Distribution-Aware Sampling and Weighted Model Counting for SAT * †

Supratik Chakraborty
Indian Institute of Technology, Bombay

Daniel J. Fremont
University of California, Berkeley

Kuldeep S. Meel
Department of Computer Science, Rice University

Sanjit A. Seshia
University of California, Berkeley

Moshe Y. Vardi
Department of Computer Science, Rice University

Abstract

Given a CNF formula and a weight for each assignment of values to variables, two natural problems are weighted model counting and distribution-aware sampling of satisfying assignments. Both problems have a wide variety of important applications. Due to the inherent complexity of the exact versions of the problems, interest has focused on solving them approximately. Prior work in this area scaled only to small problems in practice, or failed to provide strong theoretical guarantees, or employed a computationally-expensive most-probable-explanation (MPE) queries that assumes prior knowledge of a factored representation of the weight distribution. We identify a novel parameter, *tilt*, which is the ratio of the maximum weight of satisfying assignment to minimum weight of satisfying assignment and present a novel approach that works with a black-box oracle for weights of assignments and requires only an NP-oracle (in practice, a SAT-solver) to solve both the counting and sampling problems when the tilt is small. Our approach provides strong theoretical guarantees, and scales to problems involving several thousand variables. We also show that the assumption of small tilt can be significantly relaxed while improving computational efficiency if a factored representation of the weights is known.

1 Introduction

Given a set of weighted elements, computing the cumulative weight of all elements that satisfy a set of constraints is a fundamental problem that arises in many contexts. Known variously as weighted model counting, discrete integration and partition function computation, this problem has applications in machine learning, probabilistic reasoning, statistics, planning, and combinatorics, among other areas (Roth 1996; Sang et al. 2004; Domshlak and Hoffmann 2007; Xue, Choi, and Darwiche 2012). A closely related problem is that of sampling elements satisfying a set of constraints, where the probability of choosing an element is proportional to its weight. The latter problem, known as distribution-aware sampling or weighted sampling, also has important applications in probabilistic reasoning, machine learning, statistical physics, constrained random verification and other domains (Jerrum and Sinclair 1996; Bacchus, Dalmao, and Pitassi 2003; Naveh et al. 2006; Madras and Piccioni 1999). Unfortunately, the exact versions of both problems are computationally hard. Weighted model counting can be used to count the number of satisfying assignments of a CNF formula; hence it is #P-hard (Valiant 1979). As shown by (Toda 1989), $\mathsf{P}^{\#\mathsf{P}}$ contains the entire polynomial-time hierarchy, and thus #P-hard problems are structurally harder than NP-complete problems. Fortunately, approximate solutions to weighted model counting and weighted sampling are adequate for most applications. As a result, there has been significant interest in designing practical approximate algorithms for these problems. Before discussing approaches to the approximate weighted sampling and counting problems, we should pause to note that a fully polynomial randomized approximation scheme (FPRAS) for weighted sampling would yield an FPRAS for #P-complete inference problems (Jerrum and Sinclair 1996; Madras and Piccioni 1999) – a possibility that lacks any evidence so far. Therefore, even approximate versions of weighted sampling and counting are computationally challenging problems with applications to a wide variety of domains.

Since constraints arising from a large class of real-world problems can be modeled as propositional CNF (henceforth CNF) formulas, we focus on CNF formulas and assume that the weights of truth assignments are given by a weight function $w(\cdot)$. Roth showed that approximately counting the models of a CNF formula is NP-hard even when the structure of the formula is severely restricted (Roth 1996). By a result of Jerrum, Valiant and Vazirani 1986, we also know that approximate model counting and almost uniform sampling (a special case of approximate weighted sampling) are polynomially inter-reducible. Therefore, it is unlikely that there exist polynomial-time algorithms for either approximate weighted model counting or approximate weighted sampling (Karp, Luby, and Madras 1989). Recently, a new class of algorithms that use pairwise-independent random parity constraints and MPE (*most-probable-explanation*)-

*The full version is available at `http://arxiv.org/abs/1404.2984`

† 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 the Board of Research in Nuclear Sciences, India, by the Data Analysis and Visualization Cyberinfrastructure funded by NSF under grant OCI-0959097, and by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

queries have been proposed for solving both problems (Ermon et al. 2013c; 2014; 2013a). These algorithms provide strong theoretical guarantees (an FPRAS relative to the MPE-oracle), and have been shown to scale to medium-sized problems in practice. While this represents a significant step in our quest for practically efficient algorithms with strong guarantees for approximate weighted model counting and approximate weighted sampling, the use of MPE-queries presents issues that need to be addressed in practice. First, MPE is an optimization problem – significantly harder in practice than a feasibility query (CNF satisfiability) (Park and Darwiche 2004). Indeed, even the approximate version of MPE has been shown to be NP-hard (Ermon et al. 2013a). Second, the use of MPE-queries along with parity constraints poses scalability hurdles (Ermon et al. 2014), and it has been argued in (Ermon et al. 2013b) that the MPE-query-based weighted model counting algorithm proposed in (Ermon et al. 2013c) is unlikely to scale well to large problem instances. This motivates us to ask if we can design approximate algorithms for weighted model counting and weighted sampling that do not invoke MPE-queries at all, and do not assume any specific representation of the weight distribution.

Our primary contributions are twofold. First, we identify a novel parameter, *tilt*, defined as the ratio of the maximum weight of a satisfying assignment to the minimum weight of a satisfying assignment, which characterizes the hardness of the approximate weighted model counting and weighted sampling problems. Second, we provide an affirmative answer to the question posed above, when the tilt is small. Specifically, we show that two recently-proposed algorithms for approximate (unweighted) model counting (Chakraborty, Meel, and Vardi 2013b) and almost-uniform (unweighted) sampling (Chakraborty, Meel, and Vardi 2014; 2013a) can be adapted to work in the setting of weighted assignments, using only a SAT-solver (NP-oracle) and a black-box weight function $w(\cdot)$ when the tilt is small. A black-box weight function is useful in applications where a factored representation is not easily available, such as in probabilistic program analysis and constrained random simulation. We also present arguments why it might be reasonable to assume a small tilt for some important classes of problems; a detailed classification of problems based on their tilt, however, is beyond the scope of this work. For distributions with large tilt, we propose an adaptation of our algorithm, which requires a pseudo-Boolean satisfiability solver instead of an (ordinary) SAT-solver as an oracle.

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. Given a set of variables $S \subseteq X$ and an assignment σ of truth values to the variables in X, we write $\sigma|_S$ to denote the projection of σ onto 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. For notational convenience, whenever the formula F is clear from the context, we omit mentioning it. Let $\mathcal{D} \subseteq X$ be a subset of the support such that there are no two satisfying assignments that differ only in the truth values of variables in $\mathcal{D}$. In other words, in every satisfying assignment, 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; for example, $(a \vee \neg b) \wedge (\neg a \vee b)$ has three, namely $\{a\}$, $\{b\}$, and $\{a, b\}$. Clearly, if $\mathcal{I}$ is an independent support of F, so is every superset of $\mathcal{I}$.

Let $w(\cdot)$ be a function that takes as input an assignment σ and yields a real number $w(\sigma) \in (0, 1]$ called the *weight* of σ. Given a set Y of assignments, we use $w(Y)$ to denote $\Sigma_{\sigma \in Y} w(\sigma)$. Our main algorithms (see Section 4) make no assumptions about the nature of the weight function, treating it as a black-box. In particular, we do not assume that the weight of an assignment can be factored into the weights of projections of the assignment on specific subsets of variables. The exception to this is Section 6, where we consider possible improvements when the weights are given by a known function, or "white-box". Three important quantities derived from the weight function are $w_{max} = \max_{\sigma \in R_F} w(\sigma)$, $w_{min} = \min_{\sigma \in R_F} w(\sigma)$, and the *tilt* $\rho = w_{max}/w_{min}$. Following standard definitions, the MPE (*most probable explanation*) is w_{max}. Thus an MPE-query for a CNF formula F and weight distribution $w(\cdot)$ seeks the value of w_{max}. Our algorithms require an upper bound on the tilt, denoted r, which is provided by the user. To maximize the efficiency of the algorithms, it is desirable to obtain as tight a bound on the tilt as possible.

We write $\mathsf{Pr}[X : \mathcal{P}]$ for the probability of outcome X when sampling from a probability space $\mathcal{P}$. For brevity, we omit $\mathcal{P}$ when it is clear from the context. The expected value of the outcome X is denoted $\mathsf{E}[X]$.

Special classes of hash functions, called *k-wise independent* hash functions, play a crucial role in our work (Bellare, Goldreich, and Petrank 1998). Let n, m and k be positive integers, and let $H(n, m, k)$ denote a family of k-wise independent hash functions mapping $\{0,1\}^n$ to $\{0,1\}^m$. We use $h \xleftarrow{R} H(n, m, k)$ to denote the probability space obtained by choosing a hash function h uniformly at random from $H(n, m, k)$. The property of k-wise independence guarantees that for all $\alpha_1, \ldots \alpha_k \in \{0,1\}^m$ and for all distinct $y_1, \ldots y_k \in \{0,1\}^n$, $\mathsf{Pr}\left[\bigwedge_{i=1}^{k} h(y_i) = \alpha_i : h \xleftarrow{R} H(n, m, k)\right] = 2^{-mk}$. For every $\alpha \in \{0,1\}^m$ and $h \in H(n, m, k)$, let $h^{-1}(\alpha)$ denote the set $\{y \in \{0,1\}^n \mid h(y) = \alpha\}$. Given $R_F \subseteq \{0,1\}^n$ and $h \in H(n, m, k)$, we use $R_{F,h,\alpha}$ to denote the set $R_F \cap h^{-1}(\alpha)$.

Our work uses an efficient family of hash functions, denoted as $H_{xor}(n, m, 3)$. Let $h : \{0,1\}^n \to \{0,1\}^m$ 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 $\{h(y) \mid h(y)[i] = a_{i,0} \oplus (\bigoplus_{l=1}^{n} a_{i,l} \cdot y[l]), a_{i,j} \in \{0,1\}, 1 \leq i \leq m, 0 \leq j \leq n\}$, 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 (Gomes, Sabharwal,

and Selman 2007) that this family of hash functions is 3-wise independent. For every $m \in \{1, \ldots |S| - 1\}$, 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 $i \in \{1, \ldots m\}$. 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 $i \in \{1, \ldots m\}$.

Given a CNF formula F, an *exact weighted model counter* returns $w(R_F)$. An *approximate weighted model counter* relaxes this requirement to some extent: given a *tolerance* $\varepsilon > 0$ and *confidence* $1 - \delta \in (0, 1]$, the value v returned by the counter satisfies $\mathsf{Pr}[\frac{w(R_F)}{1+\varepsilon} \leq v \leq (1+\varepsilon)w(R_F)] \geq 1 - \delta$. A related kind of algorithm is a *weighted-uniform probabilistic generator*, which outputs a witness $w \in R_F$ such that $\mathsf{Pr}[w = y] = w(y)/w(R_F)$ for every $y \in R_F$. An *almost weighted-uniform generator* relaxes this requirement, ensuring that for all $y \in R_F$, we have $\frac{w(y)}{(1+\varepsilon)w(R_F)} \leq \mathsf{Pr}[w = y] \leq \frac{(1+\varepsilon)w(y)}{w(R_F)}$. Probabilistic generators are allowed to occasionally "fail" by not returning a witness (when R_F is non-empty), with the failure probability upper bounded by δ.

3 Related Work

Marrying strong theoretical guarantees with scalable performance is the holy grail of research in the closely related areas of weighted model counting and weighted sampling. The tension between the two objectives is evident from a survey of the literature. Earlier algorithms for weighted model counting can be broadly divided into three categories: those that give strong guarantees but scale poorly in practice, those that give weak guarantees but scale well in practice, and some recent attempts to bridge this gap. Techniques in the first category attempt to compute the weighted model count exactly by enumerating partial solutions (Sang, Beame, and Kautz 2005) or by converting the CNF formula to alternative representations (Darwiche 2004; Choi and Darwiche 2013). Unfortunately, none of these approaches scale to large problem instances. Techniques in the second category employ variational methods, sampling-based methods or other heuristic methods. Variational methods (Wainwright and Jordan 2008; Gogate and Dechter 2011) work extremely well in practice, but do not provide guarantees except in very special cases. Sampling-based methods are usually based on importance sampling (e.g. (Gogate and Dechter 2011)), which provide weak one-sided bounds, or on Markov Chain Monte Carlo (MCMC) sampling (Jerrum and Sinclair 1996; Madras 2002). MCMC sampling is perhaps the most popular technique for both weighted sampling and weighted model counting. Several MCMC algorithms like simulated annealing and the Metropolis-Hastings algorithm have been studied extensively in the literature (Kirkpatrick, Gelatt, and Vecchi 1983; Madras 2002). While MCMC sampling is guaranteed to converge to a target distribution under mild requirements, convergence is often impractically slow (Jerrum and Sinclair 1996). Therefore, practical MCMC sampling-based tools use heuristics that destroy the theoretical guarantees. Several other heuristic techniques that provide weak one-sided bounds have also been proposed in the literature (Gomes, Sabharwal, and Selman 2006).

