Deterministic Volume Estimation of Truncated Hypercubes

Kyra Gunluk

April 2, 2026

Abstract

Deterministic Volume Computation.

1 Introduction

Computing the volume is an ancient and difficult problem, even for convex bodies. Dyer and Frieze [DF88a] showed that even for an explicit polyhedron specified by a unimodular constraint matrix, computing the volume is #P-hard. In fact, even approximating the volume of a convex body is notoriously difficult. In the general membership oracle model, it was shown by [Ele86, BF87, FB88] that any deterministic algorithm that computes a polynomial (in the dimension) relative approximation of the volume of a convex body must incur exponential complexity. Against this backdrop, the randomized polynomial-time approximation scheme of Dyer, Frieze and Kannan [DFK89] was a surprising breakthrough that heralded an age of new techniques for randomized algorithms and analysis for volume computation that have also been extended to logconcave integration [Lov90, AK91, DF91, LS90, LS93, KLS95, KLS97, LV06, LV07, KV06, CV18, KV25].

Our interest here is in deterministic algorithms when the input is specified explicitly (and therefore the membership oracle model lower bounds do not apply). Even for explicitly specified polyhedra, the known efficient algorithms are essentially the same as in the membership oracle model and rely heavily on randomization — they are based on sampling a sequence of distributions by Markov chains. If P=BPP, one expects a polytime deterministic algorithm in the Turing Machine model. The main motivation for the present paper is to understand classes of convex bodies whose volume can be estimated by efficient deterministic algorithms.

Going back to the original #P-hardness proof of Dyer and Frieze [DF88b], the core problem shown to be hard is very simple: Let P be the polytope obtained by intersecting the unit hypercube $[0,1]^n$ with a *single* halfspace with a nonnegative constraint vector, $a^\top x \le b$. Computing the volume of this family of polytopes, equivalently, computing the probability that a random point in a unit hypercube satisfies a given linear constraint, is already #P-hard!

Theorem 1.1 (Polytope Volume is#P-hard [DF88a]). *Let $P(A,b) = \{x \in \mathbb{R}^n : Ax \leq b\}$ be a polyhedron. A, b have rational entries where $A = (a_{ij})$, $i = 1, 2, \ldots, m$, $j = 1, 2, \ldots, n$, and $b = (b_i)$, $i = 1, 2, \ldots, m$. Computing $\mathrm{vol}(P(A,b))$ is #P-hard even when A is totally unimodular; it is also #P-hard for $P = [0,1]^n \cap \{x : a^\top x \leq b\}$ for a given integer vector a.*

1.1 Results

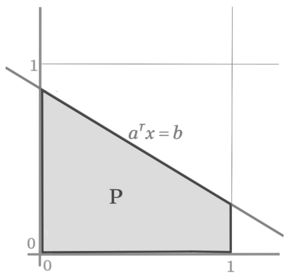

Figure 1.1

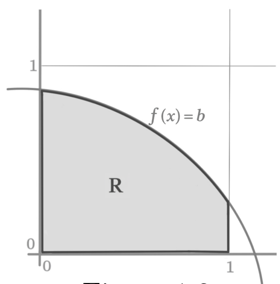

Figure 1.2

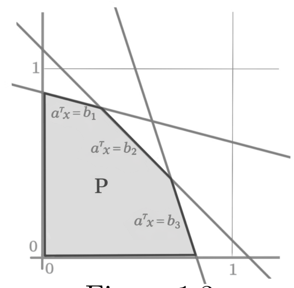

Figure 1.3

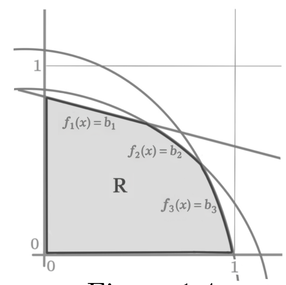

Figure 1.4

Our first result is a deterministic fully polynomial-time approximation scheme for a hypercube truncated by a single arbitrary halfspace, as depicted in Figure 1.1.

Theorem 1.2. *Given $a \in \mathbb{R}^n, b \in \mathbb{R}, \varepsilon > 0$, we can find $Z' \in \mathbb{R}_{\geq 0}$ approximating the volume of polytope $P = [0,1]^n \cap \{x : a^\top x \leq b\}$ s.t.*

$$\mathsf{vol}(P) \leq Z' \leq (1+\varepsilon)\mathsf{vol}(P)$$

using $O\left(\frac{n^3}{\varepsilon}(\log\frac{n}{\varepsilon} + L)^4\right)$ arithmetic operations on $O(L + \log n)$-bit numbers, where $L \in \mathbb{R}_{\geq 0}$ is the max bit size of any input number.

While there are known formulas giving the volume of a hypercube clipped by a single hyperplane [BS79], they rely on exponentially many arithmetic operations.

Our next result is the more general setting of a hypercube truncated by a decomposable convex constraint, as depicted in Figure 1.2. These special constraints include p-norm balls $\{x : \|x\|_p \leq b\}$ for any $b \geq 0$.

Theorem 1.3. *Given nondecreasing, univariate, convex functions $f_j : \mathbb{R} \to \mathbb{R}_{\geq 0}$ for $j \in [n]$, $b \in \mathbb{R}, \varepsilon > 0$, we can find $Z' \in \mathbb{R}_{\geq 0}$ approximating the volume of $Q = [0,1]^n \cap \{x : \sum_{j=1}^n f_j(x_j) \leq b\}$ s.t.*

$$\mathsf{vol}(Q) \leq Z' \leq (1+\varepsilon)\mathsf{vol}(Q)$$

using $O(n^5(L_o + \log(\frac{n}{\varepsilon}))^2 L/\varepsilon)$ arithmetic operations on $O(L + \log n)$-bit numbers, where $L \in \mathbb{R}_{\geq 0}$ is the max input bit size and $L_o = \log \mathsf{vol}(Q)$.

We then extend these results to a fixed number of constraints, as depicted in Figure 1.3. We give a deterministic fully polynomial-time approximation scheme for a hypercube truncated by k

halfspaces defined by normals with nonnegative coefficients. The complexity is polynomial in the dimension for any fixed k.

Theorem 1.4. *Given $A \in \mathbb{R}_{\geq 0}^{k \times n}$ and $b \in \mathbb{R}_{\geq 0}^{k}$, $\varepsilon > 0$, we can find $Z' \in \mathbb{R}_{\geq 0}$ approximating the volume of polytope $P = [0,1]^n \cap \{x : Ax \leq b\}$ s.t.*

$$\mathsf{vol}(P) \leq Z' \leq (1+\varepsilon)\mathsf{vol}(P)$$

in $n^{O(k^2)}(\log \frac{n}{\varepsilon} + L)^{O(k)}/\varepsilon^{O(k)}$ arithmetic operations on $O(L + \log n)$-bit numbers, where $L \in R_{\geq 0}$ is the max input bit size.

Note that there is also a formula for the volume of the cube intersected with multiple halfspaces, extended by [CK15]. Again, this is an exponential computation.

Our last result is an extension to the intersection of a hypercube with multiple decomposable *convex* constraints, as depicted in Figure 1.4.

Theorem 1.5. *Given nondecreasing convex functions $f_{ij} : \mathbb{R} \to \mathbb{R}_{\geq 0}$ for $i \in [k]$, $j \in [n]$, $b \in \mathbb{R}^n_{\geq 0}$, $\varepsilon > 0$, we can find $Z' \in \mathbb{R}_{\geq 0}$ approximating the volume of $Q = [0,1]^n \cap \{x : \sum_{j=1}^n f_{ij}(x_j) \leq b_i \quad \forall i \in [k]\}$ s.t.*

$$\mathsf{vol}(Q) \leq Z' \leq (1+\varepsilon)\mathsf{vol}(Q)$$

in $n^{O(k^2)}(\log \frac{n}{\varepsilon} + L + L_o)^{O(k)}/\varepsilon^{O(k)}$ arithmetic operations on $O(L + L_o + \log n)$-bit numbers, where $L \in \mathbb{R}_{\geq 0}$ is the max input bit size and L_o is the output bit size of $\mathsf{vol}(Q)$.

[Cos24] gives a quasi-polynomial deterministic approximation algorithm for computing the volume of a hypercube intersected with two sets. Both of these sets are of the form $\{x : \sum_{i=1}^n f_i(x_i) \leq b\}$ where each f_i is a polynomial function rather than our requirement of nonnegative, nondecreasing, and convex. [GN24] gives an FPTAS for determining the volume of the truncated fractional matching polytope for a hypergraph with bounded maximum degree. [NS26] gives a polynomial-time deterministic algorithm for approximating the volume of the Kostka polytope within an exponential multiplicative factor. [BR24] gives a deterministic algorithm for approximating the volume of the nonnegative orthant intersected with m hyperplanes. While the runtime of this algorithm is polynomial, the approximation factor is exponential in the number of constraints.

Our results are the first deterministic polynomial-time approximation schemes for volume for truncated hypercubes; they generalize and improve the complexity of several special cases studied in the literature.

1.2 Techniques

The algorithms in this paper rely on choosing a lattice for which the fraction of lattice points within our polytope (or convex body) P approximates its volume $\mathsf{vol}(P)$. We find such a lattice by first scaling our polytope P by a factor of $u \in \mathbb{Z}_+$ (it now lies in the $[0,u]^n$ cube) and using points with integer coordinates. For sufficiently large u, we can guarantee that the number of integer points is a good approximation of $u^n \cdot \mathsf{vol}(P)$. We then adapt known fully polynomial-time approximation schemes for counting integer solutions under read once constraints.

Our algorithm for Theorem 1.2 relies on counting knapsack solutions. We call on the following result by [SVV12].

Theorem 1.6 (Stefankovic-Vempala-Vigoda)**.** *Let $w_1, \ldots, w_n$ and C be an instance of a knapsack problem. Let Z be the number of solutions of the knapsack problem. There is a deterministic algorithm COUNT_KNAPSACK(ε) which, for any $\varepsilon \in (0,1)$, outputs Z' such that*

$$(1-\varepsilon)Z \leq Z' \leq Z.$$

The algorithm runs in time $O\big(n^3\varepsilon^{-1}\log(n/\varepsilon)\big)$.

Our algorithm for Theorem 1.3 relies on counting integer solutions under a read-once constraint. We call on the following result by [GKM10].

Theorem 1.7 (Extension to Theorem 1.3 in [GKM10])**.** *Given set $S(f,b,u) = \{x \in \{0,1,\ldots,u-1\}^n : \sum_{j=1}^n f_j(x_j) \leq b\}$ and $\delta > 0$, there is a deterministic algorithm ROUND_ROBP(δ) that computes an δ relative error approximation for $|S(f,b,u)|$ in $O(n^5(\log u)^2(\log b)/\delta)$ time.*

For Theorems 1.4 and 1.5 our FPTAS to count integer points relies on adapted theorems from [GKM10].

Lemma 1.8 (Extension to [Dye03])**.** *Let $U_n = \{0,\ldots,u_1\} \times \cdots \times \{0,\ldots,u_n\}$. Consider set $S = \cap_{i\in[k]} S_i$ where $S_i = \{x \in U_n : \sum_{j=1}^n f_{ij}(x_j) \leq b_i\}$ where each $f_{ij} : \mathbb{R} \to \mathbb{R}_{\geq 0}$ is convex and nondecreasing. There exists a set*

$$S' = \cap_{i\in[k]} S'_i \quad \textit{where} \quad S'_i = \{x \in U_n : \sum_{j=1}^n \lfloor \frac{2n^2 f_{ij}(x_j)}{b_i} \rfloor \leq 2n^2\}$$

such that $|S| \leq |S'| \leq 2n^k|S|$.

Theorem 1.9 (Extension to Theorem 1.2 of [GKM10])**.** *Let $S = \cap_{i\in[k]} S(f_i, b_i, u)$ where $S(f_i,b_i,u) = \{x \in \{0,\ldots,u_1\} \times \cdots \times \{0,\ldots,u_n\} : \sum_{j=1}^n f_{ij}(x_j) \leq b_i\}$ where each $f_{ij} : \mathbb{R} \to \mathbb{R}_{\geq 0}$ is convex and non-decreasing. There is a deterministic algorithm ROUND_ROBPS(δ) which, for any $\delta > 0$, outputs V such that*

$$|S| \leq V \leq (1+\delta)|S|$$

in $n^{O(k^2)}(\frac{\log U}{\varepsilon})^{O(k)}\log b$ time where $U = \max_{j\in[n]} u_j$ and $b = \max_{i\in[k]} b_i$.

The assumption that the normals to the truncating hyperplanes have nonnegative coefficients is necessary for the techniques we use, i.e., one cannot extend the method of approximating volume through approximately counting integer points when coefficients are allowed to be negative. This is because there is no multiplicative approximation for the cardinality of such a set, unless $P = NP$. More precisely, we have the following result.

Theorem 1.10. *Given $a \in \mathbb{Z}^n, b \in \mathbb{Z}$, determining whether the set $S = \{x \in \{0,1\}^n : -b \le a^\top x \le b\}$ is empty is NP-hard and hence $|S|$ cannot be approximated to within any multiplicative factor unless P=NP.*

2 Preliminaries

The following definitions and propositions will be referenced throughout the paper.

