first used to estimate $|R_F|$ within a specified tolerance and with a specified confidence. This estimate, along with a user-provided parameter, is then used to determine a *unique* value of m in `PAWS`. Unfortunately, this does not facilitate proving that `PAWS` is an almost-uniform genera-

tor. Instead, Ermon, et al. show that PAWS behaves like an almost-uniform generator with probability greater than $1 - \delta$, for a suitable δ that depends on difficult-to-estimate input parameters. In contrast, we use the estimate of $|R_F|$ to determine a *small range* of candidate values of m . This allows us to prove that UniGen is almost-uniform generator with confidence 1.

Algorithm 1 UniGen(F, ε, S)

```

/*Assume  $S = \{x_1, \dots, x_{|S|}\}$  is an independent support of  $F$ ,
and  $\varepsilon > 6.84$  */
1:  $(\kappa, \text{pivot}) \leftarrow \text{ComputeKappaPivot}(\varepsilon)$ ;
2:  $\text{hiThresh} \leftarrow 1 + \sqrt{2}(1 + \kappa)\text{pivot}$ ;
3:  $\text{loThresh} \leftarrow \frac{1}{\sqrt{2}(1 + \kappa)}\text{pivot}$ ;
4:  $Y \leftarrow \text{BSAT}(F, \text{hiThresh})$ ;
5: if ( $|Y| \leq \text{hiThresh}$ ) then
6:   Let  $y_1, \dots, y_{|Y|}$  be the elements of  $Y$ ;
7:   Randomly choose  $j$  from  $\{1, \dots, |Y|\}$ ; return  $y_j$ ;
8: else
9:    $C \leftarrow \text{ApproxModelCounter}(F, 0.8, 0.8)$ ;
10:   $q \leftarrow \lceil \log C + \log 1.8 - \log \text{pivot} \rceil$ ;  $i \leftarrow q - 3$ ;
11:  repeat
12:     $i \leftarrow i + 1$ ;
13:    Choose  $h$  at random from  $H_{xor}(|S|, i, 3)$ ;
14:    Choose  $\alpha$  at random from  $\{0, 1\}^i$ ;
15:     $Y \leftarrow \text{BSAT}(F \wedge (h(x_1, \dots, x_{|S|}) = \alpha), \text{hiThresh})$ ;
16:  until ( $\text{loThresh} \leq |Y| \leq \text{hiThresh}$ ) or ( $i = q$ )
17:  if ( $|Y| > \text{hiThresh}$ ) or ( $|Y| < \text{loThresh}$ ) then
18:    return  $\perp$ 
19:  else
20:    Let  $y_1, \dots, y_{|Y|}$  be the elements of  $Y$ ;
21:    Randomly choose  $j$  from  $\{1, \dots, |Y|\}$ ; return  $y_j$ ;

```

Algorithm 2 ComputeKappaPivot(ε)

```

Find  $\kappa \in [0, 1]$  such that  $\varepsilon = (1 + \kappa)(7.55 + \frac{0.29}{(1 - \kappa)^2}) - 1$ ;
pivot  $\leftarrow \lceil 4.03(1 + \frac{1}{\kappa})^2 \rceil$ ;
return  $(\kappa, \text{pivot})$ 

```

The pseudo-code for UniGen is shown in Algorithm 1. UniGen takes as inputs a Boolean CNF formula F , a tolerance ε (> 6.84 , for technical reasons explained in the [6]), and a set S of sampling variables. It either returns a random witness of F or \perp (indicating failure). The algorithm assumes access to a source of random bits, and to two sub-routines: (i) $\text{BSAT}(F, N)$, which, for every $N > 0$, returns $\min(|R_F|, N)$ distinct witnesses of F , and (ii) an approximate model counter $\text{ApproxModelCounter}(F, \varepsilon', 1 - \delta')$.