Recently, Ermon et al. proposed new hashing-based algorithms for approximate weighted model counting and approximate weighted sampling (2013c; 2013a; 2013b; 2014). Their algorithms use random parity constraints as pairwise independent hash functions to partition the set of satisfying assignments of a CNF formula into cells. An oracle is then queried to obtain the maximum weight of an assignment in a randomly chosen cell. By repeating the MPE-queries polynomially many times for randomly chosen cells of appropriate expected sizes, Ermon et al. showed that they can provably compute approximate weighted model counts and also provably achieve approximate weighted sampling. The performance of Ermon et al's algorithms depend crucially on the ability to efficiently answer MPE-queries. Complexity-wise, MPE is an optimization problem and is believed to be much harder than a feasibility query (CNF satisfiability). Furthermore, even the approximation of MPE is known to be NP-hard (Ermon et al. 2013a). The problem is further compounded by the fact that the MPE-queries generated by Ermon et al's algorithms have random parity constraints built into them. Existing MPE-solving techniques work efficiently when the weight distribution of assignments is specified by a graphical model, and the underlying graph has specific structural properties (Park and Darwiche 2004). With random parity constraints, these structural properties are likely to be violated very often. In (Ermon et al. 2013b), it has been argued that an MPE-query-based weighted model-counting algorithm proposed in (Ermon et al. 2013c) is unlikely to scale well to large problem instances. Since MPE-solving is also crucial in the weighted sampling algorithm of (Ermon et al. 2013a), the same criticism applies to that algorithm as well. Several relaxations of the MPE-query-based algorithm proposed in (Ermon et al. 2013c), were therefore discussed in (Ermon et al. 2013b). While these relaxations help reduce the burden of MPE-solving, they also significantly weaken the theoretical guarantees.

In later work, Ermon et al. (2014) showed how the average size of parity constraints in their weighted model counting and weighted sampling algorithms can be reduced using a new class of hash functions. This work, however, still stays within the same paradigm as their earlier work – i.e., it uses MPE-queries and XOR constraints. Although Ermon et al's algorithms provide a 16-factor approximation in theory, in actual experiments, they use relaxations and timeouts of the MPE-solver to get upper and lower bounds on the optimal MPE solution. Unfortunately, these bounds do not come with any guarantees on the factor of approximation. Running the MPE-solver to obtain the optimal value is likely to take significantly longer, and is not attempted in Ermon et al's work. Furthermore, to get an approximation factor less than 16 using Ermon et al's algorithms requires replication of variables. This can significantly increase the running time of their algorithms. In contrast, the performance of our algorithms is much less sensitive to the approximation factor, and our experiments routinely compute 1.75-approximations of weighted model counts.

The algorithms developed in this paper are closely related

to two algorithms proposed recently by Chakraborty, Meel and Vardi (2013b; 2014) The first of these (Chakraborty, Meel, and Vardi 2013b) computes the approximate (unweighted) model-count of a CNF formula, while the second algorithm (Chakraborty, Meel, and Vardi 2014) performs near-uniform (unweighted) sampling. Like Ermon et al's algorithms, these algorithms make use of parity constraints as pairwise independent hash functions, and can benefit from the new class of hash functions proposed in (Ermon et al. 2014). Unlike Ermon et al's algorithms, however, Chakraborty et al. use a SAT-solver (NP-oracle) specifically engineered to handle parity constraints efficiently. This allows Chakraborty, Meel, and Vardi's algorithms to scale to much larger problems than those analyzed by Ermon et al., albeit in the unweighted setting. As we show later, this scalability is inherited by our algorithms in the weighted setting as well.

Algorithm 1 WeightMC$(F, \varepsilon, \delta, S, r)$

counter $\leftarrow 0$; $C \leftarrow$ emptyList; $\mathrm{w}_{\max} \leftarrow 1$;
pivot $\leftarrow 2 \times \lceil e^{3/2}\left(1+\frac{1}{\varepsilon}\right)^2 \rceil$;
$t \leftarrow \lceil 35 \log_2(3/\delta) \rceil$;
repeat
  $(c, \mathrm{w}_{\max}) \leftarrow$ WeightMCCore$(F, S, \text{pivot}, r, \mathrm{w}_{\max})$;
  counter $\leftarrow$ counter $+ 1$;
  if $c \neq \perp$ **then**
    AddToList$(C, c \cdot \mathrm{w}_{\max})$;
until (counter $< t$)
finalCount $\leftarrow$ FindMedian(C);
return finalCount;

Algorithm 2 WeightMCCore$(F, S, \text{pivot}, r, \mathrm{w}_{\max})$

$(Y, \mathrm{w}_{\max}) \leftarrow$
  BoundedWeightSAT$(F, \text{pivot}, r, \mathrm{w}_{\max}, S)$;
if $(w(Y)/\mathrm{w}_{\max} \leq \text{pivot})$ **then**
  return $w(Y)$;
else
  $i \leftarrow 0$;
  repeat
    $i \leftarrow i + 1$;
    Choose h at random from $H_{xor}(|S|, i, 3)$;
    Choose α at random from $\{0,1\}^i$;
    $(Y, \mathrm{w}_{\max}) \leftarrow$ BoundedWeightSAT$(F \wedge (h(x_1, \ldots x_{|S|}) = \alpha), \text{pivot}, r, \mathrm{w}_{\max}, S)$;
  until $((0 < w(Y)/\mathrm{w}_{\max} \leq \text{pivot})$ or $(i = n))$
  if $((w(Y)/\mathrm{w}_{\max} > \text{pivot})$ **or** $(w(Y) = 0))$ **then return** $(\perp, \mathrm{w}_{\max})$;
  else return $(\frac{w(Y) \cdot 2^{i-1}}{\mathrm{w}_{\max}}, \mathrm{w}_{\max})$;

Algorithm 3 BoundedWeightSAT$(F, \text{pivot}, r, \mathrm{w}_{\max}, S)$

$\mathrm{w}_{\min} \leftarrow \mathrm{w}_{\max}/r$; $\mathrm{w}_{\text{total}} \leftarrow 0$; $Y = \{\}$;
repeat
  $y \leftarrow$ SolveSAT(F);
  if $(y =$ UNSAT$)$ **then**
    break;
  $Y = Y \cup y$;
  $F =$ AddBlockClause$(F, y|_S)$;
  $\mathrm{w}_{\text{total}} \leftarrow \mathrm{w}_{\text{total}} + w(y)$;
  $\mathrm{w}_{\min} \leftarrow min(\mathrm{w}_{\min}, w(y))$;
until $(\mathrm{w}_{\text{total}}/(\mathrm{w}_{\min} \cdot r) > \text{pivot})$;
return $(Y, \mathrm{w}_{\min} \cdot r)$;

4 Algorithms

We now present algorithms for approximate weighted model counting and approximate weighted sampling, assuming a small bounded tilt and a black-box weight function. Recalling that the tilt concerns weights of only satisfying assignments, our assumption about it being bounded by a small number is reasonable in several practical situations. For example, when performing probabilistic inference with evidence by reduction to weighted model counting (Chavira and Darwiche 2008), every satisfying assignment of the CNF formula corresponds to an assignment of values to variables in the underlying probabilistic graphical model that is consistent with the evidence. Furthermore, the weight of a satisfying assignment is the joint probability of the corresponding assignment of variables in the probabilistic graphical model. A large tilt would therefore mean existence of two assignments that are consistent with the same evidence, but of which one is overwhelmingly more likely than the other. In several real-world situations, this is considered unlikely given that numerical conditional probability values are often obtained from human experts providing qualitative and rough quantitative data (see, e.g. Sec 8.3 of (Dıez and Druzdzel 2006)).

The algorithms presented in this section require an upper bound for the tilt ρ as part of their input. It is worth noting that although tighter upper bounds improve performance, the algorithms are sound with respect to any upper bound estimate. While an algorithmic solution to the estimation of upper bounds for ρ is beyond the scope of this work, such an estimate can often be easily obtained from the designers of probabilistic models. It is often easier for designers to estimate an upper bound for ρ than to accurately estimate w_{max}, since the former does not require precise knowledge of the probabilities of all models.

Our weighted model counting algorithm, called WeightMC, is best viewed as an adaptation of the ApproxMC algorithm proposed by Chakraborty, Meel and Vardi (2013b) for approximate unweighted model counting. Similarly, our weighted sampling algorithm, called WeightGen, can be viewed as an adaptation of the the UniGen algorithm (Chakraborty, Meel, and Vardi 2014), originally proposed for near-uniform unweighted sampling. The key idea in both ApproxMC and UniGen is to partition the set of satisfying assignments into "cells" containing roughly equal numbers of satisfying assignments, using a random hash function from the family $H_{xor}(n, m, 3)$. A random cell is then chosen and inspected to see if the number of satisfying assignments in it is smaller than a pre-computed threshold. This threshold, in turn, depends on

Algorithm 4 WeightGen(F, ε, r, S)

/*Assume $\varepsilon > 6.84$ */

1: $\mathrm{w}_{\max} \leftarrow 1$; Samples $= \{\}$;
2: $(\kappa, \mathrm{pivot}) \leftarrow$ ComputeKappaPivot(ε);
3: $\mathrm{hiThresh} \leftarrow 1 + \sqrt{2}(1+\kappa)\mathrm{pivot}$;
4: $\mathrm{loThresh} \leftarrow \frac{1}{\sqrt{2}(1+\kappa)}\mathrm{pivot}$;
5: $(Y, \mathrm{w}_{\max}) \leftarrow$ BoundedWeightSAT$(F, \mathrm{hiThresh}, r, \mathrm{w}_{\max}, S)$;
6: **if** $(w(Y)/\mathrm{w}_{\max} \leq \mathrm{hiThresh})$ **then**
7: Choose y weighted-uniformly at random from Y;
8: **return** y;
9: **else**
10: $(C, \mathrm{w}_{\max}) \leftarrow$ WeightMC$(F, 0.8, 0.2)$;
11: $q \leftarrow \lceil \log C - \log \mathrm{w}_{\max} + \log 1.8 - \log \mathrm{pivot} \rceil$;
12: Choose h at random from $H_{xor}(|S|, n, 3)$;
13: Choose α at random from $\{0,1\}^n$;
14: $i \leftarrow q - 4$;
15: **repeat**
16: $i \leftarrow i + 1$;
17: $(Y, \mathrm{w}_{\max}) \leftarrow$ BoundedWeightSAT$(F \wedge (h_i(x_1, \ldots x_{|S|}) = \alpha_i), \mathrm{hiThresh}, r, \mathrm{w}_{\max}, S)$;
18: $W \leftarrow w(Y)/\mathrm{w}_{\max}$
19: **until** $(\mathrm{loThresh} \leq W \leq \mathrm{hiThresh})$ or $(i = q)$
20: **if** $((W > \mathrm{hiThresh})$ or $(W < \mathrm{loThresh}))$ **then return** $\perp$
21: **else** Choose y weighted-uniformly at random from Y; **return** y;

Algorithm 5 ComputeKappaPivot(ε)

1: Find $\kappa \in [0, 1)$ such that $\varepsilon = (1+\kappa)(7.55 + \frac{0.29}{(1-\kappa)^2}) - 1$
2: $\mathrm{pivot} \leftarrow \lceil 4.03 \left(1 + \frac{1}{\kappa}\right)^2 \rceil$; **return** (κ, pivot)

the desired approximation factor or tolerance ε. If the chosen cell is small enough, UniGen samples uniformly from the chosen small cell to obtain a near-uniformly generated satisfying assignment. ApproxMC multiplies the number of satisfying assignments in the cell by a suitable scaling factor to obtain an estimate of the model count. ApproxMC is then repeated a number of times (depending on the desired confidence $1 - \delta$) and the median of the computed counts determined to obtain the final approximate model count. For weighted model counting and sampling, the primary modification that needs to be made to ApproxMC and UniGen is that instead of requiring "cells" to have roughly equal numbers of satisfying assignments, we now require them to have roughly equal weights of satisfying assignments.

A randomly chosen hash function from $H_{xor}(n, m, 3)$ consists of m XOR constraints, each of which has expected size $n/2$. Although ApproxMC and UniGen were shown to scale to a few thousand variables, their performance erodes rapidly beyond that point. It has recently been shown in (Chakraborty, Meel, and Vardi 2014) that by using random parity constraints on the independent support of a formula (which can be orders of magnitude smaller than the complete support), we can significantly reduce the size of XOR constraints. We use this idea in our work. For all our benchmark problems, obtaining the independent support of the CNF formulae has been easy, once we examine the domain from which the problem originated.