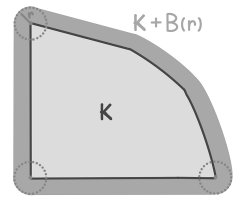

Figure 2.1

Definition 2.1 (Rounded Convex Body)**.** *For a convex body K, $r \in \mathbb{R}_{\ge 0}$, and the Euclidean distance metric $d : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}_{\ge 0}$, $K + B(r)$ is the convex body $\{y : \exists x \in K \text{ with } d(x,y) \le r\}$. See figure 2.1.*

Definition 2.2 (Dilated Convex Body)**.** *Given a convex body K and $r \in \mathbb{R}_{\ge 0}$, rK is the convex body $\{rx : x \in K\}$. Note that $\mathsf{vol}(rK) = r^n \mathsf{vol}(K)$.*

Definition 2.3 (Closure)**.** *Given a convex body K, its closure $cl(K)$ is defined as the intersection of all closed bodies containing K. Equivalently, it is the smallest closed body containing K.*

Definition 2.4 (Nondecreasing Function)**.** *A function $f : \mathbb{R}^n \to \mathbb{R}$ is said to be* nondecreasing *if for every $x, y \in \mathbb{R}^n$, we have that $\forall j \in [n], x_j \le y_j$ implies that $f(x) \le f(y)$.*

Definition 2.5. *We use $\mathbb{1}_n$ to refer to the length n vector where each entry is a 1. Similarly, $\mathbb{0}_n$ refers to the length n vector where each entry is a 0.*

Proposition 2.1. *Let $f : \mathbb{R}_{\ge 0} \to \mathbb{R}_{\ge 0}$ be a nondecreasing and convex function. Take $h(x) = f(x) - f(0)$. Then h is also nonnegative, nondecreasing and convex.*

Proof of Proposition 2.1. Since $f(x)$ is nondecreasing, we have $f(0) \le f(x)$ for all $x \ge 0$. Thus, $h(x) = f(x) - f(0) \ge 0$. Also, h is nondecreasing and convex since f_j is nondecreasing and convex, and h subtracts the same number $f(0)$ on every input. □

Proposition 2.2. *Let $b \in R_{\ge 0}$, $u \in \mathbb{Z}_+$, and $f_{ij} : \mathbb{R} \to \mathbb{R}_{\ge 0}$ be nonnegative, nondecreasing, and convex functions for $i \in [k], j \in [n]$. Let $K = \{x \in [0,u)^n : \sum_{j=1}^n f_{ij}(\frac{x}{u}) \le b_i \quad \forall i \in [k]\}$. Consider the axis-aligned partition of $[0,u)^n$ into unit cubes, and let C be the subset of these cubes that intersect K. Let $\ell = \min\{u, \min_{i \in [k]} \min_{j \in [n]} u f_{ij}^{-1}(b_i)\}$ be the smallest axis-intercept of $cl(K)$ on any axis. Then*

$$\mathsf{vol}(K) \le \mathsf{vol}(C) \le \left(1 + \frac{2n\sqrt{n}}{\ell}\right)^n \mathsf{vol}(K).$$

Proof of Proposition 2.2. Let $Z = K \cap \mathbb{Z}^n$ be the integer points in K and $C_z = \{x : z \le x < z + \mathbb{1}_n\}$. Then, $C = \cup_{z \in Z} C_z$. See Figure 2.2 for an example. We will prove that

$$K \subseteq C \subseteq K + B(\sqrt{n}) \subseteq (1 + \frac{\sqrt{n}}{r})\text{cl}(K),$$

where $r = \ell/2n$.

We first ague that $K \subseteq C$: Consider any $x \in K$, and let $\bar{x}$ be such that $\bar{x}_i = \lfloor x_i \rfloor$ for all $i \in [n]$. Note that $\bar{x} \leq x < \bar{x} + \mathbb{1}_n$. Since $x \in [0,u)^n$, it follows that $\bar{x} \in [0, u-1]^n \subseteq [0,u]^n$. By the nondecreasing property of each f_{ij}, we have $\sum_{j=1}^n f_{ij}(\bar{x}/u) \leq \sum_{j=1}^n f_{ij}(x/u) \leq b_i$ for all i. Finally, $\bar{x} \in \mathbb{Z}^n$ by construction, and so $\bar{x} \in Z$. Thus $x \in C_{\bar{x}} \subseteq C$.

Next we argue that $C \subseteq K + B(\sqrt{n})$:

Consider any $x \in C$, and take $\bar{x} \in Z \subseteq K$ such that $x \in C_{\bar{x}}$. Note that since $C_{\bar{x}}$ has length 1, the furthest distance within this cube is $\sqrt{n}$, and so the distance between x and $\bar{x}$ (a point in K) is at most $\sqrt{n}$. Thus $x \in K + B(\sqrt{n})$.

Now we argue that $K + B(\sqrt{n}) \subseteq (1 + \sqrt{n}/r)\mathrm{cl}(K)$. It is clear that $K \subseteq (1 + \sqrt{n}/r)K \subseteq (1 + \sqrt{n}/r)\mathrm{cl}(K)$, but we must also show this for $K + B(\sqrt{n}) \setminus K$.

First, note that if $\mathbb{0}_n \notin K$ then $b_i < \sum_{j=1}^n f_{ij}(x/u) \leq \sum_{j=1}^n f_{ij}(x/u)$ for all $x \in [0,u)^n$, and so $K = \emptyset$. In this case, the claim follows clearly, so we assume that $0 \in K$. In addition, K is the intersection of level sets of convex functions and a convex body, so it is convex. Recall that $\ell = \min\{u, \min_{i \in [k]} \min_{j \in [n]} u f_{ij}^{-1}(b_i)\}$ is the smallest x_j-intercept, and so the entire x_j axis between 0 and ℓ lies within $\mathrm{cl}(K)$ for every $j \in [n]$.

These facts imply that K contains the simplex $S := \{x : \sum_{j=1}^n x_j \leq \ell\}$. Consider the ball $B(z, r)$ where $r = \ell/2n$ and $z_j = \ell/2n$ for every $j \in [n]$. See figure 2.3 for these objects. Note that for any $\bar{x} \in B(z,r)$, $\bar{x}_j \leq \frac{\ell}{n}$ for all $j \in [n]$, and so $\sum_{j=1}^n \bar{x}_j \leq \sum_{j=1}^n \ell/n = \ell$. Thus $\bar{x} \in S \subseteq \mathrm{cl}(K)$.

Now consider $x \in K + B(\sqrt{n}) \setminus K$, and let x' be the point in $\mathrm{cl}(K)$ that is closest to x. Note that x' will be a boundary point of $\mathrm{cl}(K)$. By definition of $K + B(\sqrt{n})$, $d(x, x') \leq \sqrt{n}$, where $d : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is Euclidean distance. By triangle inequality

$$d(x,z) \leq d(x',z) + d(x,x') \text{ and thus } \frac{d(x,z)}{d(x',z)} \leq \frac{d(x',z) + \sqrt{n}}{d(x',z)} = 1 + \frac{\sqrt{n}}{d(x',z)} \leq 1 + \frac{\sqrt{n}}{r}.$$

where the last inequality comes from the fact that $B(z,r)$ lies fully within K, and thus boundary point $x' \in K$ is at least distance r from z. See figure 2.4 for an example.

Thus, $K \subseteq C \subseteq (1 + \sqrt{n}/r)\mathrm{cl}(K)$ and consequently

$$\mathsf{vol}(K) \leq \mathsf{vol}(C) \subseteq \left(1 + \frac{\sqrt{n}}{r}\right)^n \mathsf{vol}(\mathrm{cl}(K)) = \left(1 + \frac{\sqrt{n}}{r}\right)^n \mathsf{vol}(K).$$

□

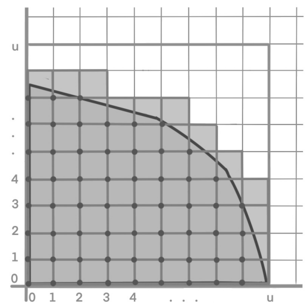

Figure 2.2

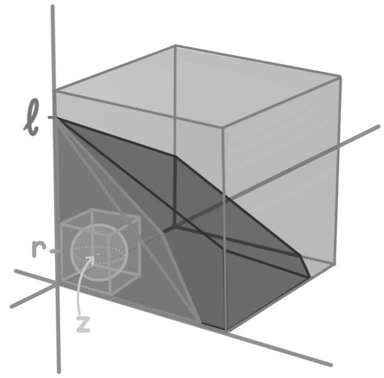

Figure 2.3

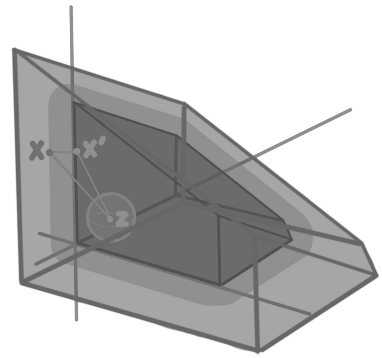

Figure 2.4

Lemma 2.3. *Let $b \in \mathbb{R}^k_{\geq 0}$ and let $f_{ij} : \mathbb{R} \to \mathbb{R}_{\geq 0}$ be nondecreasing, nonnegative and convex functions for $i \in [k], j \in [n]$. Let*

$$\begin{aligned} u &\geq \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{i\in[k],j\in[n]} \frac{1}{f_{ij}^{-1}(b_i)}\} \\ K &= \{x \in [0,u)^n : \sum_{j=1}^n f_{ij}(\frac{x_j}{u}) \leq b_i \quad \forall i \in [k]\} \\ Z &= \mathbb{Z}^n \cap K. \end{aligned}$$

Then, for any $\varepsilon > 0$, for V such that $|Z| \leq V \leq (1+\frac{\varepsilon}{9})|Z|$, we have $\mathsf{vol}(K) \leq V \leq (1+\varepsilon)\mathsf{vol}(K)$.

Proof of Lemma 2.3. Consider points Z which are all the integer points in K. For any $z \in Z$, we will consider the length 1 cube rooted at this point: $C_z = \{x : z \leq x < z + \mathbb{1}_n\}$.

Note that $C = \cup_{z\in Z} C_z$ is a subset of the partition of $[0,u]^n$ into length 1 cubes with integer vertices. Thus,

$$\mathsf{vol}(C) = \mathsf{vol}(\cup_{z\in Z} C_z) = \sum_{z\in Z} \mathsf{vol}(C_z) = \sum_{z\in Z}(1^n) = |Z|.$$

Take u as in the statement of the lemma and $\ell = \min\{u, \min_{i\in[k]} \min_{j\in[n]} u f_{ij}^{-1}(b_i)\}$. Note that ℓ is exactly the smallest axis-intercept of $\mathrm{cl}(K)$. We can apply Prop. 2.2 to get that

$$\mathsf{vol}(K) \leq \mathsf{vol}(C) \leq \left(1 + \frac{2n\sqrt{n}}{\ell}\right)^n \mathsf{vol}(K)$$

From this, we have

$$\mathsf{vol}(K) \leq \mathsf{vol}(C) = |Z| \leq V \leq (1+\frac{\varepsilon}{9})|Z| = (1+\frac{\varepsilon}{9})\mathsf{vol}(C) \leq (1+\frac{\varepsilon}{9})\left(1 + \frac{2n\sqrt{n}}{\ell}\right)^n \mathsf{vol}(K)$$

We explore two cases for $\ell = \min\{u, \min_{i\in[k]} \min_{j\in[n]} u f_{ij}^{-1}(b_i)\}$ and show that $\ell \geq 9n^{2.5}/\varepsilon$.

Case 1: $\ell = u \geq \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{i\in[k],j\in[n]} \frac{1}{f_{ij}^{-1}(b_i)}\} \geq \frac{9n^{2.5}}{\varepsilon}$.

Case 2: $\ell = \min_{i\in[k]} \min_{j\in[n]} uf_{ij}^{-1}(b_i)$. From the assumption on u, it follows that $u \geq \frac{9n^{2.5}}{\varepsilon} \frac{1}{f_{ij}^{-1}(b_i)}$ $\forall i \in [k], \forall j \in [n]$. Thus, $uf_{ij}^{-1}(b_i) \geq \frac{9n^{2.5}}{\varepsilon}$ for every i, j, and so $\ell = \min_{i\in[k]} \min_{j\in[n]} uf_{ij}^{-1}(b_i) \geq \frac{9n^{2.5}}{\varepsilon}$

Thus it always holds that $\ell \geq \frac{9n^{2.5}}{\varepsilon}$, and so

$$\begin{aligned}
V &\leq (1+\frac{\varepsilon}{9})(1+\frac{2n\sqrt{n}}{\ell})^n \mathrm{vol}(K) \\
&\leq (1+\frac{\varepsilon}{9})(1+2n\frac{2n\sqrt{n}}{\ell})\mathrm{vol}(K) && \text{By Taylor's series since} \frac{2n^{1.5}}{\ell} < 1 \\
&\leq (1+\frac{\varepsilon}{9}+\frac{4n^{2.5}}{\frac{9n^{2.5}}{\varepsilon}}+\frac{4\varepsilon n^{2.5}}{9\frac{9n^{2.5}}{\varepsilon}})\mathrm{vol}(K) \\
&\leq (1+\varepsilon)\mathrm{vol}(K)
\end{aligned}$$

Thus $\mathsf{vol}(K) \leq V \leq (1+\varepsilon)\mathsf{vol}(K)$. □

3 Hypercube Clipped by a Single Constraint

In this section we will prove Theorems 1.2 and 1.3.

We defer the proof of Theorem 1.7 to the appendix as it is a small extension of a theorem in [GKM10].

In both the case of linear and convex constraints, our algorithm acts as follows: We begin by performing affine transformations on the given convex body K to achieve K_u which is the $[0, u]^n$ cube (for some integer u) intersected by a convex constraint $g(x) \leq b$ where g is nonnegative, nondecreasing, and convex. We then approximately counting the number of integer points in this new body using either COUNT_KNAPSACK or ROUND_ROBP. We prove that this provides an approximate number, and thus volume, of axis aligned unit cubes whose smallest vertex is an integer point in K_u. We then show that this approximation also serves as a good approximation of $\mathsf{vol}(K_u) = u^n \mathsf{vol}(K)$.

3.1 Truncating Halfspace

Consider a hypercube $[0, 1]^n$ and an intersecting hyperplane $a^\top x = b$. Let polytope P be a portion of the cube clipped by this hyperplane, described by $\{x : a^\top x \leq b,\ x \in [0, 1]^n\}$.

In order to prove Theorem 1.2 we will rely on the following propositions:

Proposition 3.1. *[Canonical Position] The volume of* $P = [0, 1]^n \cap \{x : \sum_{j=1}^n a_j x_j \leq b\}$ *is equal to the volume of* $R = [0, 1]^n \cap \{x : \sum_{j=1}^n |a_j| x_j \leq b + \sum_{j\in[n]:a_j<0} |a_j|\}$.

Proposition 3.2. *Consider sets* $S = \{x \in \mathbb{R}^{n \log u} : \sum_{i=1}^{n} \sum_{j=0}^{\log u - 1} 2^j a_i x_{i,j} \leq ub, x_{ij} \in \{0,1\}\}$ *and* $Z = \mathbb{Z}^n \cap [0,u)^n \cap \{x : a^\top x \leq ub\} = \big\{x \in \mathbb{R}^n : a^\top x \leq ub, x_i \in \{0, 1, \ldots, u-1\}\big\}$ *where* $a, b, u \geq 0$ *and* u *is a power of* 2*. Then* $|S| = |Z|$.

Using the above statements, we can prove the main theorem of this section.