UniGen first computes two quantities, “pivot” and κ , that represent the expected size of a “small” cell and the tolerance of this size, respectively. The specific choices of expressions in ComputeKappaPivot are motivated by technical reasons explained in [6]. The values of κ and “pivot” are used to determine high and low thresholds (denoted “hiThresh” and “loThresh” respectively) for the size of each cell. Lines 5–7 handle the easy case when F has no more than “hiThresh” witnesses. Otherwise, UniGen invokes $\text{ApproxModelCounter}$ to obtain an estimate, C , of $|R_F|$ to within a tolerance of 0.8 and with a confidence of 0.8. Once again, the specific choices of the tolerance and confidence parameters used in computing C are motivated by technical reasons explained in [6]. The estimate C is then used to determine a range of candidate values for m . Specifically, this range is $\{q -$

$2, \dots, q\}$, where q is determined in line 11 of the pseudo-code. The loop in lines 11–16 checks whether some value in this range is good enough for m , i.e., whether the number of witnesses in a cell chosen randomly after partitioning R_F using $H_{xor}(|S|, m, 3)$, lies within “hiThresh” and “loThresh”. If so, lines 20–21 return a random witness from the chosen cell. Otherwise, the algorithm reports a failure in line 18.

A probabilistic generator is likely to be invoked multiple times with the same input constraint in constrained-random verification. Towards this end, lines 1–10 of the pseudo-code need to be executed only once for every formula F . Generating a new random witness requires executing afresh only lines 11–21. While this optimization appears similar to “leapfrogging” [5, 7], it is fundamentally different since it does not sacrifice any theoretical guarantees, unlike “leapfrogging”.

Implementation issues: In our implementation of UniGen, BSAT is implemented using CryptoMiniSAT [1] – a SAT solver that handles xor clauses efficiently. CryptoMiniSAT uses *blocking clauses* to prevent already generated witnesses from being generated again. Since the independent support of F determines every satisfying assignment of F , blocking clauses can be restricted to only variables in the set S . We implemented this optimization in CryptoMiniSAT, leading to significant improvements in performance. $\text{ApproxModelCounter}$ is implemented using ApproxMC [7]. Although the authors of [7] used “leapfrogging” in their experiments, we disable this optimization since it nullifies the theoretical guarantees of [7]. We use “random_device” implemented in C++ as the source of pseudo-random numbers in lines 7, 13, 14 and 21 of the pseudo-code, and also as the source of random numbers in ApproxMC .

Guarantees: The following theorem shows that UniGen is an almost-uniform generator with a high success probability.

Theorem 1. *If S is an independent support of F and if $\varepsilon > 6.84$, then for every $y \in R_F$, we have*

$$\frac{1}{(1 + \varepsilon)(|R_F| - 1)} \leq \Pr[\text{UniGen}(F, \varepsilon, S) = y] \leq (1 + \varepsilon) \frac{1}{|R_F| - 1}.$$

In addition, $\Pr[\text{UniGen}(F, \varepsilon, S) \neq \perp] \geq 0.52$.

The proof of Theorem 1 can be found in [6]. It can be shown that UniGen runs in time polynomial in ε^{-1} and in the size of F , relative to an NP-oracle.

The guarantees provided by Theorem 1 are significantly stronger than those provided by earlier generators that scale to large problem instances. Specifically, neither $\text{XORSample}'$ [12] nor UniWit [5] provide strong upper bounds for the probability of generation of a witness. PAWS [10] offers a *probabilistic* guarantee that the probability of generation of a witness lies within a tolerance factor of the uniform probability, while the guarantee of Theorem 1 is not probabilistic. The success probability of UniWit is bounded below by 0.125, which is significantly smaller than the lower bound of 0.52 guaranteed by Theorem 1.

Trading scalability with uniformity: The tolerance parameter ε provides a knob to balance scalability and uniformity in UniGen. Smaller values of ε lead to stronger guarantees of uniformity (by Theorem 1). The value of “hiThresh”, however, increases with decreasing values of ε , requiring BSAT to find more witnesses. Thus, each invocation of BSAT is likely to take longer as ε is reduced. A more detailed discussion ε is presented in the [6].