Both WeightMC and WeightGen assume access to a subroutine called BoundedWeightSAT that takes as inputs a CNF formula F, a "pivot", an upper bound r on the tilt and an upper bound $\mathrm{w}_{\max}$ on the maximum weight of a satisfying assignment in the independent support set S. BoundedWeightSAT returns a set of satisfying assignments of F such that the total weight of the returned assignments scaled by $1/\mathrm{w}_{\max}$ exceeds the "pivot". Since all weights are assumed to be in $(0, 1]$, the upper bound of $\mathrm{w}_{\max}$ is set to 1 in the initial invocation of BoundedWeightSAT. Subsequently, BoundedWeightSAT returns a refined upper bound of $\mathrm{w}_{\max}$ from the knowledge of r (upper bound on the tilt), and the minimum weight of all satisfying assignments seen so far. Every invocation of BoundedWeightSAT also accesses an NP-oracle, called SolveSAT, which can decide SAT. In addition, it accesses a subroutine AddBlockClause that takes as inputs a formula F and a projected assignment $\sigma|_S$, computes a blocking clause for $\sigma|_S$, and returns the formula F' obtained by conjoining F with the blocking clause.

4.1 WeightMC Algorithm

The pseudo-code for WeightMC is shown in Algorithm 1. The algorithm takes as inputs a CNF formula F, tolerance $\varepsilon \in (0, 1)$, confidence parameter $\delta \in (0, 1)$, independent support S, and upper bound r on the tilt, and returns an approximate weighted model count. WeightMC invokes an auxiliary procedure WeightMCCore that computes an approximate weighted model count by randomly partitioning the space of satisfying assignments using hash functions from the family $H_{xor}(|S|, m, 3)$. WeightMC first computes two parameters: pivot, which quantifies the size of a "small" cell, and t, which determines the number of invocations of WeightMC. The particular choice of expressions to compute these parameters is motivated by technical reasons. After invoking WeightMCCore sufficiently many times, WeightMC returns the median of the non-$\perp$ counts returned by WeightMCCore.

The pseudo-code for subroutine WeightMCCore is presented in Algorithm 2. WeightMCCore takes as inputs a CNF formula F, independent support S, parameter pivot to quantify the size of "small" cells, upper bound r on the tilt, and the current upper bound on w_{max}, and returns an approximate weighted model count and a revised upper bound on w_{max}. WeightMCCore first handles the easy case of the total weighted count of F being less than pivot in lines 1–4. Otherwise, in every iteration of the loop in lines 7–12, WeightMCCore randomly partitions the solution space of F using $H_{xor}(|S|, i, 3)$ until a randomly chosen cell is "small" i.e. the total weighted count of the cell is less than pivot. We also refine the estimate for w_{max} in every iteration of the loop in lines 7–12 using the minimum weight of solutions seen so far (computed by BoundedWeightSAT). In the event a chosen cell is "small", WeightMCCore multiplies the weighted count of the cell by the total number of cells

to obtain the estimated total weighted count. The estimated total weighted count along with a refined estimate of w_{max} is finally returned in line 14.

Theorem 1. *Given a propositional formula F, $\varepsilon \in (0,1)$, $\delta \in (0,1)$, independent support S, and upper bound r of the tilt, suppose* $\mathsf{WeightMC}(F,\varepsilon,\delta,S,r)$ *returns c. Then* $\Pr\left[(1+\varepsilon)^{-1} \cdot w(R_F)) \leq c \leq (1+\varepsilon)\cdot w(R_F))\right] \geq 1-\delta$.

Theorem 2. *Given an oracle (*SolveSAT*) for* SAT, $\mathsf{WeightMC}(F,\varepsilon,\delta,S,r)$ *runs in time polynomial in $\log_2(1/\delta), r, |F|$ and $1/\varepsilon$ relative to the oracle.*

The proofs of Theorem 1 and 2 can be found in (Chakraborty et al. 2014).

4.2 WeightGen Algorithm

The pseudo-code for WeightGen is presented in Algorithm 4. WeightGen takes as inputs a CNF formula F, tolerance $\varepsilon > 6.84$, upper bound r of the tilt, and independent support S, and returns a random (approximately weighted-uniform) satisfying assignment of F.

WeightGen can be viewed as an adaptation of UniGen (Chakraborty, Meel, and Vardi 2014) to the weighted domain. It first computes two parameters, κ and pivot, and then uses them to compute $\mathrm{hiThresh}$ and $\mathrm{loThresh}$, which quantify the size of a "small" cell. The easy case of the weighted count being less than $\mathrm{hiThresh}$ is handled in lines 6–9. Otherwise, WeightMC is called to estimate the weighted model count. This is then used to estimate a range of candidate values for i, where a random hash function from $H_{xor}(|S|,i,3)$ must be used to randomly partition the solution space. The choice of parameters for the invocation of WeightMC is motivated by technical reasons. The loop in lines 13–19 terminates when a small cell is found; a sample is then picked weighted-uniformly at random from that cell. Otherwise, the algorithm reports a failure in line 20.

Theorem 3. *Given a CNF formula F, tolerance $\varepsilon > 6.84$, upper bound r of the tilt, and independent support S, for every $y \in R_F$ we have* $\frac{w(y)}{(1+\varepsilon)w(R_F)} \leq \Pr\left[\mathsf{WeightGen}(F,\varepsilon,r,X)=y\right] \leq (1+\varepsilon)\frac{w(y)}{w(R_F)}$. *Also,* WeightGen *succeeds (i.e. does not return $\perp$) with probability at least 0.52.*

Theorem 4. *Given an oracle (*SolveSAT*) for* SAT, $\mathsf{WeightGen}(F,\varepsilon,r,S)$ *runs in time polynomial in $r, |F|$ and $1/\varepsilon$ relative to the oracle.*

The proofs of Theorem 3 and 4 can be found in (Chakraborty et al. 2014).

Implementation Details: In our implementations of WeightGen and WeightMC, BoundedWeightSAT is implemented using CryptoMiniSAT (Cry), 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. We used "random_device" implemented in C++11 as a source of pseudo-random numbers to make random choices in WeightGen and WeightMC.

4.3 Generalization

We have so far restricted S to be an independent support of F in order to ensure 3-wise independence of $H_{xor}(|S|,m,3)$ over the entire solution space R_F. Indeed, this is crucial for proving the theorems presented in this section. However, similar to (Chakraborty, Meel, and Vardi 2014), our results can be generalized to arbitrary subsets S of the support of F. For an arbitrary S, Theorems 3 and 1 generalize to weighted sampling and weighted counting over the solution space projected onto S. To illustrate the projection of the solution space (R_F) over S, consider $F = (a \vee b)$ and $S = \{b\}$. Then the projection of R_F over S, denoted by $R_{F|S}$, is $\{\{0\},\{1\}\}$. This generalization allows our algorithms to extend to formulas of the form $\exists(\cdot)F$ without incurring any additional cost. We discuss this generalization in detail in the full version of our paper (Chakraborty et al. 2014).

5 Experimental Results

To evaluate the performance of WeightGen and WeightMC, we built prototype implementations and conducted an extensive set of experiments. The suite of benchmarks was made up of problems arising from various practical domains as well as problems of theoretical interest. Specifically, we used bit-level unweighted versions of constraints arising from grid networks, plan recognition, DQMR networks, bounded model checking of circuits, bit-blasted versions of SMT-LIB (SMT) benchmarks, and ISCAS89 (Brglez, Bryan, and Kozminski 1989) circuits with parity conditions on randomly chosen subsets of outputs and next-state variables (Sang, Beame, and Kautz 2005; John and Chakraborty 2011). While our algorithms are agnostic to the weight oracle, other tools that we used for comparison require the weight of an assignment to be the product of the weights of its literals. Consequently, to create weighted problems with tilt bounded by some number r, we randomly selected $m = \max(15, n/100)$ of the n variables in a problem instance and assigned them the weight w, where $(w/(1-w))^m = r$, and assigned their negations the weight $1-w$. All other literals were assigned the weight 1. To demonstrate the agnostic nature of our algorithms with respect to the weight oracle, we also evaluated WeightMC and WeightGen with a non-factored weight representation. However, we do not present these results here due to lack of space and also due to the unavailability of a competing tool with which to compare our results on benchmarks with non-factored weights. Details of these experiments are presented in the full version of our paper. Unless mentioned otherwise, our experiments for WeightMC reported in this paper used $r = 5$, $\varepsilon = 0.8$, and $\delta = 0.2$, while our experiments for WeightGen used $r = 5$ and $\varepsilon = 16$.

Table 1: WeightMC, SDD, and WeightGen runtimes in seconds.

Benchmark	vars	#clas	Weight-MC	SDD	Weight-Gen
or-50	100	266	17	0.39	0.17
or-60	120	323	115	0.51	1.5
s526a_3_2	366	944	115	13	1.27
s526_15_7	452	1303	119	41	2.74
s953a_3_2	515	1297	12994	357	33
s1196a_15_7	777	2165	2559	1809	30
s1238a_7_4	704	1926	2885	mem	45
Squaring1	891	2839	18276	mem	164
Squaring7	1628	5837	21982	mem	175
LoginService2	11511	41411	207	mem	4.81
Sort	12125	49611	39641	T	258
Karatsuba	19594	82417	4352	T	326
EnqueueSeq	16466	58515	5763	mem	96
TreeMax	24859	103762	41	mem	4.2
LLReverse	63797	257657	1465	mem	298

To facilitate performing multiple experiments in parallel, we used a high performance cluster, with each experiment running on its own core. Each node of the cluster had two quad-core Intel Xeon processors with 4GB of main memory. We used 2500 seconds as the timeout of each invocation of BoundedWeightSAT and 20 hours as the overall timeout for WeightGen and WeightMC. If an invocation of BoundedWeightSAT timed out in line 11 of WeightMCCore or line 17 of WeightGen, we repeated the execution of the corresponding loops without incrementing the variable i in both algorithms. With this setup, WeightMC and WeightGen were able to successfully return weighted counts and generate weighted random solutions for formulas with almost $64,000$ variables.

We compared the performance of WeightMC with the SDD Package (sdd), a state-of-the-art tool which can perform exact weighted model counting by compiling CNF formulae into Sentential Decision Diagrams (Choi and Darwiche 2013). We also tried to compare our tools against Cachet, WISH, and PAWS, but the current versions of the tools made available to us were broken and we are yet, at the time of submission, to receive working tools. If we are granted access to working tools in future, we will update the full version of our paper with the corresponding comparisons. Our results are shown in Table 1, where column 1 lists the benchmarks and columns 2 and 3 give the number of variables and clauses for each benchmark. Column 4 lists the time taken by WeightMC, while column 5 lists the time taken by SDD. We also measured the time taken by WeightGen to generate samples, which we discuss later in this section, and list it in column 6. "T" and "mem" indicate that an experiment exceeded our imposed 20-hour and 4GB-memory limits, respectively. While SDD was generally superior for small problems, WeightMC was significantly faster for all benchmarks with more than 1,000 variables.

To evaluate the quality of the approximate counts returned by WeightMC, we computed exact weighted model counts using the SDD tool for a subset of our benchmarks. Figure 1 shows the counts returned by WeightMC, and the exact counts from SDD scaled up and down by $1 + \varepsilon$. The weighted model counts are represented on the y-axis, while the x-axis represents benchmarks arranged in increasing order of counts. We observe that the weighted counts returned by WeightMC always lie within the tolerance of the exact counts. Over all of the benchmarks, the L_2 norm of the relative error was 0.014, demonstrating that in practice WeightMC is substantially more accurate than the theoretical guarantees provided by Theorem 3.

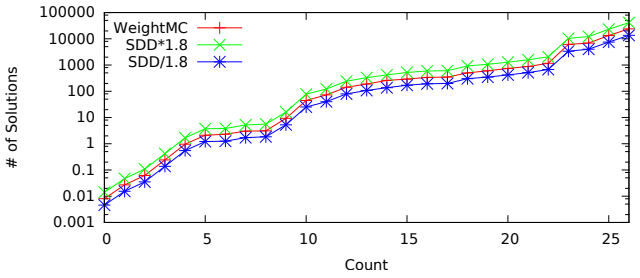

Figure 1: Quality of counts computed by WeightMC. The benchmarks are arranged in increasing order of weighted model counts.

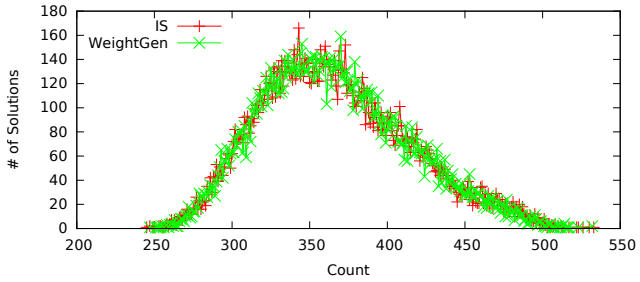

Figure 2: Uniformity comparison for case110

In another experiment, we studied the effect of different values of the tilt bound r on the runtime of WeightMC. Runtime as a function of r is shown for several benchmarks in Figure 3, where times have been normalized so that at the lowest tilt ($r = 1$) each benchmark took one (normalized) time unit. Each runtime is an average over 15 runs on the same benchmark. The theoretical linear dependence on the tilt shown in Theorem 2 can be seen to roughly occur in practice.