Proof of Theorem 1.2. Our algorithm first performs affine transformations on polytope P to achieve polytope $R = [0,1]^n \cap \{x : w^\top x \leq c\}$ where $w_j = |a_j| \geq 0$ for all j and $c = b + \sum_{j \in [n] : a_j < 0} w_j$. Proposition 3.1 tells us that $\mathsf{vol}(P) = \mathsf{vol}(R)$. Thus we can approximate the volume of polytope R, which has nonnegative coefficients, instead. From here on we can assume WLOG that $P = [0,1]^n \cap \{x : a^\top x \leq b\}$ with $a_j, b \geq 0$ for all $j \in [n]$. Note that regardless of the change in the coefficients, the bit complexity of the whole set remains unchanged. Given input bit size L, the new capacity now has bit size at most $\log(n) + L$ as it has value at most $b + \sum_{i=1}^{n} a_j \leq (n+1)(\max\{\max_j a_j, b\})$.

Let u be a power of 2 such that

$$\frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{j \in [n]} \frac{a_j}{b}\} \leq u < 2 \cdot \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{j \in [n]} \frac{a_j}{b}\}.$$

Consider P_u which is the polytope P scaled up by a factor of u in each dimension, such that it now lies in $[0,u]^n$ with hyperplane $a^\top x \leq ub$. We can find P_u by multiplying P with transformation matrix uI_n. The volume of the new polytope is exactly

$$\mathrm{vol}(P_u) = det(uI_n) \cdot \mathrm{vol}(P) = u^n \cdot \mathrm{vol}(P)$$

Thus it suffices to develop an FPTAS for the volume of P_u.

Consider the integer points in $P'_u := \{x \in [0,u)^n : a^\top x \leq ub\}$:

$$|Z| = |\mathbb{Z}^n \cap [0,u)^n \cap \{x : a^\top x \leq ub\}| = |\mathbb{Z}^n \cap P'_u|$$

Take $f_j(x) = a_j x_j$ which is clearly nonnegative, nondecreasing, and convex. Note that $\sum_{j=1}^{n} f_j(x_j/u) \leq b \iff \sum_{j=1}^{n} a_j x_j \leq ub$ and $f_j^{-1}(x_j) = b/a_j$. Also,

$$u \geq \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{j \in [n]} \frac{a_j}{b}\} = \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{j \in [n]} \frac{1}{f_j^{-1}(b)}\}.$$

Thus we can apply Lemma 2.3 which tells us that given an $\frac{\varepsilon}{9}$-approximation V for $|Z|$, then V serves as an ε-approximation for $\mathsf{vol}(P'_u) = \mathsf{vol}(\mathrm{cl}(P'_u)) = \mathsf{vol}(P_u)$.

So, it remains to find an FPTAS that counts the number of integer points in P'_u. This is exactly the number of integer knapsack solutions under the knapsack constraint $a^\top x \leq ub$ with $x \in \{0, 1, \ldots, u-1\}$.

Since u is a power of 2, Claim 3.2 gives us a 0-1 knapsack set

$$S = \{x \in \mathbb{R}^{n \log u} : \sum_{i=1}^{n} \sum_{j=0}^{\log u - 1} 2^j a_i x_{i,j} \leq ub, x_{ij} \in \{0,1\}\}$$

such that $|S| = |Z|$.

Note that counting knapsack solutions is a well known $\#P$ hard problem, so an approximate counting is the best we can do in polynomial time (unless $P = \#P$). We can rely on a Dynamic Programming FPTAS developed by [SVV12] for approximately counting $0-1$ knapsack solutions in S.

Theorem 1.6 tells us that this FPTAS runs in $O(\frac{m^3}{\delta}\log(\frac{m}{\delta}))$, however now our input size $m = n\log u$ and our error parameter $\delta = \frac{\varepsilon}{9}$. Let L be the maximum size of any of our input parameter a_i, b. It follows that $\log(u) = O(\log\frac{n^{2.5}}{\varepsilon} + 2L) = O(\log\frac{n}{\varepsilon} + L)$, and so our FPTAS has runtime $O\left(\frac{n^3}{\varepsilon}(\log\frac{n}{\varepsilon} + L)^3(\log(\frac{n}{\varepsilon}) + \log L)\right) = O\left(\frac{n^3}{\varepsilon}(\log\frac{n}{\varepsilon} + L)^4\right)$.

This algorithm gives us V such that

$$|Z| = |S| \leq V \leq (1+\frac{\varepsilon}{9})|S| = (1+\frac{\varepsilon}{9})|Z|$$

and so

$$\mathsf{vol}(P_u) \leq V \leq (1+\varepsilon)\mathsf{vol}(P_u)$$

and

$$\mathsf{vol}(P) \leq \frac{V}{u^n} \leq (1+\varepsilon)\mathsf{vol}(P)$$

Thus we have an ε approximation of the volume of our original polytope.

We explicitly state the algorithm below:

FPTAS for a Single Truncating Halfspace

Input: $a \in \mathbb{R}^n$, $b \in \mathbb{R}$, $\varepsilon > 0$

1. Set $c \leftarrow b + \sum_{j\in[n]:a_j<0}|a_j|$, $u \leftarrow 2^{\left\lceil \log_2\left(\frac{9n^{2.5}}{\varepsilon}\max\{1,\max_{j\in[n]}\frac{|a_j|}{c}\}\right)\right\rceil}$
2. $S \leftarrow \{x \in \mathbb{R}^{n\log u} : \sum_{i=1}^{n}\sum_{j=0}^{\log u-1} 2^j|a_i|x_{i,j} \leq uc, x_{ij} \in \{0,1\}\}$
3. Call COUNT_KNAPSACK(ε) on set S, obtain V such that $|S| \leq V \leq (1+\varepsilon)|S|$.
4. Return $\frac{V}{u^n}$

$\square$

Now we prove the lemma and propositions stated at the beginning of the section

Proof of Proposition 3.1. Recall that $P = [0,1]^n \cap \{x : a^\top x \leq b\}$ and $R = [0,1]^n \cap \{x : |a|^\top x \leq b + \sum_{j\in[n]:a_j<0}|a_j|\}$.

Let $J = \{j \in [n] : a_j < 0\}$. Let A be the diagonal matrix with $A_{j,j} = -1\ \forall j \in J$ and $A_{j,j} = 1\ \forall j \notin J$. Note that $\det(A) \in \{1,-1\}$. Let s be the indicator vector of the subset J, i.e., $s_j = 1 \ \forall j \in J$, $s_j = 0$ otherwise. Let polytope $Q = \{Ax + s : x \in P\}$ and notice that $\mathsf{vol}(Q) = |\det(A)|\mathsf{vol}(P) = \mathsf{vol}(P)$.

It remains to show that $Q = R$, or $x \in P \iff Ax + s \in Q \iff Ax + s \in R$. We prove $x \in [0,1]^n \iff Ax + s \in [0,1]^n$ and $x \in \{x : a^\top x \le b\} \iff Ax + s \in \{x : |a|^\top x \le b + \sum_{j\in[n]:a_j<0} |a_j|\}$ and so the claim follows. Note that

$$(Ax+s)_j = \begin{cases} -x_j + 1 & j \in J \\ x_j & j \notin J \end{cases}$$

and so $(Ax + s)_j \in [0,1]$ for $x \in [0,1]$.

Next, we see that

$$b \ge \sum_{j=1}^n a_j x_j = \sum_{j\in J} a_j x_j + \sum_{j\notin J} a_j x_j = \sum_{j\in J} -|a_j| x_j + \sum_{j\notin J} |a_j| x_j \iff$$

$$b + \sum_{j\in J} |a_j| \ge \sum_{j\in J} -|a_j|x_j + \sum_{j\notin J} |a_j|x_j + \sum_{j\in J} |a_j| = \sum_{j\in J} |a_j|(-x_j+1) + \sum_{j\notin J} |a_j| x_j = \sum_{j=1}^n |a_j|(Ax+s)_j$$

So

$$x \in \{x : a^\top x \le b\} \iff Ax + s \in \{x : |a|^\top x \le b + \sum_{j\in[n]:a_j<0} |a_j|\}$$

and we can conclude that $R = Q$ and $\mathsf{vol}(R) = \mathsf{vol}(Q) = \mathsf{vol}(P)$. $\square$

Proof of Proposition 3.2. In this claim we index $x \in S$ in a way that makes it more intuitive. In reality, $x_{i,j}$ refers to $x_{(i-1)(\log u)+j+1}$ which gives $n \log u$ many variables for $i \in \{1, 2, \ldots n\}$ and $j \in \{0, 1, \ldots, \log u - 1\}$

We can construct a one-to-one correspondence between each solution in Z and S as follows: Consider $x^* \in Z$. Let $\bar{x} \in \{0,1\}^{n \log u}$ such that $\{\bar{x}_{i,0}, \bar{x}_{i,1}, \ldots, \bar{x}_{i,\log u-1}\}$ is the unique binary representation of x_i^*: $\sum_{j=0}^{\log u-1} 2^j \bar{x}_{i,j} = x_i^*$. It follows that

$$\sum_{i=1}^n \sum_{j=0}^{\log u-1} 2^j a_i \bar{x}_{i,j} = \sum_{i=1}^n a_i \sum_{j=0}^{\log u-1} 2^j \bar{x}_{i,j} = \sum_{i=1}^n a_i x_i^* \le ub$$

and so $\bar{x} \in S$.

Now consider $\bar{x} \in S$. Let x^* such that $x_i^* = \sum_{j=0}^{\log u-1} 2^j \bar{x}_{i,j}$. Note that since $0 \le \bar{x}_{i,j} \le 1$, it follows that

$$0 \le \sum_{j=0}^{\log u-1} 2^j \bar{x}_{i,j} \le \sum_{j=0}^{\log u-1} 2^j = 2^{\log u - 1 + 1} - 1 = u - 1$$

Thus $x^* \in \{0, \ldots, u-1\}^n$. It also follows that

$$\sum_{i=1}^n a_i x_i^* = \sum_{i=1}^n a_i \sum_{j=0}^{\log u-1} 2^j \bar{x}_{i,j} = \sum_{i=1}^n \sum_{j=0}^{\log u-1} 2^j a_i \bar{x}_{i,j} \le ub$$

and so $x^* \in Z$. Thus $|Z| = |S|$. $\square$

3.2 Truncating Convex Constraint Version 1

Consider a hypercube $[0,1]^n$ and an intersecting constraint $f(x) \leq b$, where $f(x) = \sum_{j=1}^n f_j(x_j)$, and each $f_j : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a nondecreasing convex function on $[0,1]$. Let Q be the intersection of the unit cube with this constraint: $Q = \{x : f(x) \leq b,\, x \in [0,1]^n\}$.

Proof of Theorem 1.3. Let

$$u = \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{j \in n} \frac{1}{f_j^{-1}(b)}\}, \quad g_j(x_j) = f_j(x_j/u), \quad g(x) = \sum_{j=1}^n g_j(x_j).$$

Our algorithm first transforms convex body $Q = \{x : f(x) \leq b,\, x \in [0,1]^n\}$ to $Q_u = [0,u]^n \cap \{x : g(x) \leq b\}$. We can do this through multiplying Q with transformation matrix uI_n, so $\mathsf{vol}(Q_u) = u^n \mathsf{vol}(Q)$.

Consider the integer points in $Q'_u := \{x \in [0,u)^n : g(x) \leq b\}$:

$$Z = \mathbb{Z}^n \cap [0,u)^n \cap \{x : g(x) \leq b\} = \mathbb{Z}^n \cap Q'_u$$

Lemma 2.3 tells us that given an $\frac{\varepsilon}{9}$-approximation V for $|Z|$, then V serves as an ε-approximation for $\mathsf{vol}(Q'_u) = \mathsf{vol}(\mathrm{cl}(Q'_u)) = \mathsf{vol}(Q_u)$.

Note that the number of integer points in Q'_u is exactly the number of integer solutions under the read-once constraint $g(x) \leq b$ with $x \in \{0, 1, \ldots, u-1\}$. We can now apply Theorem 1.7 to give us a $\frac{\varepsilon}{9}$ approximation of the integer points. This FPTAS has runtime $O(n^5 (\log u)^2 (\log b)/\delta)$ where $\delta = \frac{\varepsilon}{9}$ and $\log b \leq L$. Recall that

$$u = \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{j \in n} \frac{1}{f_j^{-1}(b)}\} = \frac{9n^{2.5}}{\varepsilon} \frac{1}{\min\{1, \min_{j \in n} f_j^{-1}(b)\}} = \frac{9n^{2.5}}{\varepsilon \ell}.$$

Note that $\ell = \min\{1, \min_{j \in n} f_j^{-1}(b)\}$ is the smallest axis-intercept of Q, and since Q is convex, the simplex $S_\ell = \{x : \sum_{j=1}^n x_j \leq \ell\}$ is contained in Q. Thus $\mathsf{vol}(Q) \geq \mathsf{vol}(S_\ell) = \frac{\ell^n}{n!}$, and so the bit size of ℓ is $\log \ell \leq \frac{1}{n}(\log(n!) + L_o)m = O(\log n + L_o)$. So, $\log(u) = O(\log \frac{n}{\varepsilon} + L_o)$. Thus we have runtime $O(n^5 (L_o + \log(\frac{n}{\varepsilon}))^2 L/\varepsilon)$ where L is the size of the maximum output bit and L_o is the size of the maximum input bit.

We explicitly state the algorithm below:

FPTAS for Single Truncating Convex Constraint

Input: $f_j : \mathbb{R} \to \mathbb{R}_{\geq 0}$ for $j \in [n]$, $b \in \mathbb{R}_{\geq 0}$, $\varepsilon > 0$

1. Compute $u \leftarrow \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{j\in n} \frac{1}{f_j^{-1}(b)}\}$
2. $S \leftarrow \{x \in \mathbb{Z}^n : \sum_{j=1}^n f_j(\frac{x_j}{u}) \leq b, x \in [0, u-1]^n\}$
3. Call ROUND_ROBP(ε) on set S, obtain V such that $|S| \leq V \leq (1+\varepsilon)|S|$.
4. Return $\frac{V}{u^n}$

□

4 Hypercube Clipped by a Multiple Constraints

In this section we prove Theorems 1.4 and 1.5. As before, both algorithms first dilate our convex body by some factor u. We then approximately counting the number of integer points in this new convex body using an FPTAS developed in [GKM10]. This algorithm constructs small width Read-Once Branching Programs (ROBPs) to approximately count the number of points in such a set. We prove that this provides an approximate number, and thus volume, of unit cubes whose smallest coordinate is an integer point x in the dilated body. We then show that an $\frac{\varepsilon}{9}$ approximation of the volume of these cubes is an ε factor approximation of the volume of our body.