Generalization: While restriction of S to \mathcal{I} allows us to prove Theorem 1, the results generalize to arbitrary S . For an arbitrary S , Theorem 1 generalizes to sampling over projection of satisfying assignments over the set S . For example, if $F = (a \vee b)$ and $S = \{a\}$, then we represent projection of R_F over S as $R_{F|S} = \{\{0\}, \{1\}\}$. (For details see [6]).

5 Experimental Results

To evaluate the performance of UniGen, we built a prototype implementation and conducted an extensive set of experiments. Industrial constrained-random verification problem instances are typically proprietary and unavailable for published research. Therefore, we conducted experiments on CNF SAT constraints arising from several problems available in the public domain. These included bit-blasted versions of constraints arising in bounded model checking of circuits and used in [5], bit-blasted versions of SMTLib benchmarks, constraints arising from automated program synthesis, and constraints arising from ISCAS89 circuits with parity conditions on randomly chosen subsets of outputs and next-state variables.²

To facilitate running multiple experiments in parallel, we used a high-performance cluster and ran each experiment on a node of the cluster. Each node had two quad-core Intel Xeon processors with 4 GB of main memory. Recalling the terminology used in the pseudo-code of UniGen (see Section 4), we set the tolerance ε to 16, and the sampling set S to an independent support of F in all our experiments. Independent supports (not necessarily minimal ones) for all benchmarks were easily obtained from the providers of the benchmarks on request. We used 2,500 seconds as the timeout for each invocation of BSAT and 20 hours as the overall timeout for UniGen. If an invocation of BSAT timed out in line 15 of the pseudo-code of UniGen, we repeated the execution of lines 13–15 without incrementing i . With this set-up, UniGen was able to successfully generate random witnesses for formulas having up to 486,193 variables.

For performance comparisons, we also implemented and conducted experiments with UniWit – a state-of-art near-uniform generator [5]. Our choice of UniWit as a reference for comparison is motivated by two factors. First, XORSample’ is known to perform poorly vis-a-vis UniWit [5]; hence, comparing with XORSample’ is not meaningful. Second, the implementation of PAWS made available by the authors of [10] currently does not accept CNF formulae as inputs. It accepts only a graphical model of a discrete distribution as input, making a direct comparison with UniGen difficult. Since PAWS and UniWit share the same scalability problem related to large random xor-clauses, we chose to focus only on UniWit. Since the “leapfrogging” heuristic used in [5] nullifies the guarantees of UniWit, we disabled this optimization. For fairness of comparison, we used the same timeouts in UniWit as used in UniGen, i.e. 2,500 seconds for every invocation of BSAT, and 20 hours overall for every invocation of UniWit.

Table 1 presents the results of our performance-comparison experiments. Column 1 lists the CNF benchmark, and columns 2 and 3 give the count of variables and size of independent support used, respectively. The results of experiments with UniGen are presented in the next 3 columns. Column 4 gives the observed probability of success of UniGen when generating 1,000 random witnesses. Column 5 gives the average

²The tool (with source code) is available at <http://www.cs.rice.edu/CS/Verification/Projects/UniGen/>

time taken by UniGen to generate one witness (averaged over a large number of runs), while column 6 gives the average number of variables per xor-clause used for randomly partitioning R_F . The next two columns give results of our experiments with UniWit. Column 7 lists the average time taken by UniWit to generate a random witness, and column 8 gives the average number of variables per xor-clause used to partition R_F . A ‘–’ in any column means that the corresponding experiment failed to generate any witness in 20 hours.