Since a probabilistic generator is likely to be invoked many times with the same formula and weights, it is useful to perform the counting on line 10 of WeightGen only once, and reuse the result for every sample. Reflecting this, column 6 in Table 1 lists the time, averaged over a large number of runs, taken by WeightGen to generate one sample given that the weighted model count on line 10 has already been found. It is clear from Table 1 that WeightGen scales to formulas with thousands of variables.

To measure the accuracy of WeightGen, we implemented an *Ideal Sampler*, henceforth called IS, and compared the

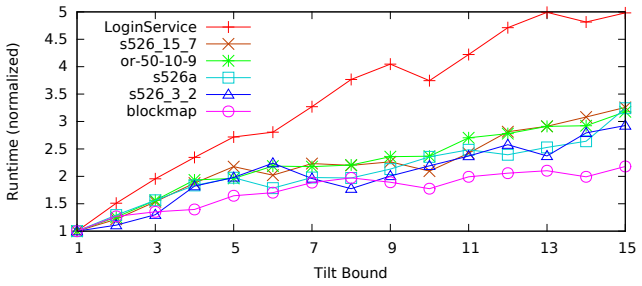

Figure 3: Runtime of WeightMC as a function of tilt bound.

distributions generated by WeightGen and IS for a representative benchmark. Given a CNF formula F, IS first generates all the satisfying assignments, then computes their weights and uses these to sample the ideal distribution. We generated a large number N ($= 6 \times 10^6$) of sample witnesses using both IS and WeightGen. In each case, the number of times various witnesses were generated was recorded, yielding a distribution of the counts. Figure 2 shows the distributions generated by WeightGen and IS for one of our benchmarks (case110) with 16,384 solutions. The almost perfect match between the distribution generated by IS and WeightGen held also for other benchmarks. Thus, as was the case for WeightMC, the accuracy of WeightGen is better in practice than that established by Theorem 3.

6 White-Box Weight Functions

As noted above, the runtime of WeightMC is proportional to the tilt of the weight function, which means that the algorithm becomes impractical when the tilt is large. If the assignment weights are given by a known polynomial-time-computable function instead of an oracle, we can do better. We abuse notation slightly and denote this weight function by $w(X)$, where X is the set of support variables of the Boolean formula F. The essential idea is to partition the set of satisfying assignments into regions within which the tilt is small. Defining $R_F(a,b) = \{\sigma \in R_F | a < w(\sigma) \leq b\}$, we have $w(R_F) = w(R_F(w_{min}, w_{max}))$. If we use a partition of the form $R_F(w_{min}, w_{max}) = R_F(w_{max}/2, w_{max}) \cup R_F(w_{max}/4, w_{max}/2) \cup \cdots \cup R_F(w_{max}/2^N, w_{max}/2^{N-1})$, where $w_{max}/2^N \leq w_{min}$, then in each partition region the tilt is at most 2. Note that we do not need to know the actual values of w_{min} and w_{max}: any bounds L and H such that $0 < L \leq w_{min}$ and $w_{max} \leq H$ will do (although if the bounds are too loose, we may partition R_F into more regions than necessary). If assignment weights are poly-time computable, we can add to F a constraint that eliminates all assignments not in a particular region. So we can run WeightMC on each region in turn, passing 2 as the upper bound on the tilt, and sum the results to get $w(R_F)$. This idea is implemented in PartitionedWeightMC (Algorithm 6).

The correctness and runtime of PartitionedWeightMC are established by the following theorems, whose proofs are presented in (Chakraborty et al. 2014).

Algorithm 6 PartitionedWeightMC$(F, \varepsilon, \delta, S, L, H)$

$N \leftarrow \lceil \log_2 H/L \rceil + 1$; $\delta' \leftarrow \delta/N$; $c \leftarrow 0$
for all $1 \leq m \leq N$ **do**
  $G \leftarrow F \wedge (H/2^m < w(X) \leq H/2^{m-1})$
  $d \leftarrow$ WeightMC$(G, \varepsilon, \delta', S, 2)$
  if $(d = \perp)$ **then return** $\perp$
  $c \leftarrow c + d$
return c

Theorem 5. *If* PartitionedWeightMC$(F, \varepsilon, \delta, S, L, H)$ *returns* c *(and all arguments are in the required ranges), then* $\Pr\left[(1+\varepsilon)^{-1} w(R_F) \leq c \leq (1+\varepsilon) w(R_F))\right] \geq 1 - \delta$.

Theorem 6. *With access to an* NP *oracle, the runtime of* PartitionedWeightMC$(F, \varepsilon, \delta, S, L, H)$ *is polynomial in* $|F|$, $1/\varepsilon$, $\log(1/\delta)$, *and* $\log r = \log(H/L)$.

The reduction of the runtime's dependence on the tilt bound r from linear to logarithmic can be a substantial saving. If the assignment weights are products of literal weights, as is the case in many applications, the best *a priori* bound on the tilt ρ given only the literal weights is exponential in n. Thus, unless the structure of the problem allows a better bound on ρ to be used, WeightMC will not be practical. In this situation PartitionedWeightMC can be used to maintain polynomial runtime.

When implementing PartitionedWeightMC in practice the handling of the weight constraint $H/2^m < w(X) \leq H/2^{m-1}$ is critical to efficiency. If assignment weights are sums of literal weights, or equivalently products of literal weights (we just take logarithms), then the weight constraint is a pseudo-Boolean constraint. In this case we may replace the SAT-solver used by WeightMC with a pseudo-Poolean satisfiability (PBS) solver. While a number of PBS-solvers exist (Manquinho and Roussel 2012), none have the specialized handling of XOR clauses that is critical in making WeightMC practical. The design of such solvers is a clear direction for future work. We also note that the choice of 2 as the tilt bound for each region is arbitrary, and the value may be adjusted depending on the application: larger values will decrease the number of regions, but increase the difficulty of counting within each region. Finally, note that the same partitioning idea can be used to reduce WeightGen's dependence on r to be logarithmic.

7 Conclusion

In this paper, we considered approximate approaches to the twin problems of distribution-aware sampling and weighted model counting for SAT. For approximation techniques that provide strong theoretical two-way bounds, a major limitation is the reliance on potentially-expensive most probable explanation (MPE) queries. We identified a novel parameter, *tilt*, to categorize weighted counting and sampling problems for SAT. We showed how to remove the reliance on MPE-queries, while retaining strong theoretical guarantees. First, we provided model counting and sampling algorithms that work with a black-box model of giving weights to assignments, requiring access only to an NP-oracle, which is

efficient for small tilt values. Experimental results demonstrate the effectiveness of this approach in practice. Second, we provided an alternative approach that promises to be efficient for large tilt values, requiring, however, a white-box weight model and access to a pseudo-Boolean solver. As a next step, we plan to empirically evaluate this latter approach using pseudo-Boolean solvers designed to handle parity constraints efficiently.

Acknowledgments: The authors would like to thank Armando Solar-Lezama for generously providing a large set of benchmarks. We thank the anonymous reviewers for their detailed, insightful comments and suggestions that helped improve this manuscript significantly.

References

Bacchus, F.; Dalmao, S.; and Pitassi, T. 2003. Algorithms and complexity results for #SAT and Bayesian inference. In *Proc. of FOCS*, 340–351.

Bellare, M.; Goldreich, O.; and Petrank, E. 1998. Uniform generation of NP-witnesses using an NP-oracle. *Information and Computation* 163(2):510–526.

Brglez, F.; Bryan, D.; and Kozminski, K. 1989. Combinational profiles of sequential benchmark circuits. In *Proc. of ISCAS*, 1929–1934.

Chakraborty, S.; Fremont, D. J.; Meel, K. S.; Seshia, S.; and Vardi, M. Y. 2014. Distribution-aware sampling and weighted model counting for SAT (Technical Report). `http://arxiv.org/abs/1403.6246`.

Chakraborty, S.; Meel, K. S.; and Vardi, M. Y. 2013a. A scalable and nearly uniform generator of SAT witnesses. In *Proc. of CAV*, 608–623.

Chakraborty, S.; Meel, K. S.; and Vardi, M. Y. 2013b. A scalable approximate model counter. In *Proc. of CP*, 200–216.

Chakraborty, S.; Meel, K. S.; and Vardi, M. Y. 2014. Balancing scalability and uniformity in SAT-witness generator. In *Proc. of DAC*.

Chavira, M., and Darwiche, A. 2008. On probabilistic inference by weighted model counting. *Artificial Intelligence* 172(6):772–799.

Choi, A., and Darwiche, A. 2013. Dynamic minimization of sentential decision diagrams. In *Proc. of AAAI*, 187–194.

CryptoMiniSAT. `http://www.msoos.org/cryptominisat2/`.

Darwiche, A. 2004. New advances in compiling CNF to decomposable negation normal form. In *Proc. of ECAI*, 328–332.

Dıez, F. J., and Druzdzel, M. J. 2006. Canonical probabilistic models for knowledge engineering. Technical report, CISIAD-06-01, UNED, Madrid, Spain.

Domshlak, C., and Hoffmann, J. 2007. Probabilistic planning via heuristic forward search and weighted model counting. *Journal of Artificial Intelligence Research* 30(1):565–620.

Ermon, S.; Gomes, C. P.; Sabharwal, A.; and Selman, B. 2013a. Embed and project: Discrete sampling with universal hashing. In *Proc of NIPS*, 2085–2093.

Ermon, S.; Gomes, C. P.; Sabharwal, A.; and Selman, B. 2013b. Optimization with parity constraints: From binary codes to discrete integration. In *Proc. of UAI*, 202–211.

Ermon, S.; Gomes, C. P.; Sabharwal, A.; and Selman, B. 2013c. Taming the curse of dimensionality: Discrete integration by hashing and optimization. In *Proc. of ICML*, 334–342.

Ermon, S.; Gomes, C. P.; Sabharwal, A.; and Selman, B. 2014. Low-density parity constraints for hashing-based discrete integration. In *Proc. of ICML*, 271–279.

Gogate, V., and Dechter, R. 2011. Samplesearch: Importance sampling in presence of determinism. *Artificial Intelligence* 175(2):694–729.

Gomes, C. P.; Sabharwal, A.; and Selman, B. 2006. Model counting: A new strategy for obtaining good bounds. In *Proc. of AAAI*, 54–61.

Gomes, C. P.; Sabharwal, A.; and Selman, B. 2007. Near uniform sampling of combinatorial spaces using XOR constraints. In *Proc. of NIPS*, 670–676.

Jerrum, M., and Sinclair, A. 1996. The Markov chain Monte Carlo method: an approach to approximate counting and integration. *Approximation algorithms for NP-hard problems* 482–520.

Jerrum, M.; Valiant, L.; and Vazirani, V. 1986. Random generation of combinatorial structures from a uniform distribution. *TCS* 43(2-3):169–188.

John, A., and Chakraborty, S. 2011. A quantifier elimination algorithm for linear modular equations and disequations. In *Proc. of CAV*, 486–503.

Karp, R.; Luby, M.; and Madras, N. 1989. Monte-Carlo approximation algorithms for enumeration problems. *Journal of Algorithms* 10(3):429–448.

Kirkpatrick, S.; Gelatt, C. D.; and Vecchi, M. P. 1983. Optimization by simulated annealing. *Science* 220(4598):671–680.

Madras, N., and Piccioni, M. 1999. Importance sampling for families of distributions. *Annals of applied probability* 9(4):1202–1225.

Madras, N. 2002. *Lectures on Monte Carlo Methods*, volume 16 of *Fields Institute Monographs*. AMS.

Manquinho, V., and Roussel, O. 2012. Seventh pseudo-Boolean competition. `http://www.cril.univ-artois.fr/PB12/`.

Naveh, Y.; Rimon, M.; Jaeger, I.; Katz, Y.; Vinov, M.; Marcus, E.; and Shurek, G. 2006. Constraint-based random stimuli generation for hardware verification. In *Proc. of IAAI*, 1720–1727.

Park, J. D., and Darwiche, A. 2004. Complexity results and approximation strategies for map explanations. *J. Artif. Intell. Res.* 21:101–133.

Roth, D. 1996. On the hardness of approximate reasoning. *Artificial Intelligence* 82(1):273–302.

Sang, T.; Beame, P.; and Kautz, H. 2005. Performing Bayesian inference by weighted model counting. In *Proc. of AAAI*, 475–481.

Sang, T.; Bacchus, F.; Beame, P.; Kautz, H.; and Pitassi, T. 2004. Combining component caching and clause learning for effective model counting. In *Proc. of SAT*.

The SDD package. `http://reasoning.cs.ucla.edu/sdd/`.

SMTLib. `http://goedel.cs.uiowa.edu/smtlib/`.

Toda, S. 1989. On the computational power of PP and (+)P. In *Proc. of FOCS*, 514–519.