Explicitly, the algorithm will do the following:

FPTAS for Multiple Truncating Convex Constraints

Input: $f_{ij} : \mathbb{R} \to \mathbb{R}_{\geq 0}$ for $i \in [k]$, $j \in [n]$, $b \in \mathbb{R}^n_{\geq 0}$, $\varepsilon > 0$

1. Compute $u \leftarrow \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{i\in[k]} \max_{j\in[n]} \frac{1}{f_{ij}^{-1}(b)}\}$
2. $S_i \leftarrow \{x \in \mathbb{Z}^n : \sum_{j=1}^n f_{ij}(\frac{x_j}{u}) \leq b_i, x \in [0, u-1]^n\}$, $S \leftarrow \cap_{i=1}^k S_i$
3. Call ROUND_ROBPS(ε) on set S, obtain V such that $|S| \leq V \leq (1+\varepsilon)|S|$.
4. Return $\frac{V}{u^n}$

We defer the proofs of 1.8 and 1.9 the above to the appendix, as they are small extensions to proofs from [Dye03] and [GKM10] respectively.

4.1 Multiple Nonnegative Linear Constraints

Consider a hypercube $[0,1]^n$ and a collection of k intersecting hyperplanes each described by $a_i^\top x = b_i$ where $a_{ij}, b_i \geq 0\ \forall i \in [k], j \in [n]$. Let polytope P be the intersection of the cube $[0,1]^n$ truncated by these planes, i.e., $P = \{x : a_i^\top x \leq b_i\ \forall i \in [k],\ x \in [0,1]^n\}$.

Proof of Theorem 1.4. Our algorithm first transforms convex body $P = \{x : Ax \leq b, x \in [0,1]^n\}$ to $P_u = \{x : Ax \leq ub, x \in [0,u]^n\}$ through dilating by a factor of $u = \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{i\in[k]} \max_{j\in[n]} \frac{a_i}{b}\}$. It follows that $\mathsf{vol}(P_u) = u^n \mathsf{vol}(P)$.

Consider the integer points in $P'_u := \{x \in [0,u)^n : Ax \leq ub\}$:

$$Z = \mathbb{Z}^n \cap [0,u)^n \cap \{x : Ax \leq ub\} = \mathbb{Z}^n \cap P'_u$$

Take $f_{ij}(x) = a_{ij}x_j$ which is clearly nonnegative, nondecreasing, and convex. Note that

$$\sum_{j=1}^n f_{ij}(\frac{x_j}{u}) \leq b_i \iff \sum_{j=1}^n a_{ij}x_j \leq ub_i \text{ and } f_j^{-1}(x_j) = \frac{b}{a_j}.$$

From the definition of u, it follows that

$$u \geq \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{i\in[k]} \max_{j\in[n]} \frac{a_i}{b}\} = \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{i\in[k]} \max_{j\in[n]} \frac{1}{f_{ij}^{-1}(b_i)}\}.$$

We can apply Lemma 2.3 which tells us that given an $\frac{\varepsilon}{9}$-approximation V for $|Z|$, then V serves as an ε-approximation for $\mathsf{vol}(P'_u) = \mathsf{vol}(\mathrm{cl}(P'_u)) = \mathsf{vol}(P_u)$.

So, it remains to find an FPTAS that counts the number of integer points in P'_u. This is exactly the number of integer solutions under the integer knapsack constraints $a_i^\top x \leq b_i$, $i \in [k]$ with $x \in \{0, 1, \ldots, u-1\}$.

Theorem 1.9 gives us such an FPTAS with runtime $O(n^{O(k^2)}(\log u/\varepsilon)^{O(k)} \log b)$. Recall $u = \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{i\in[k]} \max_{j\in[n]} \frac{a_i}{b}\}$. So, $\log(u) = O(\log \frac{n}{\varepsilon} + L)$ and note that $\log b = O(L)$. Thus we have runtime $n^{O(k^2)}(\log \frac{n}{\varepsilon} + L)^{O(k)}/\varepsilon^{O(k)}$ where L is the size of the maximum input bit. $\square$

4.2 Multiple Convex Constraints

Let $f_i(x) = \sum_{j=1}^n f_{ij}(x_j)$ where $f_{ij} : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ are nondecreasing convex functions on $[0,1]$. Let

$$Q = \{x \in [0,1]^n : f_i(x) \leq b_i \quad \forall i \in [k]\}.$$

Proof of Theorem 1.5. Let

$$u = \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{i\in[k]} \max_{j\in[n]} \frac{1}{f_{ij}^{-1}(b)}\}, \quad g_{ij}(x_j) = f_{ij}(\frac{x_j}{u}), \quad g_i(x) = \sum_{j=1}^n g_{ij}(x_j).$$

Our algorithm first transforms convex body $Q = \{x : f_i(x) \le b_i \ \forall i \in [k],\ x \in [0,1]^n\}$ to $Q_u = \{x : g_i(x) \le b_i \ \forall i \in [k],\ x \in [0,u]^n\}$ through dilating by a factor of u. We do this by multiplying with transformation matrix uI_n, so it follows that $\mathsf{vol}(Q_u) = u^n \mathsf{vol}(Q)$.

Consider the integer points in $Q'_u := \{x \in [0,u)^n : g_i(x) \le b_i \ \forall i \in [k]\}$:

$$Z = \mathbb{Z}^n \cap [0,u)^n \cap \{x : g_i(x) \le b_i \quad \forall i \in [k]\} = \mathbb{Z}^n \cap Q'_u$$

Lemma 2.3 tells us that given an $\frac{\varepsilon}{9}$-approximation V for $|Z|$, then V serves as an ε-approximation for $\mathsf{vol}(Q'_u) = \mathsf{vol}(\mathrm{cl}(Q'_u)) = \mathsf{vol}(Q_u)$.

It remains to find an FPTAS that counts the number of integer points in Q'_u. This is exactly the number of integer solutions under the integer read-once constraints $f_i(x) \le b_i$ with $x \in \{0, 1, \ldots, u-1\}$, $i \in [k]$.

We use the adapted FPTAS developed by [GKM10]. Theorem 1.9 gives us such an FPTAS with runtime $O(n^{O(k^2)}(\log u/\varepsilon)^{O(k)} \log W)$ where $W = \max_{i\in[k]} b_i$. Thus $\log W \le L + \log k = O(L)$. Recall that

$$u = \frac{9n^{2.5}}{\varepsilon} \max\{1, \max_{i\in[k]} \max_{j\in[n]} \frac{1}{f_{ij}^{-1}(b)}\} = \frac{9n^{2.5}}{\varepsilon} \frac{1}{\min\{1, \min_{i\in[k]} \min_{j\in n} f_{ij}^{-1}(b)\}}.$$

Note that $\ell = \min\{1, \min_{i\in[k]} \min_{j\in n} f_j^{-1}(b)\}$ is the smallest axis-intercept of Q, and since Q is convex, the simplex $S_\ell = \{x : \sum_{j=1}^n x_j \le \ell\}$ is contained in Q. Thus $\mathsf{vol}(Q) \ge \mathsf{vol}(S_\ell) = \frac{\ell^n}{n!}$, and so the bit size of ℓ is $\log \ell \le \frac{1}{n}(\log(n!) + L_o)m = O(\log n + L_o)$. So, $\log(u) = O(\log \frac{n}{\varepsilon} + L_o)$. Thus we have runtime $n^{O(k^2)}(\log \frac{n}{\varepsilon} + L + L_o)^{O(k)}/\varepsilon^{O(k)}$ where L is the size of the maximum input bit and L_o is the size of the maximum output bit. $\square$

5 Hardness of Approximately Counting solutions to Two Linear Inequalities

Proof of Theorem 1.10. We prove that this problem is hard by reducing to a variant of the Subset-Sum Problem. Consider a set $S = \{s_1, \ldots, s_n\}$, where each $s_i \in \mathbb{Z}$, and target $T = 0$. The problem of deciding if there is a subset $I \subseteq [n]$ such that $\sum_{i\in I} s_i = 0$ is NP-hard.

This problem has a solution if and only if there exists an $x \in \{0,1\}^n$ such that $\sum_{i=1}^n s_i x_i = 0$. Here x_i can be though of as the indicator variable for wether or not $i \in I$.

Note that the sum $\sum_{i=1}^n s_i x_i$ always has integral value for $x \in \{0,1\}^n$. So, if $\sum_{i=1}^n s_i x_i \neq 0$, it must be true that $\sum_{i=1}^n s_i x_i \ge 1$ or $\sum_{i=1}^n s_i x_i \le -1$. Similarly, any solution satisfying $-1 < \sum_{i=1}^n s_i x_i < 1$ must have sum exactly equal to 0.

Let $a_1 \in \mathbb{R}^n$ to be the vector such that $a_{1i} = 2s_i$ and let $b = 1$. Consider set $S = \{x \in \{0,1\}^n : -b \le a^\top x \le b\}$ There exists solution $x \in S$ if and only if there is $x \in \{0,1\}^n$ such that $-1 \le \sum_{i=1}^n a_i x_i \le 1$, or equivalently $-1 < -\frac{1}{2} \le \sum_{i=1}^n s_i x_i \le \frac{1}{2} < 1$. Thus, determining if there is a point in S is as hard as solving the subset sum variant. $\square$

References

[AK91] David Applegate and Ravi Kannan. Sampling and integration of near log-concave functions. In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, May 5-8, 1991, New Orleans, Louisiana, USA*, pages 156–163, 1991.

[BF87] I. Bárány and Z. Füredi. Computing the volume is difficult. *Discrete Comput. Geom.*, 2(4):319–326, 1987.

[BR24] Alexander Barvinok and Mark Rudelson. A quick estimate for the volume of a polyhedron. *Israel Journal of Mathematics*, 262(1):449–473, 2024.

[BS79] D. L. Barrow and P. W. Smith. Spline notation applied to a volume problem. *The American Mathematical Monthly*, 86(1):50–51, 1979.

[CK15] Yunhi Cho and Seonhwa Kim. Volume of hypercubes clipped by hyperplanes and combinatorial identities. *The Electronic Journal of Linear Algebra*, 36, 12 2015.

[Cos24] Marius Costandin. A deterministic algorithm of quasi-polynomial complexity for clipped cubes volume approximation, 2024.

[CV18] Ben Cousins and Santosh S. Vempala. Gaussian cooling and $\mathrm{o}^{*}(\mathrm{n}^{3})$ algorithms for volume and gaussian volume. *SIAM J. Comput.*, 47(3):1237–1273, 2018.

[DF88a] M. E. Dyer and A. M. Frieze. On the complexity of computing the volume of a polyhedron. *SIAM J. Comput.*, 17(5):967–974, 1988.

[DF88b] M. E. Dyer and A. M. Frieze. On the complexity of computing the volume of a polyhedron. *SIAM Journal on Computing*, 17(5):967–974, 1988.

[DF91] M. E. Dyer and A. M. Frieze. Computing the volume of a convex body: a case where randomness provably helps. In *Proc. of AMS Symposium on Probabilistic Combinatorics and Its Applications*, pages 123–170, 1991.

[DFK89] M. E. Dyer, A. M. Frieze, and R. Kannan. A random polynomial time algorithm for approximating the volume of convex bodies. In *STOC*, pages 375–381, 1989.

[Dye03] Martin Dyer. Approximate counting by dynamic programming. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '03, page 693–699, New York, NY, USA, 2003. Association for Computing Machinery.

[Ele86] G. Elekes. A geometric inequality and the complexity of computing volume. *Discrete and Computational Geometry*, pages 289–292, 1986.

[FB88] Z. Furedi and I. Barany. Approximation of the sphere by polytopes having few vertices. *Proceedings of the AMS*, 102(3), 1988.

[GKM10] Parikshit Gopalan, Adam R. Klivans, and Raghu Meka. Polynomial-time approximation schemes for knapsack and related counting problems using branching programs. *CoRR*, abs/1008.3187, 2010.

[GN24] Heng Guo and Vishvajeet N. Deterministic approximation for the volume of the truncated fractional matching polytope, 2024.

[KLS95] R. Kannan, L. Lovász, and M. Simonovits. Isoperimetric problems for convex bodies and a localization lemma. *Discrete and Computational Geometry*, 13:541–559, 1995.

[KLS97] R. Kannan, L. Lovász, and M. Simonovits. Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Structures and Algorithms*, 11:1–50, 1997.

[KV06] A. T. Kalai and S. Vempala. Simulated annealing for convex optimization. *Math. Oper. Res.*, 31(2):253–266, 2006.

[KV25] Yunbum Kook and Santosh S Vempala. Sampling and integration of logconcave functions by algorithmic diffusion. In *ACM Symposium on the Theory of Computing*, 2025.

[Lov90] L. Lovász. How to compute the volume? *Jber. d. Dt. Math.-Verein, Jubiläumstagung 1990*, pages 138–151, 1990.

[LS90] L. Lovász and M. Simonovits. Mixing rate of Markov chains, an isoperimetric inequality, and computing the volume. In *ROCS*, pages 482–491, 1990.

[LS93] L. Lovász and M. Simonovits. Random walks in a convex body and an improved volume algorithm. In *Random Structures and Alg.*, volume 4, pages 359–412, 1993.

[LV06] L. Lovász and S. Vempala. Fast algorithms for logconcave functions: sampling, rounding, integration and optimization. In *FOCS*, pages 57–68, 2006.

[LV07] L. Lovász and S. Vempala. The geometry of logconcave functions and sampling algorithms. *Random Structures and Algorithms*, 30(3):307–358, 2007.

[MZ09] Raghu Meka and David Zuckerman. Pseudorandom generators for polynomial threshold functions. *CoRR*, abs/0910.4122, 2009.

[NS26] Hariharan Narayanan and Piyush Srivastava. Deterministically approximating the volume of a kostka polytope. *International Mathematics Research Notices*, 2026(6):rnag036, 03 2026.

[SVV12] Daniel Stefankovic, Santosh S. Vempala, and Eric Vigoda. A deterministic polynomial-time approximation scheme for counting knapsack solutions. *SIAM J. Comput.*, 41(2):356–366, 2012.

A Appendix

A.1 Preliminaries

First, we note some definitions regarding Read-Once Branching Programs (ROBPs) as introduced by [MZ09]:

Definition A.1 (ROBPs). *For $u = (u_1, \ldots, u_n) \in \mathbb{Z}_+^n$, $S, T \in \mathbb{Z}_+$, an (S, u, T) ROBP M is a layered multi-graph with a layer for each $0 \leq l \leq T$, and at most S states in each layer. The first layer has a single (start) vertex s, and each vertex in the last layer is labeled accepting or rejecting. Each vertex v in layer $l-1$ has exactly $u_l + 1$ edges labeled $\{0, 1, \ldots, u_l\}$ to layer l.*

We also introduce the following notation

1. $\mathbf{L}(M, l)$ denotes the vertices in layer l of M

2. For a string $\{0, 1, \ldots, u_{l+1}\} \times \cdots \times \{0, 1, \ldots, u_j\}$ and a vertex $v \in L(M, l)$, $\mathbf{M}(v, z)$ denotes the state in layer j reached by starting from v and following edges labeled z.