It is clear from Table 1 that the average run-time for generating a random witness by UniWit can be two to three orders of magnitude larger than the corresponding run-time for UniGen. This is attributable to two reasons. The first stems from fewer variables in xor-clauses and blocking clauses when small independent supports are used. Benchmark “tutorial3” exemplifies this case. Here, UniWit failed to generate any witness because all calls to BSAT in UniWit, with xor-clauses and blocking clauses containing large numbers of variables, timed out. In contrast, the calls to BSAT in UniGen took much less time, due to short xor-clauses and blocking clauses using only variables from the independent support. The other reason for UniGen’s improved efficiency is that the computationally expensive step of identifying a good range of values for m (see Section 4 for details) needs to be executed only once per benchmark. Subsequently, whenever a random witness is needed, UniGen simply iterates over this narrow range of m . In contrast, generating every witness in UniWit (without leapfrogging) requires sequentially searching over all values afresh to find a good choice for m . Referring to Table 1, UniWit requires more than 20,000 seconds on average to find a good value for m and generate a random witness for benchmark “s953a_3_2”. Unlike in UniGen, there is no way to amortize this large time over multiple runs in UniWit, while preserving the guarantee of near-uniformity.

Table 1 also shows that the observed success probability of UniGen is almost always 1, much higher than what Theorem 1 guarantees and better than those from UniWit. It is clear from our experiments that UniGen can scale to problems involving almost 500K variables, while preserving guarantees of almost uniformity. This goes much beyond the reach of any other random-witness generator that gives strong guarantees on the distribution of witnesses.




Figure 1: Uniformity comparison for case110

Theorem 1 guarantees that the probability of generation of every witness lies within a specified tolerance of the uniform probability. In practice, however, the distribution of witnesses generated by UniGen is much more closer to a uniform distribution. To illustrate this, we implemented a *uniform sampler*, henceforth called US, and compared the

Table 1: Runtime performance comparison of UniGen and UniWit

Benchmark	X	S	UniGen			UniWit		
			Succ Prob	Avg Run Time (s)	Avg XOR leng	Avg Run Time (s)	Avg XOR len	Succ Prob
Squaring8	1101	72	1.0	1.77	36	5212.19	550	1.0
Squaring10	1099	72	1.0	1.83	36	4521.11	550	0.5
Squaring7	1628	72	1.0	2.44	36	2937.5	813	0.87
s1196a_7_4	708	32	1.0	6.91	16	833.1	353	0.37
s1238a_7_4	704	32	1.0	7.27	16	1570.27	352	0.35
s953a_3_2	515	45	1.0	12.51	23	22414.86	257	*
EnqueueSeqSK	16466	42	1.0	32.41	21	–	–	–
LoginService2	11511	36	0.99	6.14	18	–	–	–
LLReverse	63797	25	1.0	33.91	13	3460.58	31888	0.63
Sort	12125	52	1.0	74.59	26	–	–	–
Karatsuba	19594	41	1.0	84.03	21	–	–	–
tutorial3	486193	31	1.0	732.25	16	–	–	–

A “*” entry indicates insufficient data for estimating success probability

distributions of witnesses generated by UniGen and by US for some representative benchmarks. Given a CNF formula F , US first determines $|R_F|$ using an exact model counter (such as sharpSAT). To mimic generating a random witness, US simply generates a random number i in $\{1 \dots |R_F|\}$ using the same source of randomness as in UniGen. For every problem instance on which the comparison was done, we generated a large number $N (= 4 \times 10^6)$ of sample witnesses using each of US and UniGen. In each case, the number of times various witnesses were generated was recorded, yielding a distribution of the counts. Figure 1 shows the distributions of counts generated by UniGen and by US for one of our benchmarks (case110) with 16,384 witnesses. The horizontal axis represents counts and the vertical axis represents the number of witnesses appearing a specified number of times. Thus, the point (242, 450) represents the fact that each of 450 distinct witnesses were generated 242 times in 4×10^6 runs. Observe that the distributions resulting from UniGen and US can hardly be distinguished in practice. This holds not only for this benchmark, but for all other benchmarks we experimented with.

Additional experimental results demonstrating the trade-off between tolerance (ϵ) and performance are presented in [6]. Overall, our experiments confirm that UniGen is two to three orders of magnitude more efficient than state-of-the-art generators, has probability of success almost 1, and the distribution of generated witnesses can hardly be distinguished from that of a uniform sampler in practice.