Valiant, L. 1979. The complexity of enumeration and reliability problems. *SIAM Journal on Computing* 8(3):410–421.

Wainwright, M. J., and Jordan, M. I. 2008. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning* 1(1-2):1–305.

Xue, Y.; Choi, A.; and Darwiche, A. 2012. Basing decisions on sentences in decision diagrams. In *Proc. of AAAI*, 842–849.

APPENDIX

Using notation introduced in Section 2, let $w(y)$ denote the weight of solution y and R_F denote the set of witnesses of the Boolean formula F. We denote the weight of the set R_F by $w(R_F)$. For brevity, we write $\mathcal{W}(y)$ for the expression $w(y)/\mathrm{w_{max}}$ (where $\mathrm{w_{max}}$ is the variable appearing in several of our algorithms).

Recall that WeightMC is a probabilistic algorithm that takes as inputs a Boolean CNF formula F, a tolerance ε, confidence parameter δ, a subset S of the support of F, and an upper bound r on the ratio ρ. We extend the results in (**?**) and show that if X is the support of F, and if $S \subseteq X$ is an independent support of F, then WeightMC($F, \varepsilon, \delta, S, r$) behaves *identically* (in a probabilistic sense) to WeightMC(F, ε, δ, X, r). Once this is established, the remainder of the proof proceeds by making the simplifying assumption $S = X$. The proofs of Lemmas 1 and 2 extend the earlier results by (**?**) for unweighted sample space.

Clearly, the above claim holds trivially if $X = S$. Therefore, we focus only on the case when $S \subsetneq X$. For notational convenience, we assume $X = \{x_1, \ldots x_n\}$, $0 \leq k < n$, $S = \{x_1, \ldots x_k\}$ and $D = \{x_{k+1}, \ldots x_n\}$ in all the statements and proofs in this section. We also use $\vec{X}$ to denote the vector $(x_1, \ldots x_n)$, and similarly for $\vec{S}$ and $\vec{D}$.

Lemma 1. *Let $F(\vec{X})$ be a Boolean function, and S an independent support of F. Then there exist Boolean functions $g_0, g_1, \ldots g_{n-k}$, each with support S such that*

$$F(\vec{X}) \leftrightarrow \left(g_0(\vec{S}) \wedge \bigwedge_{j=1}^{n-k} (x_{k+j} \leftrightarrow g_j(\vec{S})) \right)$$

Proof. Since S is an independent support of F, the set $D = X \setminus S$ is a dependent support of F. From the definition of a dependent support, there exist Boolean functions $g_1, \ldots g_k$, each with support S, such that $F(\vec{X}) \rightarrow \bigwedge_{j=1}^{n-k}(x_{k+j} \leftrightarrow g_j(\vec{S}))$.

Let $g_0(\vec{S})$ be the characteristic function of the projection of R_F on S. More formally, $g_0(\vec{S}) \equiv \bigvee_{(x_{k+1}, \ldots x_n) \in \{0,1\}^{n-k}} F(\vec{X})$. It follows that $F(\vec{X}) \rightarrow g_0(\vec{S})$. Combining this with the result from the previous paragraph, we get the implication $F(\vec{X}) \rightarrow \left(g_0(\vec{S}) \wedge \bigwedge_{j=1}^{n-k}(x_{k+j} \leftrightarrow g_j(\vec{S})) \right)$

From the definition of $g_0(\vec{S})$ given above, we have $g_0(\vec{S}) \rightarrow F(\vec{S}, x_{k+1}, \ldots x_n)$, for some values of $x_{k+1}, \ldots x_n$. However, we also know that $F(\vec{X}) \rightarrow \bigwedge_{j=1}^{n-k}(x_{k+j} \leftrightarrow g_j(\vec{S}))$. It follows that $\left(g_0(\vec{S}) \wedge \bigwedge_{j=1}^{n-k}(x_{k+j} \leftrightarrow g_j(\vec{S})) \right) \rightarrow F(\vec{X})$. □

Referring to the pseudocode of WeightMC in Section 4, we observe that the only steps that depend directly on S are those in line 9, where h is chosen randomly from $H_{xor}(|S|, i, 3)$, and line 11, where the set Y is computed by calling BoundedWeightSAT($F \wedge (h(x_1, \ldots x_{|S|}) = \alpha), pivot, r, \mathrm{w_{max}}$). Since all subsequent steps of the algorithm depend only on Y, it suffices to show that if S is an independent support of F, the probability distribution of Y obtained at line 11 is *identical* to what we would obtain if S was set equal to the entire support X.

The following lemma formalizes the above statement. As before, we assume $X = \{x_1, \ldots x_n\}$ and $S = \{x_1, \ldots x_k\}$.

Lemma 2. *Let S be an independent support of $F(\vec{X})$. Let h and h' be hash functions chosen uniformly at random from $H_{xor}(k, i, 3)$ and $H_{xor}(n, i, 3)$, respectively. Let α and α' be tuples chosen uniformly at random from $\{0,1\}^i$. Then, for every $Y \in \{0,1\}^n$, $pivot > 0$, $r \geq 1$, and $\mathrm{w_{max}} \geq 1$, we have*

$$\mathsf{Pr}\left[\mathsf{BoundedWeightSAT}\left(F(\vec{X}) \wedge (h(\vec{S}) = \alpha), pivot, r, \mathrm{w_{max}}\right) = Y\right]$$
$$= \mathsf{Pr}\left[\mathsf{BoundedWeightSAT}\left(F(\vec{X}) \wedge (h'(\vec{X}) = \alpha'), pivot, r, \mathrm{w_{max}}\right) =\right.$$

Proof. Since h' is chosen uniformly at random from $H_{xor}(n, i, 3)$, recalling the definition of the latter we have $F(\vec{X}) \wedge (h'(\vec{X}) = \alpha') \equiv F(\vec{X}) \wedge \bigwedge_{l=1}^{i} \left((a_{l,0} \oplus \bigoplus_{j=1}^{n} a_{l,j} \cdot x[j]) \leftrightarrow \alpha'[l] \right)$, where the coefficients $a_{l,j}$ are chosen i.i.d. uniformly from $\{0,1\}$.

Since S is an independent support of F, from Lemma 1, there exist Boolean functions $g_1, \ldots g_{n-k}$, each with support S, such that $F(\vec{X}) \rightarrow \bigwedge_{j=1}^{n-k}(x_{k+j} \leftrightarrow g_j(\vec{S}))$. Therefore, $F(\vec{X}) \wedge (h'(\vec{X}) = \alpha')$ holds iff $F(\vec{X}) \wedge \bigwedge_{l=1}^{i} \left((a_{l,0} \oplus \bigoplus_{j=1}^{k} a_{l,j} \cdot x[j] \oplus B) \leftrightarrow \alpha'[l] \right)$ does, where $B \equiv \bigoplus_{j=k+1}^{n} a_{l,j} \cdot g_{j-k}(\vec{S})$. Rearranging terms, we get $F(\vec{X}) \wedge \bigwedge_{l=1}^{i} \left((a_{l,0} \oplus \bigoplus_{j=1}^{k} a_{l,j} \cdot x[j]) \leftrightarrow (\alpha'[l] \oplus B) \right)$. Now since α' is chosen uniformly at random from $\{0,1\}^i$ and since B is independent of α', we have that $\alpha'[l] \oplus B$ is a random binary variable with equal probability of being 0 and 1. So $\alpha'[l] \oplus B$ has the same distribution as $\alpha[l]$, and the result follows. □

Lemma 2 allows us to continue with the remainder of the proof assuming $S = X$. It has already been shown in (Gomes, Sabharwal, and Selman 2007) that $H_{xor}(n, m, 3)$ is a 3-independent family of hash functions. We use this fact in a key way in the remainder of our analysis. The following result about Chernoff-Hoeffding bounds, proved in (Chakraborty, Meel, and Vardi 2013a), plays an important role in our discussion.

Lemma 3. *Let Γ be the sum of r-wise independent random variables, each of which is confined to the interval $[0,1]$, and suppose $\mathsf{E}[\Gamma] = \mu$. For $0 < \beta \leq 1$, if $2 \leq r \leq 3$ and $\mu \geq \lceil e^{3/2}\beta^2 \rceil$, then $\mathsf{Pr}[\,|\Gamma - \mu| \geq \beta\mu\,] \leq e^{-3/2}$.*

8 Analysis of WeightMC

In this section we denote the quantity $\log_2 \mathcal{W}(R_F) - \log_2 pivot + 1$ by m. For simplicity of exposition, we assume henceforth that m is an integer. A more careful analysis removes this restriction with only a constant factor scaling of the probabilities.

Lemma 4. *Let algorithm* WeightMCCore*, when invoked from* WeightMC*, return c with i being the final value of the loop counter in* WeightMCCore*. Then* $\Pr\left[(1+\varepsilon)^{-1} \cdot \mathcal{W}(R_F) \leq c \leq (1+\varepsilon) \cdot \mathcal{W}(R_F) \;\middle|\; c \neq \perp \wedge i \leq m\right] \geq 1 - e^{-3/2}$.

Proof. Referring to the pseudocode of WeightMCCore, the lemma is trivially satisfied if $\mathcal{W}(R_F) \leq pivot$. Therefore, the only non-trivial case to consider is when $\mathcal{W}(R_F) > pivot$ and WeightMCCore returns from line 14. In this case, the count returned is $2^i \cdot \mathcal{W}(R_{F,h,\alpha})$, where α, i and h denote (with abuse of notation) the values of the corresponding variables and hash functions in the final iteration of the repeat-until loop in lines 7–12 of the pseudocode. From the pseudocode of WeightMCCore, we know that $pivot = \lceil e^{3/2}(1+1/\varepsilon)^2 \rceil$. The lemma is now proved by showing that for every i in $\{0, \ldots m\}$, $h \in H(n,i,3)$, and $\alpha \in \{0,1\}^i$, we have $\Pr\left[(1+\varepsilon)^{-1} \cdot \mathcal{W}(R_F) \leq 2^i \mathcal{W}(R_{F,h,\alpha}) \leq (1+\varepsilon) \cdot \mathcal{W}(R_F)\right] \geq 1 - e^{-3/2}$.

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

Let $\Gamma_\alpha = \sum_{y \in R_F} \gamma_{y,\alpha}$ and $\mu_\alpha = \mathsf{E}[\Gamma_\alpha]$. Clearly, $\Gamma_\alpha = \mathcal{W}(R_{F,h,\alpha})$ and $\mu_\alpha = \sum_{y \in R_F} \mathsf{E}[\gamma_{y,\alpha}] = 2^{-i}\mathcal{W}(R_F)$. Therefore, using Lemma 3 with $\beta = \varepsilon/(1+\varepsilon)$, we have $\Pr\left[\mathcal{W}(R_F)\left(1 - \frac{\varepsilon}{1+\varepsilon}\right) \leq 2^i \mathcal{W}(R_{F,h,\alpha}) \leq (1 + \frac{\varepsilon}{1+\varepsilon})\mathcal{W}(R_F)\right] \geq 1 - e^{-3/2}$. Simplifying and noting that $\frac{\varepsilon}{1+\varepsilon} < \varepsilon$ for all $\varepsilon > 0$, we obtain $\Pr\left[(1+\varepsilon)^{-1} \cdot \mathcal{W}(R_F) \leq 2^i \mathcal{W}(R_{F,h,\alpha}) \leq (1+\varepsilon) \cdot \mathcal{W}(R_F)\right] \geq 1 - e^{-3/2}$. □

Lemma 5. *Given $\mathcal{W}(R_F) > pivot$, the probability that an invocation of* WeightMCCore *from* WeightMC *returns non-$\perp$ with $i \leq m$, is at least $1 - e^{-3/2}$.*

Proof. Let p_i $(0 \leq i \leq n)$ denote the conditional probability that WeightMCCore terminates in iteration i of the repeat-until loop (lines 7–12 of the pseudocode) with $0 < \mathcal{W}(R_{F,h,\alpha}) \leq pivot$, given $\mathcal{W}(R_F) > pivot$. Since the choice of h and α in each iteration of the loop are independent of those in previous iterations, the conditional probability that WeightMCCore returns non-$\perp$ with $i \leq m$, given $\mathcal{W}(R_F) > pivot$, is $p_0 + (1-p_0)p_1 + \cdots + (1-p_0)(1-p_1)\cdots(1-p_{m-1})p_m$. Let us denote this sum by P. Thus, $P = p_0 + \sum_{i=1}^{m} \prod_{k=0}^{i-1}(1-p_k)p_i \geq \left(p_0 + \sum_{i=1}^{m-1} \prod_{k=0}^{i-1}(1-p_k)p_i\right) p_m + \prod_{s=0}^{m-1}(1-p_s)p_m = p_m$. The lemma is now proved by showing that $p_m \geq 1 - e^{-3/2}$.