3. For $z \in \{0, 1, \ldots, u_1\} \times \cdots \times \{0, 1, \ldots, u_n\}$, let $\mathbf{M}(z) = 1$ if $M(s, z)$ is an accept state, 0 otherwise.

4. For $v \in L(M, l)$, $\mathbf{A_M}(v) = \{z \in \{0, 1, \ldots, u_{l+1}\} \times \cdots \times \{0, 1, \ldots, u_n\} : M(v, z)$ is accepting in $M\}$

5. For $v \in L(M, l)$, $\mathbf{P_M}(v) = \frac{|A_M(v)|}{\prod_{j=l}^{n}(u_j+1)}$ is the fraction of suffixes $z \in \{z \in \{0, 1, \ldots, u_{l+1}\} \times \cdots \times \{0, 1, \ldots, u_n\}$ that lead to an accepting state. This is equivalently the probability that a suffix sampled uniformly at random leads to an accepting state.

6. Input S is referred to as the **width** of M. The *width* of layer l of M is exactly $|L(M, l)|$, the number of states in the layer.

We use this notation to define an important class of ROBPs which we will discuss.

Definition A.2 (Interval ROBPs). *An (S, u, T)-interval ROBP M is an ROBP such that there exists a total order $\prec$ on the vertices of layer $L(M, l)$ such that for $w, v \in L(M, l)$ with $w \prec v$ then $A_M(w) \subseteq A_M(v)$. In addition, vertices v in layer $l-1$ have out edges labeled $\{0, 1, \ldots, u_l\}$ that respect the ordering $\prec$: if $M(v, k)$ denotes the kth neighbor of v for $k \leq u_i$, then $M(v, u_l) \preceq M(v, u_l - 1) \preceq \cdots \preceq M(v, 0)$.*

An interval ROBP defines a function $M : [0, u_1] \times [0, u_2] \times \cdots \times [0, u_n] \to \{0, 1\}$ where on input x, we begin at the start vertex and output the label of the final vertex reached when traversing edges of M according to labels in x.

Note that given an (S,u,T)-interval ROBP M, and a vertex $v \in L(M, l-1)$, the edges out of v can be described succinctly by a subset of at most S edges irrespective of how large u is. Let $E(v,w) = \{0 \le k \le u_l : M(v,k) = w\}$ which is the set of edges from v to w. Then $E(v,w)$ is an interval, meaning $E(v,w) = \{\ell_{v,w}, \ldots, r_{v,w}\}$ for some $\ell_{v,w}, r_{v,w} \in \mathbb{Z}_+$.

Thus, to describe $E(v,w)$ we only need to know $\ell_{v,w}$ and $r_{v,w}$. This allows us to succinctly describe an interval ROBP $\hat{M}$ by storing just the end points of the edge sets $E(v,w)$ for $v, w \in \hat{M}$. This bounds the number of edge sets by the number of vertices in the subsequent layer, which is at most width S.

A.2 Read Once Branching Program for a single read-once convex constraint

Proof of Theorem 1.7. In [GKM10], Theorem 1.3 (integer knapsack) states that given a knapsack instance $\mathrm{KNAP}(a,b,u)$ with weight $W = \sum_i a_i u_i + b$, $U = \max_i u_i$ and $\varepsilon > 0$, there is a deterministic $O\big(n^5 (\log U)^2 (\log W)/\delta^2\big)$ algorithm that computes an δ-relative error approximation for $|\mathrm{KNAP}(a,b,u)|$. Here $\mathrm{KNAP}(a,b,u)$ is the set $\{x \in \mathbb{Z} : \sum_{j=1}^n a_j x_j \le b, 0 \le x_j \le u_j \quad \forall j \in [n]\}$. We substitute $\mathrm{KNAP}(a,b,u)$ with $S(f,b,u) = \{x \in \mathbb{Z} : \sum_{j=1}^n f_j(x_j) \le b, 0 \le x_j \le u-1 \quad \forall j \in [n]\}$ where each f_j is nonnegative, convex, and nondecreasing. These functions obey all the properties used in the proof of the theorem in [GKM10]. We provide a detailed replica of the proof with the appropriate substitutions made.

We will construct an approximate branching program for our set $S(f,b,u)$. We will do so by selecting a subset of vertices in the ROBP to remain, and altering the edge sets accordingly. Note that throughout the proof we will assume without loss of generality that $f_j(0) = 0$ for all $j \in [n]$, which follows from proposition 2.1.

Let M denote the exact branching program for $\sum_{j=1}^n f_j(x_j) \le b$, which consists of $n+1$ layers numbered 0 to n. Layer 0 has single start state s, and for $l \le n$, layer $L(M,l)$ has a state for each partial sum $\sum_{j=1}^l f_j(x_j)$. Given a vertex v in layer l, and $x_{l+1} \in \{0,1,\ldots,u-1\}$, the x_{l+1}^{th} neighbor of v is $M(v, x_{l+1}) = v + f_{l+1}(x_{l+1})$. A vertex $v \in L(M,n)$ is accepting if partial sum $v \le b$, and rejecting otherwise.

Note that this program is clearly read once. It is also interval with total ordering $\prec$ defined as $>$. We can see this because for any $v, w \in L(M,l)$ with $v > w$ and any suffix $z_j \in \{0,1,\ldots,u-1\}^{n-l}$, we have $v + \sum_{j=l+1}^n f_j(z_j) > w + \sum_{j=l+1}^n f_j(z_j)$. Thus if $z_j \in A_M(v)$, or $v + \sum_{j=l+1}^n f_j(z_j) \le b$, then clearly $w + \sum_{j=l+1}^n f_j(z_j) \le b$ so $w \in A_M(v)$, and so $A_M(v) \subseteq A_M(w)$. Also, for any vertex $v \in L(M,l)$ and pair of suffixes $0 \le \ell < r \le u-1$, we have $M(v,\ell) = v + f_{l+1}(\ell) \le v + f_{l+1}(r) = M(v,r)$ since each f_{l+1} is nondecreasing. Thus, $M(v,\ell) \succeq M(v,r)$.

By our construction of states, that the width of layer $L(M,l)$ is bounded by the total number of partial sums between our minimum and maximum possibilities: $\{0,1,\ldots,\sum_{j=1}^l f_j(u-1)\}$. Thus the maximum width is $W := \sum_{j=1}^n f_j(u-1)$. This is potentially exponential, so we will construct a new ROBP with smaller width. A small width program will allow us to evaluate the number of accepting solutions in polynomial time.

Our algorithm will construct a series of interval ROBPs $\hat{M}_n = M, \hat{M}_{n-1}, \ldots, \hat{M}_0$. Our final ROBP $\hat{M} := \hat{M}_0$ will have polynomially bounded width. We obtain $\hat{M}_i$ from $\hat{M}_{i+1}$ by *rounding* the states in $L(\hat{M}_{i+1}, i+1)$. More precisely, we set $L(\hat{M}_i, i+1) = \{b_1, \ldots, b_{\ell+1}\} \subseteq L(\hat{M}_{i+1}, i+1)$ where the b_j's are defined as follows: let $b_1 = 0$, which is the minimum value state in layer $i+1$, and equivalently the one with the most accepting suffixes. Given b_j, let

$$b_{j+1} = \min\{v : v > b_j \text{ and } 0 < P_{\hat{M}_{i+1}}(v) < P_{\hat{M}_{i+1}}(b_j)/(1+\eta)\}. \tag{A.1}$$

For η to be set later. Now consider b_ℓ, the last vertex assigned. Let $b_{\ell+1} = \min\{v : P_{\hat{M}_{i+1}}(v) = 0\}$. Note that $b_1 \succ b_2 \succ \cdots \succ b_\ell \succ b_{\ell+1}$ as the probabilities decrease. Let $I_1 = \{b_1, \ldots, b_2 - 1\}, I_2 = \{b_2, \ldots, b_3 - 1\}, \ \ldots, \ I_{\ell+1} = \{b_{\ell+1}, \ldots\}$, where $\ell \leq n(\log u)/\eta$ as $P_{\hat{M}_i}(b_1) \leq 1$ and $P_{\hat{M}_i}(b_\ell) \geq u^{-n}$. We refer to $b_1, \ldots, b_{\ell+1}$ as breakpoints since they break our original vertices into intervals. Next we redirect the transitions going from level i to level $i+1$. If we have an edge labeled $z \in \{0, \ldots, u-1\}$ entering a vertex $v \in I_j$, then we redirect the edge to vertex v_j. The redirection will be done implicitly: for any vertex v in level i and a vertex b_j, we only compute and store the end points of the interval $E(v, b_j) = \{0 \leq k \leq u-1 : \hat{M}_i(v,k) = b_j\}$.

Our branching programs have the following approximating properties:

Claim A.1 (Corresponding to the first half of Lemma A.2. of [GKM10])**.** *For $v \in L(\hat{M}^i, l)$ and $0 \leq \ell < r \leq u-1$, $\hat{M}^i(v,l) \leq \hat{M}^i(v,r)$.*

Proof. Note that this claim is only interesting if $l > i$, otherwise all vertices $v, \hat{M}_i(v,\ell), \hat{M}_i(v,r)$ have the same structure in the original interval ROBP M and thus the claim holds by the interval property. Thus we can assume that layer $l+1$ is rounded and consists of breakpoints. Take $\hat{M}_i(v,\ell) = b_x$ and $\hat{M}_i(v,r) = b_y$.

Assume towards contradiction that $b_x > b_y$.

When creating edge labeled ℓ from vertex v, we found the value of $v_\ell = v + f_{l+1}(\ell)$, and similarly for edge labeled r, we found the value of $v_r = v + f_{l+1}(r)$. We then found intervals $I_x = \{b_x, \ldots, b_{x+1} - 1\}$ and $I_y = \{b_y, \ldots, b_{y+1} - 1\}$ such that $b_x \leq v_\ell < b_{x+1}$ and $b_y \leq v_r < b_{y+1}$. Since $\ell < r$ and f_{l+1} is nondecreasing, we know that $v_\ell = v + f_{l+1}(\ell) \leq v + f_{l+1}(r) = v_r$. Thus we see that $b_{y+1} > v_r \geq v_l \geq b_x > b_y$, which implies that there is a breakpoint (b_x) between (and not equal to) consecutive breakpoints b_{y+1}, b_y which is a contradiction, and so the claim holds. $\square$

Claim A.2 (Corresponding to the second half of Lemma A.2. of [GKM10])**.** *Let $v, v' \in L(\hat{M}_i, l)$, $v \leq v'$. For suffix z, $\hat{M}_i(v,z) \leq \hat{M}_i(v',z)$.*

Proof. Note that we need only prove this claim for suffixes of length one, as the property for any longer suffix will follow inductively.

Also note as in the previous theorem, this claim is only interesting if $l > i$, otherwise vertices $v, v', \hat{M}_i(v,z), \hat{M}_i(v',z)$ have the same structure in the original interval ROBP M. Thus we can assume that layer $l+1$ is rounded and consists of breakpoints.

Consider two vertices v, v' in layer l of our partially rounded ROBP such that $v \le v'$. We aim to show that suffixes maintain the ordering. Take $\hat{M}_i(v, z) = b_x$ and $\hat{M}_i(v', z) = b_y$.

Assume towards contradiction hat $b_x > b_y$. Take $v_z = v + f_{l+1}(z)$, $v'_z = v' + f_{l+1}(z)$ and note that since $v \le v'$, we have $v_z \le v'_z$. Thus, we see that $b_{y+1} > v'_z \ge v_z \ge b_x > b_y$, which implies that there is a breakpoint (b_x) between (and not equal to) consecutive breakpoints b_{y+1}, b_y which is a contradiction, and so the claim holds. □

By inductively applying the previous two claims to each layer, we see that each $\hat{M}_i$ remains an interval ROBP respecting the same ordering as M. We now analyze how the probability of acceptance (and thus the total number of solutions) in our new ROBP $\hat{M}$ compares to that of the original ROBP M.

Claim A.3 (Corresponding to the first half Lemma A.3. of [GKM10])**.** *For $v \in \hat{M}_i$, we have $A_{\hat{M}_{i+1}}(v) \subseteq A_{\hat{M}_i}(v)$.*

Proof. Note that for any vertex $v \in L(\hat{M}_i, l)$ where $l > i$, both ROBPs make identical transitions starting from v, and so $A_{\hat{M}_{i+1}}(v) = A_{\hat{M}_i}(v)$.

Let $l = i$. Let $b_z = \hat{M}_i(v, z)$ be the child of v along edge $z \in \{0, 1, \ldots, u-1\}$. Note that since $\hat{M}_i$ is obtained from $\hat{M}_{i+1}$ by rounding layer $i+1$, these children are breakpoints, and the "true" children of v are $v_z = M_{i+1}(v, z)$ in layer $i+1$ of $\hat{M}_{i+1}$ such that $v_z \ge b_z$.

Since each $\hat{M}_i$ remains an interval ROBP and the structure below b_z is the same in $\hat{M}_i$ and $\hat{M}_{i+1}$, we have $A_{\hat{M}_{i+1}}(v_z) \subseteq A_{\hat{M}_{i+1}}(b_z) = A_{\hat{M}_i}(b_z)$. Thus, taking the union over each suffix z placed ahead of each string in $A_{\hat{M}_{i+1}}(v_z)$ and $A_{\hat{M}_i}(b_z)$, we maintain this containment, and $A_{\hat{M}_{i+1}}(v) \subseteq A_{\hat{M}_i}(v)$.

Now for any $l < i$, the claim follows because $\hat{M}_i$ and $\hat{M}_{i+1}$ make identical transitions up to layer i. Thus we can take the union over suffixes used up to layer i and achieve the same property. □

Claim A.4 (Corresponding to the second half of Lemma A.3. of [GKM10])**.** *For any $v \in L(\hat{M}_i, l)$ where $l \le i$, we have $P_M(v) \le P_{\hat{M}_i}(v) \le P_M(v)(1+\eta)^{n-i}$.*

Proof. Note that since $A_{\hat{M}_{i+1}}(v) \subseteq A_{\hat{M}_i}(v)$, it is clear that $P_{\hat{M}_{i+1}}(v) \le P_{\hat{M}_i}(v)$ for all v, and thus $P_M(v) = P_{\hat{M}_n}(v) \le P_{\hat{M}_{n-1}}(v) \le \cdots \le P_{\hat{M}_i}(v)$.