6 Conclusion

Striking a balance between scalability and uniformity is a difficult challenge when designing random witness generators for constrained-random verification. UniGen is the first such generator for Boolean CNF formulae that scales to hundreds of thousands of variables and still preserves strong guarantees of uniformity.

Acknowledgments: We profusely thank Mate Soos for his generous suggestions and modifications in CryptoMiniSAT, without which the experimentation section would have been incomplete. We thank Ajith John for his help in experimental setup.

7 References

[1] CryptoMiniSAT. <http://www.msoos.org/cryptominisat2/>.
 [2] F. Bacchus, S. Dalmao, and T. Pitassi. Algorithms and complexity results for #SAT and Bayesian inference. In *Proc. of FOCS*, 2003.
 [3] M. Bellare, O. Goldreich, and E. Petrank. Uniform generation of NP-witnesses using an NP-oracle. *Information and*

Computation, 163(2):2000, 1998.
 [4] L. Bening and H. Foster. Principles of verifiable RTL design – a functional coding style supporting verification processes. In *Springer*, 2001.
 [5] S. Chakraborty, K. Meel, and M. Vardi. A scalable and nearly uniform generator of SAT witnesses. In *Proc. of CAV*, 2013.
 [6] S. Chakraborty, K. S. Meel, and M. Y. Vardi. Balancing scalability and uniformity in SAT-witness generator (Technical Report). <http://arxiv.org/abs/1403.6246>.
 [7] S. Chakraborty, K. S. Meel, and M. Y. Vardi. A scalable approximate model counter. In *Proc. of CP*, 2013.
 [8] R. Dechter, K. Kask, E. Bin, and R. Emek. Generating random solutions for constraint satisfaction problems. In *Proc. of AAAI*, 2002.
 [9] S. Deng, Z. Kong, J. Bian, and Y. Zhao. Self-adjusting constrained random stimulus generation using splitting evenness evaluation and xor constraints. In *Proc. of ASP-DAC*, 2009.
 [10] S. Ermon, C. P. Gomes, A. Sabharwal, and B. Selman. Embed and project: Discrete sampling with universal hashing. In *Proc. of NIPS*, 2013.
 [11] V. Gogate and R. Dechter. A new algorithm for sampling csp solutions uniformly at random. In *Proc. of CP*, 2006.
 [12] C. Gomes, A. Sabharwal, and B. Selman. Near uniform sampling of combinatorial spaces using XOR constraints. In *Proc. of NIPS*, 2007.
 [13] C. P. Gomes, J. Hoffmann, A. Sabharwal, and B. Selman. Short XORs for model counting; from theory to practice. In *Proc. of SAT*, 2007.
 [14] M. A. Iyer. Race: A word-level atpg-based constraints solver system for smart random simulation. In *ITC*, 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] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
 [17] N. Kitchen. *Markov Chain Monte Carlo Stimulus Generation for Constrained Random Simulation*. PhD thesis, University of California, Berkeley, 2010.
 [18] N. Kitchen and A. Kuehlmann. Stimulus generation for constrained random simulation. In *Proc. of ICCAD*, 2007.
 [19] J. H. Kukula and T. R. Shiple. Building circuits from relations. In *Proc. of CAV*, 2000.
 [20] M. W. Moskewicz, C. F. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an efficient sat solver. In *Proc. of DAC*, 2001.
 [21] 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 IAAI*, 2006.
 [22] S. M. Plaza, I. L. Markov, and V. Bertacco. Random stimulus generation using entropy and xor constraints. In *Proc. of DAC*, 2008.
 [23] D. Roth. On the hardness of approximate reasoning. *Artificial Intelligence*, 82(1), 1996.
 [24] M. Sipser. A complexity theoretic approach to randomness. In *Proc. of STOC*, 1983.
 [25] J. Yuan, A. Aziz, C. Pixley, and K. Albin. Simplifying boolean constraint solving for random simulation vector generation. *TCAD*, 23(3), 2004.