It was shown in the proof of Lemma 4 that $\Pr\left[(1+\varepsilon)^{-1} \cdot \mathcal{W}(R_F) \leq 2^i \mathcal{W}(R_{F,h,\alpha}) \leq (1+\varepsilon) \cdot \mathcal{W}(R_F)\right] \geq 1 - e^{-3/2}$ for every $i \in \{0, \ldots, m\}$, $h \in H(n,i,3)$ and $\alpha \in \{0,1\}^i$. Substituting m for i, re-arranging terms and noting that the definition of m implies $2^{-m}\mathcal{W}(R_F) = pivot/2$, we get $\Pr\left[(1+\varepsilon)^{-1}(pivot/2) \leq \mathcal{W}(R_{F,h,\alpha}) \leq (1+\varepsilon)(pivot/2)\right] \geq 1 - e^{-3/2}$. Since $0 < \varepsilon \leq 1$ and $pivot > 4$, it follows that $\Pr[0 < \mathcal{W}(R_{F,h,\alpha}) \leq pivot] \geq 1 - e^{-3/2}$. Hence, $p_m \geq 1 - e^{-3/2}$. □

Lemma 6. *Let an invocation of* WeightMCCore *from* WeightMC *return c. Then* $\Pr\left[c \neq \perp \wedge (1+\varepsilon)^{-1} \cdot w(R_F) \leq c \cdot \mathrm{w_{max}} \leq (1+\varepsilon) \cdot w(R_F)\right] \geq (1 - e^{-3/2})^2 > 0.6$.

Proof. It is easy to see that the required probability is at least as large as $\Pr\left[c \neq \perp \wedge i \leq m \wedge (1+\varepsilon)^{-1} w(R_F) \leq c \cdot \mathrm{w_{max}} \leq (1+\varepsilon) \cdot w(R_F)\right]$. Dividing by $\mathrm{w_{max}}$ and applying Lemmas 4 and 5, this probability is $\geq (1 - e^{-3/2})^2$. □

We now turn to proving that the confidence can be raised to at least $1 - \delta$ for $\delta \in (0,1]$ by invoking WeightMCCore $\mathcal{O}(\log_2(1/\delta))$ times, and by using the median of the non-$\perp$ counts thus returned. For convenience of exposition, we use $\eta(t,m,p)$ in the following discussion to denote the probability of at least m heads in t independent tosses of a biased coin with $\Pr[heads] = p$. Clearly, $\eta(t,m,p) = \sum_{k=m}^{t} \binom{t}{k} p^k (1-p)^{t-k}$.

Theorem 1. *Given a propositional formula F and parameters ε $(0 < \varepsilon \leq 1)$ and δ $(0 < \delta \leq 1)$, suppose* WeightMC$(F, \varepsilon, \delta, X, r)$ *returns c. Then* $\Pr\left[(1+\varepsilon)^{-1} \cdot w(R_F)) \leq c \leq (1+\varepsilon) \cdot w(R_F))\right] \geq 1 - \delta$.

Proof. Throughout this proof, we assume that WeightMCCore is invoked t times from WeightMC, where $t = \lceil 35 \log_2(3/\delta) \rceil$ (see pseudocode for ComputeIterCount in Section **??**). Referring to the pseudocode of WeightMC, the final count returned is the median of the non-$\perp$ counts obtained from the t invocations of WeightMCCore. Let Err denote the event that the median is not in $\left[(1+\varepsilon)^{-1} \cdot \mathcal{W}(R_F), (1+\varepsilon) \cdot \mathcal{W}(R_F)\right]$. Let "$\#non\perp = q$" denote the event that q (out of t) values returned by WeightMCCore are non-$\perp$. Then, $\Pr[Err] = \sum_{q=0}^{t} \Pr[Err \mid \#non\perp = q] \cdot \Pr[\#non\perp = q]$.

In order to obtain $\Pr[Err \mid \#non\perp = q]$, we define a 0-1 random variable Z_i, for $1 \leq i \leq t$, as follows. If the i^{th} invocation of WeightMCCore returns c, and if c is either $\perp$ or a non-$\perp$ value that does not lie in the interval $[(1+\varepsilon)^{-1} \cdot \mathcal{W}(R_F), (1+\varepsilon) \cdot \mathcal{W}(R_F)]$, we set Z_i to 1; otherwise, we set it to 0. From Lemma 6, $\Pr[Z_i = 1] = p < 0.4$.

If Z denotes $\sum_{i=1}^{t} Z_i$, a necessary (but not sufficient) condition for event Err to occur, given that q non-$\perp$s were returned by WeightMCCore, is $Z \geq (t - q + \lceil q/2 \rceil)$. To see why this is so, note that $t - q$ invocations of WeightMCCore must return $\perp$. In addition, at least $\lceil q/2 \rceil$ of the remaining q invocations must return values outside the desired interval. To simplify the exposition, let q be an even integer. A more careful analysis removes this restriction and results in an additional constant scaling factor for $\mathsf{Pr}[Err]$. With our simplifying assumption, $\mathsf{Pr}[Err \mid \#non\perp = q] \leq \mathsf{Pr}[Z \geq (t - q + q/2)] = \eta(t, t - q/2, p)$. Since $\eta(t, m, p)$ is a decreasing function of m and since $q/2 \leq t - q/2 \leq t$, we have $\mathsf{Pr}[Err \mid \#non\perp = q] \leq \eta(t, t/2, p)$. If $p < 1/2$, it is easy to verify that $\eta(t, t/2, p)$ is an increasing function of p. In our case, $p < 0.4$; hence, $\mathsf{Pr}[Err \mid \#non\perp = q] \leq \eta(t, t/2, 0.4)$.

It follows from the above that $\mathsf{Pr}[Err] = \sum_{q=0}^{t} \mathsf{Pr}[Err \mid \#non\perp = q] \cdot \mathsf{Pr}[\#non\perp = q] \leq \eta(t, t/2, 0.4) \cdot \sum_{q=0}^{t} \mathsf{Pr}[\#non\perp = q] = \eta(t, t/2, 0.4)$. Since $\binom{t}{t/2} \geq \binom{t}{k}$ for all $t/2 \leq k \leq t$, and since $\binom{t}{t/2} \leq 2^t$, we have $\eta(t, t/2, 0.4) = \sum_{k=t/2}^{t} \binom{t}{k} (0.4)^k (0.6)^{t-k} \leq \binom{t}{t/2} \sum_{k=t/2}^{t} (0.4)^k (0.6)^{t-k} \leq 2^t \sum_{k=t/2}^{t} (0.6)^t (0.4/0.6)^k \leq 2^t \cdot 3 \cdot (0.6 \times 0.4)^{t/2} \leq 3 \cdot (0.98)^t$. Since $t = \lceil 35 \log_2(3/\delta) \rceil$, it follows that $\mathsf{Pr}[Err] \leq \delta$. □

Theorem 2. *Given an oracle for* SAT, $\mathsf{WeightMC}(F, \varepsilon, \delta, S, r)$ *runs in time polynomial in* $\log_2(1/\delta), r, |F|$ *and* $1/\varepsilon$ *relative to the oracle.*

Proof. Referring to the pseudocode for WeightMC, lines 1–3 take $\mathcal{O}(1)$ time. The repeat-until loop in lines 4–9 is repeated $t = \lceil 35 \log_2(3/\delta) \rceil$ times. The time taken for each iteration is dominated by the time taken by WeightMCCore. Finally, computing the median in line 10 takes time linear in t. The proof is therefore completed by showing that WeightMCCore takes time polynomial in $|F|, r$ and $1/\varepsilon$ relative to the SAT oracle.

Referring to the pseudocode for WeightMCCore, we find that BoundedWeightSAT is called $\mathcal{O}(|F|)$ times. Observe that when the loop in BoundedWeightSAT terminates, $\mathrm{w_{min}}$ is such that each $y \in R_F$ whose weight was added to $\mathrm{w_{total}}$ has weight at least $\mathrm{w_{min}}$. Thus since the loop terminates when $\mathrm{w_{total}}/\mathrm{w_{min}} > r \cdot pivot$, it can have iterated at most $(r \cdot pivot) + 1$ times. Therefore each call to BoundedWeightSAT makes at most $(r \cdot pivot) + 1$ calls to the SAT oracle, and takes time polynomial in $|F|$, r, and $pivot$ relative to the oracle. Since $pivot$ is in $\mathcal{O}(1/\varepsilon^2)$, the number of calls to the SAT oracle, and the total time taken by all calls to BoundedWeightSAT in each invocation of WeightMCCore is polynomial in $|F|, r$ and $1/\varepsilon$ relative to the oracle. The random choices in lines 9 and 10 of WeightMCCore can be implemented in time polynomial in n (hence, in $|F|$) if we have access to a source of random bits. Constructing $F \wedge h(z_1, \ldots z_n) = \alpha$ in line 11 can also be done in time polynomial in $|F|$. □

9 Analysis of WeightGen

For convenience of analysis, we assume that $\log(\mathcal{W}(R_F) - 1) - \log pivot$ is an integer, where $pivot$ is the quantity computed by algorithm ComputeKappaPivot (see Section 4). A more careful analysis removes this assumption by scaling the probabilities by constant factors. Let us denote $\log(\mathcal{W}(R_F) - 1) - \log pivot$ by m. The expression used for computing $pivot$ in algorithm ComputeKappaPivot ensures that $\mathrm{pivot} \geq 17$. Therefore, if an invocation of WeightGen does not return from line 8 of the pseudocode, then $\mathcal{W}(R_F) \geq 18$. Note also that the expression for computing κ in algorithm ComputeKappaPivot requires $\varepsilon \geq 1.71$ in order to ensure that $\kappa \in [0, 1)$ can always be found.

In the case where $\mathcal{W}(R_F) \leq 1 + (1 + \kappa)pivot$, BoundedWeightSAT returns all witnesses of F and WeightGen returns a perfect weighted-uniform sample on line 8. So we restrict our attention in the lemmas below to the other case, where as noted above we have $\mathcal{W}(R_F) \geq 18$. The following lemma shows that q, computed in line 11 of the pseudocode, is a good estimator of m.

Lemma 7. $\mathsf{Pr}[q - 3 \leq m \leq q] \geq 0.8$

Proof. Recall that in line 10 of the pseudocode, an approximate weighted model counter is invoked to obtain an estimate, C, of $w(R_F)$ with tolerance 0.8 and confidence 0.8. By the definition of approximate weighted model counting, we have $\mathsf{Pr}[\frac{C}{1.8} \leq w(R_F) \leq (1.8)C] \geq 0.8$. Defining $c = C/w_{max}$, we have $\mathsf{Pr}[\log c - \log(1.8) \leq \log \mathcal{W}(R_F) \leq \log c + \log(1.8)] \geq 0.8$. It follows that $\mathsf{Pr}[\log c - \log(1.8) - \log pivot - \log(\frac{1}{1-1/\mathcal{W}(R_F)}) \leq \log(\mathcal{W}(R_F) - 1) - \log pivot \leq \log c - \log pivot + \log(1.8) - \log(\frac{1}{1-1/\mathcal{W}(R_F)})] \geq 0.8$. Substituting $q = \lceil \log C - \log \mathrm{w_{max}} + \log 1.8 - \log pivot \rceil = \lceil \log c + \log 1.8 - \log pivot \rceil$, and using the bounds $\mathrm{w_{max}} \leq 1$, $\log 1.8 \leq 0.85$, and $\log(\frac{1}{1-1/\mathcal{W}(R_F)}) \leq 0.12$ (since $\mathcal{W}(R_F) \geq 18$ at line 10 of the pseudocode, as noted above), we have $\mathsf{Pr}[q - 3 \leq m \leq q] \geq 0.8$. □

The next lemma provides a lower bound on the probability of generation of a witness. Let $w_{i,y,\alpha}$ denote the probability $\mathsf{Pr}\left[\frac{pivot}{1+\kappa} \leq \mathcal{W}(R_{F,h,\alpha}) \leq 1 + (1 + \kappa)pivot \wedge h(y) = \alpha\right]$, with $h \xleftarrow{R} H_{xor}(n, i, 3)$. The proof of the lemma also provides a lower bound on $w_{m,y,\alpha}$.

Lemma 8. *For every witness* $y \in R_F$, $\mathsf{Pr}[y \text{ is output}] \geq \frac{0.8(1-e^{-3/2})\mathcal{W}(y)}{(1.06+\kappa)(\mathcal{W}(R_F)-1)}$

Proof. Let U denote the event that witness $y \in R_F$ is output by WeightGen on inputs F, ε, r, and X. Let $p_{i,y}$ denote the probability that we exit the loop at line 19 with a particular value of i and $y \in R_{F,h,\alpha}$, where $\alpha \in \{0, 1\}^i$ is the value chosen on line 13. Then, $\mathsf{Pr}[U] = \cup_{i=q-3}^{q} \frac{\mathcal{W}(y)}{\mathcal{W}(Y)} p_{i,y}$, where Y is the set returned by BoundedWeightSAT on line 17. Let $f_m = \mathsf{Pr}[q - 3 \leq m \leq q]$. From Lemma 7, we know that $f_m \geq 0.8$. From line 20, we also know that $\frac{1}{1+\kappa} pivot \leq \mathcal{W}(Y) \leq 1 + (1 + \kappa)pivot$. Therefore,

$\mathsf{Pr}[U] \geq \frac{\mathcal{W}(y)}{1+(1+\kappa)pivot} \cdot p_{m,y} \cdot f_m$. The proof is now completed by showing $p_{m,y} \geq \frac{1}{2^m}(1-e^{-3/2})$, as then we have $\mathsf{Pr}[U] \geq \frac{0.8(1-e^{-3/2})}{(1+(1+\kappa)pivot)2^m} \geq \frac{0.8(1-e^{-3/2})}{(1.06+\kappa)(\mathcal{W}(R_F)|-1)}$. The last inequality uses the observation that $1/pivot \leq 0.06$.