We will show that for every $i < n$ and $v \in M^i$, $P_{M^i}(v) \le P_{M^{i+1}}(v)(1+\eta)$, and so $P_{\hat{M}_i}(v) \le P_{\hat{M}_n}(v)(1+\eta)^{n-i}$ as needed.

Let $v \in L(\hat{M}_i, l)$. The above is trivial when $l \ge i+1$, since $A_{\hat{M}_i}(v) = A_{\hat{M}_{i+1}}(v)$ for such v. Indeed, it suffices to consider the case when $l = i$, since for $l < i$, $\hat{M}_i$ and M are identical up to layer i. Hence we can express both $P_{\hat{M}_{i+1}}(v)$ and $P_{\hat{M}_i}(v)$ as the same convex combination of acceptance probabilities of vertices in layer i.

Let $l = i$. Fix a vertex $v \in L(\hat{M}_i, i)$. For any $z \in \{0, 1, \ldots, u-1\}$, let $v_z = \hat{M}_{i+1}(v, z)$ be the "true" children of v and let $b_{x_z} = \hat{M}_i(v, z)$ for $z \in \{0, 1, \ldots, u-1\}$ be the breakpoint children. Recall that $\hat{M}_i$ is obtained from $\hat{M}_{i+1}$ by rounding layer i, and $v_z \in I_{x_z} = \{b_{x_z}, \ldots, b_{x_z+1} - 1\}$.

If $x_z < \ell$, then breakpoint b_{x_z+1} is be the smallest partial sum such that $0 < P_{\hat{M}_{i+1}}(b_{x_z+1}) < P_{\hat{M}_{i+1}}(b_{x_z})/(1+\eta)$. Note that $b_{x_z+1}-1$ wasn't chosen as the breakpoint, but the interval property tells us that $0 < P_{\hat{M}_{i+1}}(b_{x_z+1}-1)$. Thus, it must be true that $P_{\hat{M}_{i+1}}(b_{x_z+1}-1) \geq P_{\hat{M}_{i+1}}(b_{x_z})/(1+\eta)$, and so $P_{\hat{M}_{i+1}}(b_{x_z}) \leq (1+\eta)P_{\hat{M}_{i+1}}(b_{x_z+1}-1) \leq (1+\eta)P_{\hat{M}_{i+1}}(v_z)$.

If $x_z = \ell$, then all $v > b_\ell$ satisfy either $P_{\hat{M}_{i+1}}(v) = 0$ or $P_{\hat{M}_{i+1}}(b_\ell)/(1+\eta) \leq P_{\hat{M}_{i+1}}(v)$ or else v could form a new breakpoint. If $P_{\hat{M}_{i+1}}(v) = 0$, then $v > b_{\ell+1}$, so it must be true that $P_{\hat{M}_{i+1}}(b_{x_z}) = P_{\hat{M}_{i+1}}(b_{\ell+1}) \leq (1+\eta)P_{\hat{M}_{i+1}}(v)$.

Finally if $x_z = \ell+1$ then $P_{\hat{M}_{i+1}}(b_{x_z}) = P_{\hat{M}_{i+1}}(b_\ell) = 0 = (1+\eta)P_{\hat{M}_{i+1}}(v)$ So,

$$P_{\hat{M}_i}(v) = \tfrac{1}{u}\left(\sum_{z=0}^{u-1} P_{\hat{M}_i}(b_{x_z})\right) \leq (1+\eta)\frac{1}{u}\sum_{z=0}^{u-1} P_{\hat{M}_{i+1}}(v_z) = (1+\eta)P_{\hat{M}_{i+1}}(v).$$

□

We next show that $\hat{M}$ can be constructed efficiently.

Claim A.5 (Corresponding to Lemma A.4 of [GKM10])**.** *Each vertex $v_j \in L(\hat{M}_i, i+1)$ can be computed in time $O(n(\log u)(\log b)/\eta)$.*

Proof. The proof is by induction: we maintain the invariant that for every i, we know the vertices b_z of $L(\hat{M}_i, i+1)$ and their acceptance probabilities $P_{\hat{M}_i}(\cdot)$.

Suppose we have the above setup for $l > i$ and have computed $b_1, \ldots, b_j \in L(\hat{M}_i, i+1)$. In order to find the next breakpoint, we must have access to the probabilities of vertices in this layer.

We show that for a given $v \in L(\hat{M}_{i+1}, i+1)$, $P_{\hat{M}_{i+1}}(v)$ can be computed in time $O(n(\log u)/\eta)$. Let $L(\hat{M}_{i+1}, i+2) = \{b_1 < b_2 < \cdots < b_{\ell+1}\}$ then, $E(v, b_j) = \{0 \leq k \leq u-1 : b_j \leq v + f_{i+2}(k) \leq b_{j+1}\}$ where $b_{\ell+2} = v_W$ and so

$$P_{\hat{M}_{i+1}}(v) = \sum_{w \in L(M^{i+1}, i+2)} \frac{|E(v,w)|}{u} P_{\hat{M}_{i+1}}(w).$$

Thus, we can compute $P_{\hat{M}_{i+1}}(v)$ in time $O(n(\log u)/\eta)$ as $|L(\hat{M}_{i+1}, i+2)| \leq n(\log u)/\eta + 1$. Thus, to find b_{j+1}, we can now do binary search on values $P_{\hat{M}_{i+1}}(v)$. We require $0 < P_{\hat{M}_{i+1}}(v)$, so we need only search through vertices in the range $[0, b]$ since $b+1$ can never lead to an accepting suffix due to nonnegativity of each f_j. This requires $O(\log b)$ computations of $P_{\hat{M}_{i+1}}(v)$. Finally notice that $P_{\hat{M}_{i+1}}(b+1) = 0$ and $P_{\hat{M}_{i+1}}(b) > 0$ because the suffix string of 0s accepts: $b + \sum_{j=i+2}^{n} f_j(0) = b + 0 \leq b$. Thus we can set $b_{\ell+1} = b+1$. Once we have computed $L(\hat{M}_i, i+1)$ we store these vertices and their probabilities of acceptance. □

Thus we can construct $\hat{M}$ from M in time $O(n^3(\log u)^2(\log b)/\eta^2)$. We can now finish the proof of our counting result:

We set $\eta = \delta/2n$ and use the above arguments to construct the branching program $\hat{M}$ and compute the value of $P_{\hat{M}}(s)$ where s is the start state. We now apply Lemma A.3 to conclude that

$$P_M(s) \le P_{\hat{M}}(s) \le P_M(s)(1+\eta)^n \le (1+\delta)P_M(s)$$

where the last inequality holds for small enough δ. Finally, note that the number of knapsack solutions is precisely $u^n P_M(s)$. Hence we output $u^n P_{M^0}(s)$.

□

A.3 Dyer's Rounding for multiple sets

Proof of theorem 1.8. Let $U_n = \{0, \dots, u_1\} \times \cdots \times \{0, \dots, u_n\}$. Consider set $S = \cap_{i\in[k]} S_i$ where $S_i = \{x \in U_n : \sum_{j=1}^n f_{ij}(x_j) \le b_i\}$ where each $f_{ij} : \mathbb{R} \to \mathbb{R}_{\ge 0}$ is convex and nondecreasing. We will assume without loss of generality that $f_j(0) = 0$ for all $j \in [n]$, which follows from proposition 2.1.

We expand on the results in [Dye03] to find an approximate counting of $|S|$. We create new sets $S_i' = \{x \in U_n : \sum_{j=1}^n h_{ij}(x_j) \le c_i\}$ where each c_i is polynomially bounded, and the size of $S' = \cap_{i\in[k]} S_i'$ is bounded between the size of S, and some multiple of the size of S.

Specifically, we construct rounded set

$$S' = \bigcap_{i=1}^k S_i', \text{ where } S_i' = \{x \in U_n : \sum_{j=1}^n h_{ij}(x_j) \le 2n^2\}$$

such that the rounded functions are

$$h_{ij}(x_j) = \lfloor \frac{2n^2 f_{ij}(x_j)}{b_i} \rfloor \quad \text{for } 0 \le x_j \le u_j, j \in [n], i \in [k]$$

Observe that for any $x \in S_i$, we have that

$$\sum_{j=1}^n h_{ij}(x_j) = \sum_{j=1}^n \lfloor \frac{2n^2 f_{ij}(x_j)}{b_i} \rfloor \le \sum_{j=1}^n \frac{2n^2 f_{ij}(x_j)}{b_i} = \frac{2n^2}{b_i} \sum_{j=1}^n f_{ij}(x_j) \le \frac{2n^2}{b_i} b_i = 2n^2$$

And so $x \in S_i'$, and thus $S_i \subseteq S_i'$.

Now let $C_i = \{x : f_{ij}(x_j) \le \frac{b_i}{n}\ \forall j \in [n]\}$. Note that for all $x \in C_i$, it follows that $\sum_{j=1}^n f_{ij}(x_j) \le \sum_{j=1}^n \frac{b_i}{n} = n\frac{b_i}{n} = b_i$, and thus $C_i \subseteq S_i \subseteq S_i'$

For $x \in S' \setminus S$, let $I(x) = \{i : x \in S_i' \setminus S_i\}$. For each $i \in I(x)$, there must exist some $p_i^x \in [n]$ such that $f_{ip_i^x}(x_{p_i^x}) > \frac{b_i}{n}$, else $x \in C_i \subseteq S_i$. Note $h_{ip_i^x}(x_{p_i^x}) = \lfloor \frac{2n^2 f_{ip_i^x}(x_{p_i^x})}{b_i} \rfloor \ge \lfloor \frac{2n^2 b_i/n}{b_i} \rfloor = \lfloor \frac{2n^2}{n} \rfloor = 2n$.

Construct $g(x) : S' \to U_n$ as follows: $g(x) = \begin{cases} x & \text{if } x \in S \\ y & \text{where } y_j = x_j \text{ for } j \ne p_i^x, \text{ and } y_{p_i^x} = \lfloor \frac{x_{p_i^x}}{2} \rfloor \end{cases}$

By convex and nondecreasing properties of each f it follows that $f_{ij}(\lfloor \frac{x_j}{2} \rfloor) \le f_{ij}(\frac{x_j}{2}) \le \frac{1}{2} f_{ij}(x_j)$. Thus, for all $x \notin S$, $y = g(x)$, we have

$$\sum_{j=1}^n f_{ij}(y_j) = \frac{b_i}{2n^2} \sum_{j=1}^n \frac{2n^2 f_{ij}(y_j)}{b_i}$$

$$
\begin{aligned}
&= \frac{b_i}{2n^2}\Bigg(\sum_{j\neq p_i^x}\big(\frac{2n^2 f_{ij}(x_j)}{b_i}\big) + \frac{2n^2 f_{ip_i^x}(\lfloor\frac{x_{p_i^x}}{2}\rfloor)}{b_i}\Bigg) \\
&\leq \frac{b_i}{2n^2}\Bigg(\sum_{j\neq p_i^x}(h_{ij}(x_j)+1) + \frac{1}{2}\cdot\frac{2n^2 f_{ip_i^x}(x_{p_i^x})}{b_i}\Bigg) && \text{By convexity of } f_{ij} \\
&\leq \frac{b_i}{2n^2}\Bigg(\sum_{j\neq p_i^x}(h_{ij}(x_j)) + n - 1 + \frac{1}{2}(h_{ip_i^x}(x_{p_i^x})+1)\Bigg) \\
&\leq \frac{b_i}{2n^2}\Bigg(\sum_{j\neq p_i^x}(h_{ij}(x_j)) + n - 1 + h_{ip_i^x}(x_{p_i^x}) - n + \frac{1}{2}\Bigg) && \text{By bound on } h_{ip_i^x}(x_{p_i^x}) \\
&\leq \frac{b_i}{2n^2}\Bigg(2n^2 - \frac{1}{2}\Bigg) < b_i
\end{aligned}
$$

Thus, $g(S'_i) = S_i$ for every i, and for each $y \in S_i$, $|g^{-1}(y)\cap S'_i| \leq 2n+1$ since $y \in f^{-1}(y)$ and for any $1 \leq p_i^x \leq n$, there are at most two possible values of $x_{p_i^x}$ ($2y_{p_i^x}$ and $2y_{p_i^x}+1$).

We may have set $y_{p_i^x} = \lfloor\frac{x_{p_i^x}}{2}\rfloor$ for multiple i if each of these constraints are violated by x. In this case, the inverse mapping changes some set of coordinates P with $0 \leq |P| \leq k$, and so

$$
|f^{-1}(y)| \leq 1 + 2n + 2\binom{n}{2} + \cdots + 2\binom{n}{k} \leq 2n^k
$$

and therefore we have $|S'| \leq 2n^k|S|$. Also note that since each $S_i \subseteq S'_i$, it follows that $S \subseteq S'$, and

$$
|S| \leq |S'| \leq 2n^k|S|
$$

□

A.4 Read Once Branching Program for multiple read-once convex constraint

In this section we prove theorem 1.9, which is an extension of work in [GKM10], which gives us a proof of statement when $u_j = 1$ for all $j \in [n]$.

Let $U_n = \{0,\ldots,u_1\}\times\cdots\times\{0,\ldots,u_n\}$. We will develop a deterministic approximate counting algorithm for set $S = \cap_{i=1}^k S_i$, $S_i = \{x \in U_n : \sum_{j=1}^n f_{ij}(x_j) \leq b_i\}$. Note that throughout the proof we will assume without loss of generality that $f_j(0) = 0$ for all $j \in [n]$, which follows from proposition 2.1.

Recall the definitions of Interval-ROBPs from section A.2. We introduce a few more concepts:

Definition A.3 (small-space sources, Kamp et al.). *A width w small-space source is described by a (w,n)-branching program D with an additional probability distribution p_v on the outgoing edges associated with vertices $v \in D$. Samples from the source are generated by taking a random walk on D according to p_v's and outputting the labels of the edges traversed.*

We will often abuse notation and denote the distribution generated by a small-space source and the source itself by D. Also, we will assume that the distribution D is given to us explicitly as a small-space source. Several natural distributions such as all symmetric distributions and product distributions can be generated by a small-space source. The following claim is directly from [GKM10]:

Claim A.6 (Equivalent to Claim 2.4 in [GKM10])**.** *Given a (W, n)-ROBP M, the uniform distribution over M's accepting inputs, $\{x : M(x) = 1\}$ is a width W small-space source.*

Let D denote a small space generator of width at most S. For $A \subseteq U_n$ we use $D(A)$ to denote the measure of A under D. Let $\mathcal{W}^1, \ldots, \mathcal{W}^n$ be sets of vertices in D with $\mathcal{W}^l$ being the l^{th} layer of D. For a vertex $w \in \mathcal{W}^l$, let D^w be the distribution over $\{0, \ldots, u_{l+1}\} \times \cdots \times \{0, \ldots, u_n\}$ induced by taking a random walk in D starting from w. Given a vertex $v \in L(M, l)$ and $w \in \mathcal{W}^l$, let $P_{M,w}(v)$ denote the probability of accepting if we start from v and make transitions in M according to a suffix sampled from distribution D^w.