To calculate $p_{m,y}$, we first note that since $y \in R_F$, the requirement "$y \in R_{F,h,\alpha}$" reduces to "$y \in h^{-1}(\alpha)$". For $\alpha \in \{0,1\}^n$, we define $w_{m,y,\alpha} = \mathsf{Pr}\left[\frac{pivot}{1+\kappa} \leq \mathcal{W}(R_{F,h,\alpha}) \leq 1+(1+\kappa)\ pivot \wedge h(y) = \alpha : h \xleftarrow{R} H_{xor}(n,m,3)\right]$. Then we have $p_{m,y} = \Sigma_{\alpha\in\{0,1\}^m}(w_{m,y,\alpha} \cdot 2^{-m})$. So to prove the desired bound on $p_{m,y}$ it suffices to show 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 $\gamma_{z,\alpha} = \mathcal{W}(z)$ if $h(z) = \alpha$, and $\gamma_{z,\alpha} = 0$ otherwise. Let us fix α and choose h uniformly at random from $H_{xor}(n,m,3)$. The random choice of h induces a probability distribution on $\gamma_{z,\alpha}$ such that $\mathsf{E}[\gamma_{z,\alpha}] = \mathcal{W}(z)\,\mathsf{Pr}[\gamma_{z,\alpha} = \mathcal{W}(z)] = \mathcal{W}(z)\,\mathsf{Pr}[h(z) = \alpha] = \mathcal{W}(z)/2^m$. 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 $\mu_\alpha = \mathsf{E}[\Gamma_\alpha]$. Clearly, $\Gamma_\alpha = \mathcal{W}(R_{F,h,\alpha}) - \mathcal{W}(y)$ and $\mu_\alpha = \sum_{z\in\mathcal{W}(R_F)\setminus\{y\}} \mathsf{E}[\gamma_{z,\alpha}] = (\mathcal{W}(R_F) - \mathcal{W}(y))/2^m$. Since $pivot = (\mathcal{W}(R_F) - 1)/2^m \leq (\mathcal{W}(R_F) - \mathcal{W}(y))/2^m$, we have $\mathsf{Pr}[\frac{pivot}{1+\kappa} \leq \mathcal{W}(R_{F,h,\alpha}) \leq 1+(1+\kappa)pivot] \geq \mathsf{Pr}[\frac{\mathcal{W}(R_F)-\mathcal{W}(y)}{(1+\kappa)2^m} \leq \mathcal{W}(R_{F,h,\alpha}) \leq 1+(1+\kappa)\frac{\mathcal{W}(R_F)-1}{2^m}] \geq \mathsf{Pr}[\frac{\mathcal{W}(R_F)-\mathcal{W}(y)}{2^m(1+\kappa)} \leq \mathcal{W}(R_{F,h,\alpha}) - \mathcal{W}(y) \leq (1+\kappa)\frac{(\mathcal{W}(R_F)-\mathcal{W}(y))}{2^m}]$. Since $pivot = \lceil e^{3/2}(1+1/\kappa)^2\rceil$ and the variables $\gamma_{z,\alpha}$ are 2-wise independent and in the range $[0,1]$, we may apply Lemma 3 with $\beta = \kappa/(1+\kappa)$ to obtain $\mathsf{Pr}[\frac{\text{pivot}}{1+\kappa} \leq \mathcal{W}(R_{F,h,\alpha}) \leq 1+(1+\kappa)\text{pivot}] \geq 1-e^{-3/2}$. Since h is chosen at random from $H_{xor}(n,m,3)$, we also have $\mathsf{Pr}[h(y)=\alpha] = 1/2^m$. It follows that $w_{m,y,\alpha} \geq (1-e^{-3/2})/2^m$. □

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

Lemma 9. *For $i < m$, both $w_{i,y,\alpha}$ and $p_{i,y}$ are bounded above by $\frac{1}{\mathcal{W}(R_F)-1}\frac{1}{\left(1-\frac{1+\kappa}{2^{m-i}}\right)^2}$.*

Proof. We will use the terminology introduced in the proof of Lemma 8. Clearly, $\mu_\alpha = \frac{\mathcal{W}(R_F)-\mathcal{W}(y)}{2^i}$. Since each $\gamma_{z,\alpha}$ takes values in $[0,1]$, $\mathsf{V}[\gamma_{z,\alpha}] \leq \mathsf{E}[\gamma_{z,\alpha}]$. Therefore, $\sigma^2_{z,\alpha} \leq \sum_{z\neq y, z\in R_F} \mathsf{E}[\gamma_{z,\alpha}] \leq \sum_{z\in R_F} \mathsf{E}[\gamma_{z,\alpha}] = \mathsf{E}[\Gamma_\alpha] \leq 2^{-m}(\mathcal{W}(R_F) - \mathcal{W}(y))$. So $\mathsf{Pr}[\frac{pivot}{1+\kappa} \leq \mathcal{W}(R_{F,h,\alpha}) \leq 1+(1+\kappa)\text{pivot}] \leq \mathsf{Pr}[\mathcal{W}(R_{F,h,\alpha}) - \mathcal{W}(y) \leq (1+\kappa)\text{pivot}]$. From Chebyshev's inequality, we know that $\mathsf{Pr}[|\Gamma_\alpha - \mu_{z,\alpha}| \geq \lambda\sigma_{z,\alpha}] \leq 1/\lambda^2$ for every $\lambda > 0$. $\mathsf{Pr}[\mathcal{W}(R_{F,h,\alpha}) - \mathcal{W}(y) \leq (1+\kappa)\frac{(\mathcal{W}(R_F)-\mathcal{W}(y))}{2^i}]$ $\leq$ $\mathsf{Pr}\left[|(\mathcal{W}(R_{F,h,\alpha}) - \mathcal{W}(y)) - \frac{\mathcal{W}(R_F)-1}{2^i}| \geq (1-\frac{1+\kappa}{2^{m-i}})\frac{\mathcal{W}(R_F)-\mathcal{W}(y)}{2^i}\right] \leq \frac{1}{\left(1-\frac{(1+\kappa)}{2^{m-i}}\right)^2} \cdot \frac{2^i}{\mathcal{W}(R_F)-1}$. Since h is chosen at random from $H_{xor}(n,m,3)$, we also have $\mathsf{Pr}[h(y) = \alpha] = 1/2^i$. It follows that $w_{i,y,\alpha} \leq \frac{1}{\mathcal{W}(R_F)-1}\frac{1}{\left(1-\frac{1+\kappa}{2^{m-i}}\right)^2}$. The bound for $p_{i,y}$ is easily obtained by noting that $p_{i,y} = \Sigma_{\alpha\in\{0,1\}^i}\left(w_{i,y,\alpha}\cdot 2^{-i}\right)$. □

Lemma 10. *For every witness $y \in R_F$, $\mathsf{Pr}[y \text{ is output}] \leq \frac{(1+\kappa)\mathcal{W}(y)}{\mathcal{W}(R_F)-1}(2.23 + \frac{0.48}{(1-\kappa)^2})$*

Proof. We will use the terminology introduced in the proof of Lemma 8. Using $\frac{pivot}{1+\kappa} \leq \mathcal{W}(Y)$, we have $\mathsf{Pr}[U] = \cup_{i=q-3}^{q}\frac{\mathcal{W}(y)}{\mathcal{W}(Y)}p_{i,y} \leq \frac{1+\kappa}{pivot}\mathcal{W}(y)\sum_{i=q-3}^{q} p_{i,y}$. Now we subdivide the calculation of $\mathsf{Pr}[U]$ into three cases depending on the value of m.

Case 1 : $q-3 \leq m \leq q$.

Now there are four values that m can take.

1. $m = q-3$. We know that $p_{i,y} \leq \mathsf{Pr}[h(y)=\alpha] = \frac{1}{2^i}$, so $\mathsf{Pr}[U|m=q-3] \leq \frac{1+\kappa}{pivot}\cdot\frac{\mathcal{W}(y)}{2^{q-3}}\frac{15}{8}$. Substituting the values of $pivot$ and m gives $\mathsf{Pr}[U|m=q-3] \leq \frac{15(1+\kappa)\mathcal{W}(y)}{8(\mathcal{W}(R_F)-1)}$.
2. $m = q-2$. For $i \in [q-2,q]$ $p_{i,y} \leq \mathsf{Pr}[h(y)=\alpha] = \frac{1}{2^i}$ Using Lemma 9, we get $p_{q-3,y} \leq \frac{1}{\mathcal{W}(R_F)-1}\frac{1}{\left(1-\frac{1+\kappa}{2}\right)^2}$. Therefore, $\mathsf{Pr}[U\ |m=q-2] \leq \frac{1+\kappa}{pivot}\mathcal{W}(y)\frac{1}{\mathcal{W}(R_F)-1}\frac{4}{(1-\kappa)^2} + \frac{1+\kappa}{pivot}\mathcal{W}(y)\frac{1}{2^{q-2}}\frac{7}{4}$. Noting that $pivot = \frac{\mathcal{W}(R_F)-1}{2^m} > 10$, we obtain $\mathsf{Pr}[U|m=q-2] \leq \frac{(1+\kappa)\mathcal{W}(y)}{\mathcal{W}(R_F)-1}(\frac{7}{4} + \frac{0.4}{(1-\kappa)^2})$
3. $m = q-1$. For $i \in [q-1,q]$, $p_{i,y} \leq \mathsf{Pr}[h(y)=\alpha] = \frac{1}{2^i}$. Using Lemma 9, we get $p_{q-3,y} + p_{q-2,y} \leq \frac{1}{\mathcal{W}(R_F)-1}\left(\frac{1}{\left(1-\frac{1+\kappa}{2^2}\right)^2} + \frac{1}{\left(1-\frac{1+\kappa}{2}\right)^2}\right) = \frac{1}{\mathcal{W}(R_F)-1}\left(\frac{16}{(3-\kappa)^2} + \frac{4}{(1-\kappa)^2}\right)$. Therefore, $\mathsf{Pr}[U|m=q-1] \leq \frac{1+\kappa}{pivot}\mathcal{W}(y)\left(\frac{1}{\mathcal{W}(R_F)-1}\left(\frac{16}{(3-\kappa)^2} + \frac{4}{(1-\kappa)^2}\right) + \frac{1}{2^{q-1}}\frac{3}{2}\right)$. Since $pivot = \frac{\mathcal{W}(R_F)-1}{2^m} > 10$ and $\kappa \leq 1$, $\mathsf{Pr}[U|m=q-1] \leq \frac{(1+\kappa)\mathcal{W}(y)}{\mathcal{W}(R_F)-1}(1.9 + \frac{0.4}{(1-\kappa)^2})$.
4. $m = q$. We have $p_{q,y} \leq \mathsf{Pr}[h(y)=\alpha] = \frac{1}{2^q}$, and using Lemma 9 we get $p_{q-3,y} + p_{q-2,y} + p_{q-1,y} \leq \frac{1}{\mathcal{W}(R_F)-1}\left(\frac{1}{\left(1-\frac{1+\kappa}{2^3}\right)^2} + \frac{1}{\left(1-\frac{1+\kappa}{2^2}\right)^2} + \frac{1}{\left(1-\frac{1+\kappa}{2}\right)^2}\right) = \frac{1}{\mathcal{W}(R_F)-1}\left(\frac{64}{(7-\kappa)^2} + \frac{16}{(3-\kappa)^2} + \frac{4}{(1-\kappa)^2}\right)$. So $\mathsf{Pr}[U|m=q] \leq \frac{1+\kappa}{pivot}\mathcal{W}(y)\left(\frac{1}{\mathcal{W}(R_F)-1}\left(\frac{64}{(7-\kappa)^2} + \frac{16}{(3-\kappa)^2} + \frac{4}{(1-\kappa)^2}\right) + 1\right)$. Using $pivot = \frac{\mathcal{W}(R_F)-1}{2^m} > 10$ and $\kappa \leq 1$, we obtain $\mathsf{Pr}[U|m=q] \leq \frac{(1+\kappa)\mathcal{W}(y)}{\mathcal{W}(R_F)-1}(1.58 + \frac{0.4}{(1-\kappa)^2})$.

Since $\mathsf{Pr}[U|q-3 \leq m \leq q] \leq \max_{q-3\leq i\leq q}(\mathsf{Pr}[U|m=i])$, we have $\mathsf{Pr}[U|q-3 \leq m \leq q] \leq \frac{1+\kappa}{\mathcal{W}(R_F)-1}(1.9 + \frac{0.4}{(1-\kappa)^2})$ from the $m = q-1$ case above.

Case 2 : $m < q-3$. Since $p_{i,y} \leq \mathsf{Pr}[h(y)=\alpha] = \frac{1}{2^i}$, we have $\mathsf{Pr}[U|m<q-3] \leq \frac{1+\kappa}{pivot}\mathcal{W}(y) \cdot \frac{1}{2^{q-3}}\frac{15}{8}$. Substituting the value of $pivot$ and maximizing $m-q+3$, we get $\mathsf{Pr}[U|m<q-3] \leq \frac{15(1+\kappa)\mathcal{W}(y)}{16(\mathcal{W}(R_F)-1)}$.

Case 3 : $m > q$. Using Lemma 9, we know that $\mathsf{Pr}[U|m>q] \leq \frac{1+\kappa}{pivot}\frac{\mathcal{W}(y)}{\mathcal{W}(R_F)-1}\sum_{i=q-3}^{q}\frac{1}{\left(1-\frac{1+\kappa}{2^{m-i}}\right)^2}$. The R.H.S. is maximized when $m = q+1$. Hence $\mathsf{Pr}[U|m>q] \leq \frac{1+\kappa}{pivot}\frac{\mathcal{W}(y)}{\mathcal{W}(R_F)-1}\times\sum_{i=q-3}^{q}\frac{1}{\left(1-\frac{1+\kappa}{2^{q+1-i}}\right)^2}$. Noting that $pivot = \frac{\mathcal{W}(R_F)-1}{2^m} > 10$ and expanding the above summation we have $\mathsf{Pr}[U|m>q] \leq \frac{(1+\kappa)\mathcal{W}(y)}{\mathcal{W}(R_F)-1}\frac{1}{10}\left(\frac{256}{(15-\kappa)^2}+\frac{64}{(7-\kappa)^2}+\frac{16}{(3-\kappa)^2}+\frac{2}{(1-\kappa)^2}\right)$. Using $\kappa \leq 1$ for the first three summation terms, we obtain $\mathsf{Pr}[U|m>q] \leq \frac{(1+\kappa)\mathcal{W}(y)}{\mathcal{W}(R_F)-1}(0.71+\frac{0.4}{(1-\kappa)^2})$

Summing up all the above cases, $\mathsf{Pr}[U] = \mathsf{Pr}[U|m<q-3]\times\mathsf{Pr}[m<q-3]+\mathsf{Pr}[U|q-3\leq m\leq q]\times\mathsf{Pr}[q-3\leq m\leq q]+\mathsf{Pr}[U|m>q]\times\mathsf{Pr}[m>q]$. From Lemma 7 we have $\mathsf{Pr}[m<q-1] \leq 0.2$ and $\mathsf{Pr}[m>q] \leq 0.2$, so $\mathsf{Pr}[U] \leq \frac{(1+\kappa)\mathcal{W}(y)}{\mathcal{W}(R_F)-1}(2.23+\frac{0.48}{(1-\kappa)^2})$

□

Combining Lemmas 8 and 10, the following lemma is obtained.

Lemma 11. *For every witness $y \in R_F$, if $\varepsilon > 1.71$, then* $\frac{w(y)}{(1+\varepsilon)w(R_F)} \leq \mathsf{Pr}\left[\mathsf{WeightGen}(F,\varepsilon,r,X)=y\right] \leq (1+\varepsilon)\frac{w(y)}{w(R_F)}$.

Proof. In the case where $\mathcal{W}(R_F) \leq 1+(1+\kappa)pivot$, the result holds because WeightGen returns a perfect weighted-uniform sample. Otherwise, using Lemmas 8 and 10 and substituting $(1+\varepsilon) = (1+\kappa)(2.36+\frac{0.51}{(1-\kappa)^2}) = \frac{18}{17}(1+\kappa)(2.23+\frac{0.48}{(1-\kappa)^2})$, via the inequality $\frac{1.06+\kappa}{0.8(1-e^{-3/2})} \leq \frac{18}{17}(1+\kappa)(2.23+\frac{0.48}{(1-\kappa)^2})$ we have the bounds $\frac{\mathcal{W}(y)}{(1+\varepsilon)(\mathcal{W}(R_F)-1)} \leq \mathsf{Pr}\left[\mathsf{WeightGen}(F,\varepsilon,r,X)=y\right] \leq \frac{18}{17}(1+\varepsilon)\frac{\mathcal{W}(y)}{\mathcal{W}(R_F)-1}$. Using $\mathcal{W}(R_F) \geq 18$, we obtain the desired result. □

Lemma 12. *Algorithm* WeightGen *succeeds (i.e. does not return $\perp$) with probability at least* 0.62.

Proof. If $\mathcal{W}(R_F) \leq 1+(1+\kappa)$pivot, the theorem holds trivially. Suppose $\mathcal{W}(R_F) > 1+(1+\kappa)$pivot and let P_{succ} denote the probability that a run of the algorithm succeeds. Let p_i with $q-3 \leq i \leq q$ denote the conditional probability that WeightGen (F, ε, r, X) terminates in iteration i of the repeat-until loop (lines 15–19) with $\frac{pivot}{1+\kappa} \leq \mathcal{W}(R_{F,h,\alpha}) \leq 1+(1+\kappa)pivot$, given that $\mathcal{W}(R_F) > 1+(1+\kappa)pivot$. Then $P_{\text{succ}} = \sum_{i=q-3}^{q} p_i \prod_{j=q-3}^{i}(1-p_j)$. Letting $f_m = \mathsf{Pr}[q-3\leq m\leq q]$, by Lemma 7 we have $P_{\text{succ}} \geq p_m f_m \geq 0.8 p_m$. The theorem is now proved by using Lemma 3 to show that $p_m \geq 1-e^{-3/2} \geq 0.776$.

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

Let $\Gamma_\alpha = \sum_{y\in R_F}\nu_{y,\alpha}$ and $\mu_\alpha = \mathsf{E}[\Gamma_\alpha]$. Clearly, $\Gamma_\alpha = \mathcal{W}(R_{F,h,\alpha})$ and $\mu_\alpha = \sum_{y\in R_F}\mathsf{E}[\nu_{y,\alpha}] = 2^{-m}\mathcal{W}(R_F)$. Since $pivot = \lceil e^{3/2}(1+1/\epsilon)^2\rceil$, we have $2^{-m}\mathcal{W}(R_F) \geq e^{3/2}(1+1/\varepsilon)^2$, and so using Lemma 3 with $\beta = \kappa/(1+\kappa)$ we obtain $\mathsf{Pr}\left[\frac{\mathcal{W}(R_F)}{2^m}\cdot\left(1-\frac{\kappa}{1+\kappa}\right) \leq \mathcal{W}(R_{F,h,\alpha}) \leq (1+\frac{\kappa}{1+\kappa})\frac{\mathcal{W}(R_F)}{2^m}\right] > 1-e^{-3/2}$. Simplifying and noting that $\frac{\kappa}{1+\kappa} < \kappa$ for all $\kappa > 0$, we have $\mathsf{Pr}\left[(1+\kappa)^{-1}\cdot\frac{\mathcal{W}(R_F)}{2^m} \leq \mathcal{W}(R_{F,h,\alpha}) \leq (1+\kappa)\cdot\frac{\mathcal{W}(R_F)}{2^m}\right] > 1-e^{-3/2}$. Also, $\frac{pivot}{1+\kappa} = \frac{1}{1+\kappa}\frac{\mathcal{W}(R_F)-1}{2^m} \leq \frac{\mathcal{W}(R_F)}{(1+\kappa)2^m}$ and $1+(1+\kappa)pivot = 1+\frac{(1+\kappa)(\mathcal{W}(R_F)-1)}{2^m} \geq \frac{(1+\kappa)\mathcal{W}(R_F)}{2^m}$. Therefore, $p_m = \mathsf{Pr}[\frac{pivot}{1+\kappa} \leq \mathcal{W}(R_{F,h,\alpha}) \leq 1+(1+\kappa)pivot] \geq \mathsf{Pr}\left[(1+\kappa)^{-1}\cdot\frac{\mathcal{W}(R_F)}{2^m} \leq \mathcal{W}(R_{F,h,\alpha}) \leq (1+\kappa)\cdot\frac{\mathcal{W}(R_F)}{2^m}\right] \geq 1-e^{-3/2}$. □

By combining Lemmas 11 and 12, we get the following:

Theorem 3. *Given a CNF formula F, tolerance $\varepsilon > 1.71$, tilt bound r, and independent support S, for every $y \in R_F$ we have* $\frac{w(y)}{(1+\varepsilon)w(R_F)} \leq \mathsf{Pr}\left[\mathsf{WeightGen}(F,\varepsilon,r,X)=y\right] \leq (1+\varepsilon)\frac{w(y)}{w(R_F)}$. *Also,* WeightGen *succeeds (i.e. does not return $\perp$) with probability at least* 0.62.

Theorem 4. *Given an oracle for* SAT, $\mathsf{WeightGen}(F,\varepsilon,r,S)$ *runs in time polynomial in $r, |F|$ and $1/\varepsilon$ relative to the oracle.*

Proof. Referring to the pseudocode for WeightGen, the runtime of the algorithm is bounded by the runtime of the constant number (at most 5) of calls to BoundedWeightSAT and one call to WeightMC (with parameters $\delta = 0.2, \varepsilon = 0.8$). As shown in Theorem 1, the call to WeightMC can be done in time polynomial in $|F|$ and r relative to the oracle. Every invocation of BoundedWeightSAT can be implemented by at most $(r\cdot pivot)+1$ calls to a SAT oracle (as in the proof of Theorem 2), and the total time taken by all calls to BoundedWeightSAT is polynomial in $|F|, r$ and $pivot$ relative to the oracle. Since $pivot = \mathcal{O}(1/\varepsilon^2)$, the runtime of WeightGen is polynomial in r, $|F|$ and $1/\varepsilon$ relative to the oracle. □

10 Analysis of Partitioned WeightMC

Theorem 5. *If* $\mathsf{PartitionedWeightMC}(F,\varepsilon,\delta,S,L,H)$ *returns c (and all arguments are in the required ranges), then*

$$\Pr\left[c \neq \perp \wedge (1+\varepsilon)^{-1}w(R_F) \leq c \leq (1+\varepsilon)w(R_F))\right] \geq 1-\delta.$$

Proof. For future reference note that since $N \geq 1$ and $\delta < 1$, we have $(1-\delta')^N = (1-\delta/N)^N \geq 1-\delta$. Define $G_m = F \wedge (H/2^m < w(X) \leq H/2^{m-1})$, the formula passed to WeightMC in iteration m. Clearly, we have $w(R_F) = \sum_{m=1}^{N} w(R_{G_m})$. Since $w(\cdot)$ is poly-time computable, the NP oracle used in WeightMC can decide the satisfiability of G_m, and so WeightMC will return a value d_m. Now since $H/2^m$ and $H/2^{m-1}$ are lower and upper bounds respectively on the weights of any solution to G_m, by Theorem 1 we have

$$\Pr\left[d_m \neq \perp \wedge (1+\varepsilon)^{-1} w(R_{G_m}) \leq d_m \leq (1+\varepsilon) w(R_{G_m})\right] \geq 1-\delta'$$

for every m, and so

$$\begin{aligned}
&\Pr\left[c \neq \perp \wedge (1+\varepsilon)^{-1} w(R_F) \leq c \leq (1+\varepsilon) w(R_F)\right] \\
&= \Pr\left[c \neq \perp \wedge (1+\varepsilon)^{-1} \sum_m w(R_{G_m}) \leq c \leq (1+\varepsilon) \sum_m w(R_{G_m})\right] \\
&\geq (1-\delta')^N \geq 1-\delta
\end{aligned}$$

as desired. □

Theorem 6. *With access to an* NP *oracle, the runtime of* PartitionedWeightMC$(F, \varepsilon, \delta, S, L, H)$ *is polynomial in* $|F|$, $1/\varepsilon$, $\log(1/\delta)$, *and* $\log r = \log(H/L)$.

Proof. Put $r = H/L$. By Theorem 2, each call to WeightMC runs in time polynomial in $|G|$, $1/\varepsilon$ and $\log(1/\delta')$ (the tilt bound is constant). Clearly $|G|$ is polynomial in $|F|$. Since $\delta' = \delta/N$ we have $\log(1/\delta') = \log(N/\delta) = \mathcal{O}(\log((\log_k r)/\delta)) = \mathcal{O}(\log\log r + \log(1/\delta))$. Therefore each call to WeightMC runs in time polynomial in $|F|$, $1/\varepsilon$, $\log(1/\delta)$, and $\log\log r$. Since there are $N = \mathcal{O}(\log r)$ calls, the result follows. □