Recall that section A.2 tells us that the ROBP M^i exactly computing the indicator function for each S_i is an interval ROBP. Recall that each layer l of M^i has at most $\sum_{j=1}^{l} f_{ij}(u_j)$ states. We represent state v by partial sum $\sum_{j=1}^{l} f_{ij}(x_j)$ where x is the string of edge labels used to reach v from start state v_0. We will use the following theorems to build an approximate ROBP $\hat{M}^i$ for each S_i individually.

Theorem A.7. *Given a (W, n)-interval ROBP M^i for set $S_i(f_i, b_i, u)$, $\delta > 0$, and a small-space distribution D over $\{0, \ldots, u_1\} \times \cdots \times \{0, \ldots, u_n\}$ of width at most $s = n^{O(k)}$, there exists an $(O(ns \log U/\delta), n)$-interval ROBP $\hat{M}^i$ such that for all z, $M^i(z) \leq \hat{M}^i(z)$ and*

$$\Pr_{x \leftarrow D}[M^i(z) = 1] \leq \Pr_{x \leftarrow D}[\hat{M}^i(z) = 1] \leq (1 + \delta) \Pr_{x \leftarrow D}[M^i(z) = 1]$$

Moreover, given an implicit description of M^i and an explicit description of D, $\hat{M}^i$ can be constructed in deterministic time $O(n^{O(k)}(\log U)^2 \log b/\delta^2)$

Proof. We will construct a new small width ROBP $\hat{M}^i$ by iteratively "rounding" each layer of M^i, starting from the last one (layer n). As we did in section A.2, we start from the exact branching program M^i and construct a sequence of programs $\hat{M}^i_n = M^i, \ldots, \hat{M}^i_0 = \hat{M}^i$, where $\hat{M}^i_j$ is obtained from $\hat{M}^i_{j+1}$ by rounding the $(j + 1)$st layer.

We do the rounding in such a way that the acceptance probabilities are well approximated under each of the possible distributions on suffixes D^w. The program $\hat{M}^i$ will have polynomial width.

We construct each new ROBP $\hat{M}^i_{\ell-1}$ as follows:
Consider the partial sums $\{v_1 = \sum_{j=1}^{l} f_{ij}(u_j) \prec v_2 \prec \cdots \prec v_W = \sum_{j=1}^{l} f_{ij}(0) = 0\} = L(\hat{M}^i_l, l)$, which is the set of vertices we wish to round in this iteration. Fix a vertex $w \in \mathcal{W}^l$. We wish to select a subset of these vertices to remain in $\hat{M}^i_{l-1}$. We will do so by contracting vertices with "similar" acceptance probabilities. We create groups of "similar" vertices by finding breakpoints

where acceptance probabilities change by a factor of $\frac{1}{1+\delta}$, and vertices between breakpoints will be contracted to the neighboring breakpoint with higher acceptance probability. However, unlike before, we will calculate probability by sampling over accepting suffixes for each $w \in \mathcal{W}^l$.

Formally, we define a set $B^l(w) = \{v_{w(1)}, \ldots, v_{w(N_w)}\} \subseteq L(M_l^i, l)$ of breakpoints for w as follows. We start with $v_{w(1)} = v_W = 0$, the vertex with highest acceptance probability (and lowest partial sum), and given $v_{w(j)}$, define $v_{w(j+1)}$ by

$$v_{w(j+1)} = \max v \text{ s.t } v \prec v_{w(j)} \text{ and } 0 < P_{\hat{M}_{l-1}^i, w}(v) < P_{\hat{M}_{l-1}^i, w}(v_{w(j)})/(1+\delta)$$

which is the smallest partial sum whose acceptance probability decreases by a factor of at least $\frac{1}{1+\delta}$ from the previous breakpoint. Now consider b_{N_w}, the last breakpoint assigned. Let $b_{N_w+1} = \min\{v : P_{\hat{M}_{i+1}, w}(v) = 0\}$. Note that $b_1 \succ b_2 \succ \cdots \succ b_{N_w} \succ b_{N_w+1}$ as the probabilities decrease. We set $L(\hat{M}_{l-1}^i, l) = B^l := \cup_{w \in \mathcal{W}^\ell} B^l(w) = \{b_1, \ldots, b_N\}$ which is the union of breakpoints for all w. Now we have intervals $I_N = \{b_N, \ldots, b_{N-1} - 1\}, \ldots, I_1 = \{b_1, \ldots, v_1\}$.

The vertices in all other layers stay the same as in M_l^i, as do all the edges except those from layer $l-1$ to l. We round these edges *upward* as follows: for $v \in L(M_l^i, l-1)$ and suffix $z \in \{0, 1, \ldots, u_l\}$, we assign child $v_z = v + f_{i(l)}(z)$ through edge labeled z to breakpoint $b_m \in L(M_{l-1}^i, l)$ such that $b_{m-1} \prec v_z \preceq b_m$. Again we will represent edges $E(v, v')$ between vertices v, v' only by the extremes $\ell_{vv'}, r_{vv'}$ such that the only $z \in \{0, 1, \ldots, u_{l+1}\}$ that satisfy $\hat{M}_l^i(v, z) = v'$ are $\ell_{vv'} \leq z \leq r_{vv'}$.

We will not explicitly build these ROBPs, we will just use them to assist in our proofs, which use induction over the ROBPs $M_n^i, M_{n-1}^i, \ldots, M_0^i$, and thus prove the claim for $M_0^i = \hat{M}^i$ as needed. We prove that each $\hat{M}_j^i$ remains an interval ROBP through the following claims: ☐

Claim A.8. *For $v \in L(\hat{M}_j^i, \ell)$ and $0 \leq \ell < r \leq u_{l+1}$, $\hat{M}_j^i(v, \ell) \succeq \hat{M}_j^i(v, r)$.*

Proof. Note that this claim is only interesting if $\ell > i$, otherwise all vertices $v, \hat{M}_j^i(v, l), \hat{M}_j^i(v, r)$ exist in original interval ROBP M^i and thus the claim holds. Thus we can assume that layer $\ell + 1$ is rounded and consists of breakpoints. Take $b_{x_\ell} = \hat{M}_j^i(v, \ell)$ and $b_{x_r} = \hat{M}_j^i(v, r)$

Note that for any "true" child $v_z = \hat{M}_{j+1}^i(v, z)$, we replace the child through edge z with breakpoint b_{x_z} such that $b_{x_z} \succeq v \succ b_{x_z}$. This is exactly the the maximum partial sum breakpoint with lower partial sum than v_z. Thus, we can follow identical logic to claim A.1 to prove the claim. ☐

Claim A.9. *Let $v, v' \in L(\hat{M}_j^i, l)$, $v \succeq v'$. For suffix z, $\hat{M}_j^i(v, z) \succeq \hat{M}_j^i(v', z)$.*

Proof. Note that we need only prove this claim for suffixes of length one, as the property for any longer suffix will follow inductively.

Also note as in the previous theorem, this claim is only interesting if $l > j$, otherwise vertices $v, v', \hat{M}_j^i(v, z), \hat{M}_j^i(v', z)$ exist in original interval ROBP M^i. Thus we can assume that layer $l + 1$ is rounded and consists of breakpoints.

Again since edges are re-assigned to breakpoints in the same fashion as section A.2, we can follow identical logic to claim A.2 to prove the claim. $\square$

We now analyze how the probability of acceptance (and thus the total number of solutions) in our new ROBP $\hat{M}^i$ compares to that of the original ROBP M^i.

Theorem A.10. *For $v \in L(\hat{M}^i_j, \ell)$, we have $A_{\hat{M}^i_{j+1}}(v) \subseteq A_{\hat{M}^i_j}(v)$. Further, for every $w \in \mathcal{W}^\ell$, we have $P_{\hat{M}^i,w}(v) \leq P_{\hat{M}^i_j,w}(v) \leq P_{\hat{M}^i,w}(v)(1+\delta)^{n-\ell}$*

Proof. The first claim is rather straightforward. Note that for any vertex $v \in L(\hat{M}^i_j, \ell)$ where $\ell > j$, both ROBPs make identical transitions from v, and so $A_{\hat{M}^i_{j+1}}(v) = A_{\hat{M}^i_j}(v)$. Let $\ell = j$. Let $b_z = \hat{M}^i_j(v, z)$ be the children of v for each $z \in \{0, 1, \ldots, u_{\ell+1}\}$. Note that since $\hat{M}^i_j$ is obtained from $\hat{M}^i_{j+1}$ by rounding layer $j+1$, these vertices are breakpoints, and for the "true" children of v, there are vertices $v_z = \hat{M}^i_{j+1}(v, z)$ in layer $j+1$ of $\hat{M}^i_{j+1}$ such that $v_z \preceq b_z$. Note that by our previous theorem, suffixes maintain ordering, and so $A_{\hat{M}^i_{j+1}}(\hat{M}^i_{j+1}(v, z)) = A_{\hat{M}^i_{j+1}}(v_z) \subseteq A_{\hat{M}^i_j}(b_z) = A_{\hat{M}^i_{j+1}}(\hat{M}^i_j(v, z))$. Thus the set of accepting suffixes can only increase for any choice of z, and so the claim holds. Now for any $\ell < j$, the claim follows because $\hat{M}^i_j$ and $\hat{M}^i_{j+1}$ are identical up to layer j, and since we have proven the claim for layer j, the claim follows here as well.

Note that since $A_{\hat{M}^i_{j+1}}(v) \subseteq A_{\hat{M}^i_j}(v)$, it is clear that $P_{\hat{M}^i_{j+1},w}(v) \leq P_{\hat{M}^i_j,w}(v)$ for all w, and thus $P_{\hat{M}^i,w}(v) = P_{\hat{M}^i_n,w}(v) \leq P_{\hat{M}^i_{n-1},w}(v) \leq \cdots \leq P_{\hat{M}^i_j,w}(v)$

Now we will prove the latter claim by induction on the ROBPs $\hat{M}^i_n, \hat{M}^i_{n-1}, \ldots, \hat{M}^i_0$, by showing that $P_{\hat{M}^i_j,w}(v) \leq P_{\hat{M}^i_{j+1},w}(v)(1+\delta)$ for every $v \in L(\hat{M}^i_j, \ell)$, $w \in \mathcal{W}^\ell$.

Note that when proving this claim for $\hat{M}^i_j$, any layer $\ell > j$ is identical in $\hat{M}^i_{j+1}$, and so the claim is obvious. Also note that if we prove the claim for vertices in layer $\ell = j$, the claim will follow for layers in $\ell < j$, because these layers are also identical in $\hat{M}^i_{j+1}$, and thus make identical transitions and so we can express both $P_{\hat{M}^i_j,w}(v)$ and $P_{\hat{M}^i_{j+1},w}(v)$ as the same convex combination of acceptance probabilities of vertices in layer j.

Let the layer $\ell = j$, and consider vertex $v \in L(\hat{M}^i_j, j)$. We find that probability $P_{\hat{M}^i_j,w}(v)$ is equal to the sum, over edges $z \in \{0, 1, \ldots, u_{j+1}\}$ out of v, of the probability we take that same edge label z out of w, times the probability of accepting if we start from $\hat{M}^i_j(v, z)$ and make transitions in $\hat{M}^i_j$ according to a suffix sampled from D^{w_z} (where $w_z = D(w, z)$ is the child of w in D through edge labeled z): ⁕ Kyra: [sentence is much too long. How to shorten it? only say it in math?]

$$P_{\hat{M}^i_j,w}(v) = \sum_{z=0}^{u_{j+1}} p_w(z) P_{\hat{M}^i_j,w_z}(\hat{M}^i_j(v, z))$$

We now aim to bound $P_{\hat{M}^i_j,w_z}(\hat{M}^i_j(v, z))$ for the children of v taken through each edge label z. Let b_1, b_4 be the adjacent breakpoints in $B^{i+1}(w_z)$ such that $b_1 \prec v + f_{i(j+1)}(z) \preceq b_4$, and let b_2, b_3 be

the adjacent breakpoints in B^{i+1} such that $b_2 \prec v + f_{i(j+1)}(z) \preceq b_3$, and note that $\hat{M}^i_j(v,z) = b_3$ by the way we connected layer j to $j+1$ in $\hat{M}^i_j$. Also note that since both layer j and $j+1$ of $\hat{M}^i_{j+1}$ are identical to $\hat{M}^i$, we have that $v + f_{i(j+1)}(z) = \hat{M}^i_{j+1}(v,z)$.
Since $B^{i+1}(w_z) \subseteq B^{i+1}$ we have $b_1 \preceq b_2 \prec \hat{M}^i_{j+1}(v,z) \preceq b_3 \preceq b_4$. By the way we chose breakpoints and assign edges when rounding layer $j+1$, we have that

$$P_{\hat{M}^i_{j+1},w_z}(b_4) \leq (1+\delta)P_{\hat{M}^i_{j+1},w_z}(\hat{M}^i_{j+1}(v,z))$$

Since $\hat{M}^i_{j+1}$ is an interval ROBP, we have that $P_{\hat{M}^i_{j+1},w_z}(b_3) \leq P_{\hat{M}^i_{j+1},w_z}(b_4)$ and thus

$$P_{\hat{M}^i_{j+1},w_z}(b_3) \leq (1+\delta)P_{\hat{M}^i_{j+1},w_z}(\hat{M}^i_{j+1}(v,z))$$

Since the breakpoints are a subset of the vertices in layer $j+1$ of $\hat{M}^i_{j+1}$, they exist in both $\hat{M}^i_j$ and $\hat{M}^i_{j+1}$, and so $P_{\hat{M}^i_{j+1},w_z}(b_3) = P_{\hat{M}^i_j,w_z}(b_3)$. Recall that $\hat{M}^i_j(v,z) = b_3$, and thus

$$P_{\hat{M}^i_j,w_z}(\hat{M}^i_j(v,z)) \leq (1+\delta)P_{\hat{M}^i_{j+1},w_z}(\hat{M}^i_{j+1}(v,z))$$

Thus, we have

$$P_{\hat{M}^i_j,w}(v) \leq \sum_{z=0}^{u_{j+1}} p_w(z)(1+\delta)P_{\hat{M}^i_{j+1},w_z}(\hat{M}^i_{j+1}(v,z)) = (1+\delta)P_{\hat{M}^i_{j+1},w}(v)$$

Thus we have proven the claim, and it follows that

$$P_{\hat{M}^i_j,w}(v) \leq P_{\hat{M}^i_{j+1},w}(v)(1+\delta) \leq \cdots \leq P_{\hat{M}^i_n,w}(v)(1+\delta)^{n-j} = P_{\hat{M}^i,w}(v)(1+\delta)^{n-\ell}$$

for every $v \in L(\hat{M}^i_j, \ell)$, $w \in \mathcal{W}^j$. □

Thus we have shown that for every $v \in L(\hat{M}^i, \ell) = L(\hat{M}^i_0, \ell)$, and every $w \in \mathcal{W}^\ell$ we have that

$$P_{\hat{M}^i,w}(v) \leq P_{M^i,w}(v)(1+\delta)^n$$

Note that we have Taylor expansion $(1+\delta)^n = 1 + n\delta + \frac{n(n-1)}{2}\delta^2 + \ldots$ so if we take $\delta = \Omega(\frac{\eta}{2n})$, which is small, the second term $n\delta$ is much larger than the remainder of the terms, so we have $(1+\delta)^n \leq 1 + 2n\delta = 1 + 2n\frac{\eta}{2n} = 1 + \eta$, and so

$$P_{M^i,w}(v) \leq P_{\hat{M}^i,w}(v) \leq P_{M^i,w}(v)(1+\eta)$$

for every v and w as we will use in the following theorems.

Theorem A.11. *We can construct ROBP $\hat{M}^i$ in time $O(n^{O(k)}(\log U)^2 \log b/\eta^2)$*

Proof. Observe that for every vertex v in layer ℓ and every $w \in \mathcal{W}^\ell$ we have that $P_{\hat{M}^i,w}(v) \leq 1$ and $P_{\hat{M}^i,w}(v) \geq U^{-n}$. Let $N_w + 1$ be the number of breakpoints in $B^\ell(w)$. For $1 < j \leq N_w$, every breakpoint $b_{w(j)}$ changes from $b_{w(j-1)}$ by at least a factor of $\frac{1}{1+\delta}$, we have that

$$U^{-n} \leq P_{\hat{M}^i,w}(b_{w(N_w)}) < P_{\hat{M}^i,w}(b_{w(1)})/(1+\delta)^{N_w} \leq (\frac{1}{1+\delta})^{N_w}$$

This implies that $U^n \geq (1+\delta)^{N_w}$, and so $\log U \cdot n \geq \log(1+\delta) \cdot N_w$, or $N_w \leq \frac{n \log U}{\log(1+\delta)}$. Taylor series approximation tells us that for $\delta < 1$, we have $\log(1+\delta) \geq \delta - \frac{\delta^2}{2}$ so $N_w \leq \frac{n \log U}{\delta - \frac{\delta^2}{2}} \leq \frac{n \log U}{\delta - \frac{\delta}{2}} \leq \frac{n \log U}{\frac{\delta}{2}} \leq \frac{2n \log U}{\delta}$, and so $|B^\ell(w)| \leq \frac{2n \log U}{\delta} + 1$, and so $|B^\ell| = |\cup_{w \in \mathcal{W}^\ell} B^\ell(w)| \leq \sum_{w \in \mathcal{W}^\ell} |B^\ell(w)| \leq s \cdot \frac{2n \log U}{\delta} + s = \frac{2ns \log U}{\delta} + s = O(\frac{4n^2 s \log U}{\eta})$, where s is the width of small space source D.

Now note that for layer $\ell = n$, we create two vertices, 0 with $P_{\hat{M},w}(0) = 1$ for all $w \in \mathcal{W}^n$ and $b_i + 1$ with $P_{\hat{M},w}(b_i + 1) = 0$ for all $w \in \mathcal{W}^n$ in constant time.

We will now discuss the runtime of building layer ℓ from layer $\ell + 1$. We maintain that the vertices of the layer below, $L(\hat{M}^i, \ell + 1) = B^{\ell+1}$ are known and stored in a binary tree along with the values $P_{\hat{M}^i,w}(b)$ for every $b \in B^{\ell+1}$ and $w \in \mathcal{W}^{\ell+1}$.

We now prove that the statement holds for any $w \in \mathcal{W}^\ell$. For fixed $w \in \mathcal{W}^\ell$, we can compute the first breakpoint vertex of layer ℓ of $\hat{M}^i$ as $b_{w(1)} = 0$. Now assume inductively that we have breakpoints $b_{w(1)}, \ldots, b_{w(j)}$, and we now aim to compute $b_{w(j+1)}$.

Recall that $v_{w(j+1)}$ is the maximal $v \prec v_{w(j)}$ satisfying $P_{\hat{M}^i,w}(v) < P_{\hat{M}^i,w}(v_{w(j)})/(1+\delta)$. Thus, to find $b_{w(j+1)}$ we can now do binary search through values of $P_{\hat{M}_{i+1},w}(v)$. We require $0 < P_{\hat{M}_{i+1},w}(v)$, so we need only search through vertices in the range $[0, b]$ since $b+1$ can never lead to an accepting suffix due to nonnegativity of each f_j. This requires $O(\log b)$ computations of $P_{\hat{M}_{i+1}}(v)$.

For each partial sum v that we encounter during the search we must calculate $P_{\hat{M}^i,w}(v)$ and check if it exceeds $P_{\hat{M}^i,w}(v_{w(j)})/(1+\delta)$.

Here we have already calculated and stored $P_{\hat{M}^i,w}(v_{w(j)})$. We can calculate

$$P_{\hat{M}^i,w}(v) = \sum_{z=0}^{u_{j+1}} p_w(z) P_{\hat{M}^i,w_z}(\hat{M}^i(v,z))$$

where $w_z = D(w,z) \in \mathcal{W}^{\ell+1}$. Note that $|L(\hat{M}^i, \ell+1)| \leq \frac{4n^2 s \log U}{\eta}$, and this bounds the number of possible children v can have in layer $\ell + 1$, and thus bounds the number of previously calculated probabilities that we need to access and sum over. Also note that since we store these known probabilities in a binary tree, we need a factor of $\log(\frac{4n^2 s \log U}{\eta})$ to access the information. Thus we can calculate this probability in time $O(\frac{n^2 s \log U}{\eta} \log(\frac{n^2 s \log U}{\eta}))$.

We repeat this process throughout the binary search iterating through the ordered partial sums, at most $\log b$ times, until we find the breakpoint $v_{w(j+1)}$, so we have runtime $O(\frac{n^2 s \log U \log b}{\eta} \log(\frac{n^2 s \log U}{\eta}))$. As shown above, we have at most $O(n^2 \log U/\eta)$ breakpoints for each $w \in \mathcal{W}^\ell$, and there are at most

$s = n^{O(k)}$ many such w, and so we can find all vertices in this layer in time $O(\frac{n^{O(k)}(\log U)^2 \log b}{\eta^2} \log(\frac{n^{O(k)} \log U}{\eta}))$. Thus we can compute the entire tree (all $n+1$ layers) in time $O(\frac{n^{O(k)}(\log U)^2 \log b}{\eta^2} \log(\frac{n^{O(k)} \log U}{\eta}))$. $\square$

Theorem A.12. *Given a collection of the (W, n)-ROBPs $\hat{M}^i$ for $i \in [k]$ and (s, n)-small space source D, we can create a (W^k, n)- ROBP $\hat{M}$ computing the intersection of these, i.e. for any $x \in \{0, 1, \ldots, u_1\} \times \{0, 1, \ldots, u_2\} \times \cdots \times \{0, 1, \ldots, u_n\}$, $\hat{M}(x) = 1$ if and only if $\hat{M}^i(x) = 1$ for all $i \in [k]$.*

Proof. First, to create the desired ROBP $\hat{M}$, we need $n+1$ layers, where layer 0 has a single vertex $(0, 0, \ldots, 0)$ which is a k-tuple of partial sum 0, as 0 is the partial sum of the start vertex in each $\hat{M}^i$. Now from layer $\ell - 1$ we will build layer ℓ as follows:

For each vertex $v = (v_1, v_2, \ldots, v_k) \in L(\hat{M}, \ell - 1)$ represented by a k-tuple of partial sums, and for each $z \in \{0, 1, \ldots, u_\ell\}$, we construct child vertex through edge labeled z, which has partial sum tuple $(\hat{M}^1(v_1, z), (\hat{M}^2(v_2, z), \ldots, (\hat{M}^k(v_k, z))$. We maintain that the i^{th} member of this tuple is associated with a vertex in $\hat{M}^i$, and that the prefix taken in $\hat{M}$ to reach vertex $v = (v_1, \ldots, v_k)$ is exactly the prefix taken in $\hat{M}^i$ to reach v_i.

Note that since there are at most W vertices in each layer of $\hat{M}^i$, we can make at most W^k partial sum tuple vertices in each layer of $\hat{M}$. Thus $\hat{M}$ has width W^k, and can be generated in time $O(W^k)$.

Now to generate the probabilities in this new ROBP, we wish to maintain that for any vertex v in any layer $L(\hat{M}, \ell)$ and any $w \in \mathcal{W}^\ell$, the probability of accepting if we start from $v = (v_1, v_2, \ldots, v_k)$ and make transitions in $\hat{M}$ according to a suffix sampled from D^w is equal to the probability of the following happening for all $i \in [k]$: we accept starting from v_i and making transitions in $\hat{M}^i$ according to a suffix sampled from D^w.

We can compute these starting from layer n.

At layer n we know that for any $w \in \mathcal{W}^n$, vertex $v = (v_1, \ldots, v_k)$ only accepts if all $v_1, \ldots v_k$ accept, i.e. they have probability 1 of accepting in their respective $\hat{M}^i$s. If there exists any i such that v_i does not have probability 1 of accepting in $\hat{M}^i$ (i.e. is not the vertex 0), then we reject and have probability 0.

For any layer above, we assume we have the probabilities $P_{\hat{M},w}$ computed for every vertex in the layers below. For vertex $v \in L(\hat{M}, \ell)$, and for each $w \in \mathcal{W}^\ell$ we can compute

$$P_{\hat{M},w}(v) = \sum_{z=0}^{u_{\ell+1}} p_w(z) P_{\hat{M},w_z}(\hat{M}(v, z))$$

which requires the summing over at most W^k known values, since there are at most W^k unique children $\hat{M}(v, z)$. Thus we calculate the probabilities for every w and every v in a layer in $O(W^{2k} s)$ time and thus the entire tree (all n layers) in the time $O(nW^{2k} s)$. $\square$

This tells us that given $\hat{M}^i$ of width $O(n^{O(k)} \log U/\eta)$ and D of width $n^{O(k)}$, we can compute $\hat{M}$ and all associated acceptance probabilities in $O(n^{O(k^2)}(\log U/\eta)^{O(k)})$ time. Since we need $O(\frac{n^{O(k)}(\log U)^2 \log b}{\eta^2} \log(\frac{n^{O(k)} \log U}{\eta}))$ to build k many $\hat{M}^i$ prior to building $\hat{M}$, our total runtime is $O(n^{O(k^2)}(\log U/\eta)^{O(k)} \log b)$.

Finally, we can prove that this ROBP rounding algorithm provides the guaranteed of theorem 1.9.

Proof of theorem 1.9. First we call on theorem 1.8 to construct altered sets S_i'. Since our sets have width at most $\sum_{j=1}^{n} h_{ij}(u_j) + 2n^2 \leq 2n^3$, it follows that the ROBP $M_{S'}$ computing the indicator function over S' has width $O(n^{3k})$. This is because we can represent each state in layer l of $M_{S'}$ as a k-tuple of partial sums $\{\sum_{j=1}^{l} h_{1j}(x_j), \ldots, \sum_{j=1}^{l} h_{kj}(x_j)\}$ for which there are at most $O(n^{3k})$ possibel values. We use the uniform distribution over solutions in $S' = \cap_{i\in[k]} S_i'$ as our small space source D. Note that claim A.6 verifies this is a small space source of width $n^{O(k)}$.

For $i \in [k]$, let M^i be a (W, n)-ROBP exactly computing the indicator function for S_i. Let $\eta = O(\varepsilon/k(n+1)^k)$. Now, for every $i \in [k]$, by theorem A.7, we can explicitly in time $O(n^{O(k)}(\log U)^2 \log b/\eta^2)$ construct a $(n^{O(k)} \log U/\eta, n)$-ROBP $\hat{M}^i$ such that

$$\Pr_{x \leftarrow D}[\hat{M}^i(x) \neq M^i(x)] \leq \eta$$

Let $\hat{M}$ be the $(n^{O(k^2)}(\log U)^k/\eta^k, n)$-ROBP computing the intersection of all $\hat{M}^i$ for $i \in [k]$ as described in Theorem A.12. We have $\hat{M}(x) = \wedge_i \hat{M}^i(x)$, and so by union bound

$$\Pr_{x \leftarrow D}[\hat{M}(x) \neq \wedge_i M^i(x)] \leq k\eta$$

On the other hand, by Theorem 1.8,

$$\Pr_{x \leftarrow D}[\wedge_i M^i(x) = 1] \geq \frac{1}{2n^k}$$

Therefore, from the above two equations and setting $\eta = \frac{\varepsilon}{4kn^k}$, we get that

$$\Pr_{x \leftarrow D}[\hat{M}(x) = 1] \leq \Pr_{x \leftarrow D}[\wedge_i M^i(x) = 1] \leq (1+\varepsilon) \Pr_{x \leftarrow D}[\hat{M}(x) = 1]$$

Thus,

$$\Pr_{x \in_u \{0,\ldots,u_1\} \times \cdots \times \{0,\ldots,u_n\}}[x \in S'] \cdot \Pr_{x \leftarrow D}[\hat{M}(x) = 1] = \frac{|S|}{(u_1+1)\ldots(u_n+1)} \cdot \Pr_{x \leftarrow D}[\hat{M}(x) = 1]$$

is an ε-relative error approximation to the fraction of solutions to all constraints $\Pr_{x \in_u \{0,\ldots,u_1\} \times \cdots \times \{0,\ldots,u_n\}}[\wedge_i M^i(x) = 1] = \Pr_{x \in_u \{0,\ldots,u_1\} \times \cdots \times \{0,\ldots,u_n\}}[x \in S'] \cdot \Pr_{x \leftarrow D}[\wedge_i M^i(x) = 1]$.

The theorem now follows since we can compute $\frac{|S|}{(u_1+1)\ldots(u_n+1)}$ in time $n^{O(k)}$ and $\Pr_{x \leftarrow D}[\hat{M}(x) = 1]$ in time $O(n^{O(k^2)}(\log U/\varepsilon)^{O(k)} \log b)$ using theorem A.12. $\square$