Asymptotics of Random-Lattice Enumeration: Beyond HKZ via Dense Sublattices

Thijs Laarhoven

NXP Semiconductors, Eindhoven, The Netherlands
mail@thijs.com

Abstract. To bridge the gap between worst-case and heuristic analyses of lattice enumeration, we study sizes of (unpruned) enumeration trees on n-dimensional Haar-random lattices. Writing $n^{cn+o(n)}$ for the tree size at the shortest vector radius, we investigate the asymptotic scaling of c, and which bases minimize it.

For quasi-HKZ bases, we rigorously analyze the random-lattice HKZ geometry and establish $c = 1/(2e) \approx 0.1839$. This matches the worst-case enumeration exponent of Hanrot–Stehlé and confirms earlier predictions for the geometry of HKZ bases on random lattices.

Searching beyond quasi-HKZ bases, we prove that with high probability, the minimum achievable tree size exponent c over all bases of a random lattice satisfies:

$$0.1464 \approx \frac{2-\sqrt{2}}{4} \leq c < 0.1474.$$

For the lower bound, we use Kim's mean-value formula for nested chains of primitive sublattices to rule out bases with small prefix determinants. This strengthens the earlier Rankin-based heuristic lower bound exponent $c \geq 1/8$, discussed by Nguyen and Hanrot–Stehlé, and rules out the existence of random-lattice bases with trees as small as predicted by the geometric series assumption.

For the upper bound, we show that with high probability a dense primitive half-rank sublattice exists with determinant at most $n^{n/8+o(n)}$. Combining dual HKZ on this sublattice with HKZ on its quotient gives a basis with $c \leq 1/(4\sqrt{e}) \approx 0.1516$. Recursively splitting each half lowers the upper bound below 0.1474, and under a natural heuristic assumption we prove that the lower bound can be attained with suitable sublattice splits.

After lattice-dependent preprocessing, with these bases we can solve the closest vector problem with preprocessing faster with polynomial-size advice: classical enumeration gives query time $n^{0.1474n+o(n)}$, collision search runs in time $n^{0.0892n+o(n)}$, and quantum enumeration in time $n^{0.0737n+o(n)}$. The two-block basis can be used for discrete Gaussian sampling at any width in time $n^{0.1783n+o(n)}$ and polynomial space, with statistical error at most e^{-n^3}.

Keywords: Random lattices · Enumeration · Dense sublattices · Discrete Gaussian sampling

1 Introduction

Currently, one of the main challenges in cryptography is the migration to quantum-safe cryptography. After decades of research, lattice-based cryptography has emerged as a prime candidate due to its favorable performance trade-offs over other proposals. The security of lattice-based schemes relies on the hardness of high-dimensional lattice problems such as the shortest vector problem (SVP) and the closest vector problem (CVP). A good understanding of the best algorithms for solving these problems is critical for designing secure and efficient lattice-based cryptographic primitives.

For solving hard lattice problems, perhaps the most important tool in the cryptanalyst's toolbox is the BKZ algorithm: a blockwise basis reduction algorithm turning a "bad" lattice basis into a "good" one [48,13]. The BKZ algorithm relies on an exact SVP subroutine to do the heavy lifting; within lower-dimensional projected sublattices (blocks), an SVP oracle finds a shortest lattice vector, which is then used to improve the overall basis. The block size controls the cost of the algorithm and the quality of the basis, and the overall costs of BKZ are often estimated to be proportional to solving SVP in random lattices whose dimension equals the block size. As such, the costs of exact SVP oracles directly impact the estimated costs of BKZ.

For solving exact SVP, there are two main lines of work. Lattice enumeration [20,48,24,13] can be seen as a brute-force search through all potential combinations of basis vectors that might produce a shortest vector. By enumerating all combinations and keeping track only of the shortest vector found so far, this algorithm only requires polynomial memory, but its running time in dimension n scales as $n^{\Theta(n)}$ [32,33]. Its competitor, lattice sieving [2,4,9], is a saturation algorithm which first samples a lot of lattice vectors, and then repeatedly tries to form shorter combinations until it finds a shortest vector. This method runs in time $2^{\Theta(n)}$, but additionally requires $2^{\Theta(n)}$ space for storing all the sampled vectors. Currently, the highest records in the SVP challenge database [47] were achieved using lattice sieving. However, for a long time lattice enumeration clearly dominated this same database, and improvements to enumeration may shift the balance back in its favor.

In this work, we study the main ingredient dominating the costs of lattice enumeration: the size of the full (unpruned) enumeration tree when searching at a radius proportional to the shortest vector norm. In particular, we focus on the sizes of these trees on Haar-random lattices of dimension n and determinant one. For a basis $\mathbf{B}$ of a random lattice $\mathcal{L}$, write:

$$T(\mathbf{B}) = \#\{\text{nodes in the full enumeration tree at radius } \lambda_1(\mathcal{L})\} = n^{cn+o(n)}. \tag{1.1}$$

We are interested in finding bases $\mathbf{B}$ which minimize c for Haar-random lattices.[1]

1.1 Related work

A classical approach for obtaining a basis with small enumeration trees is HKZ reduction [37]. To reduce an n-dimensional basis, this greedy approach selects a shortest non-zero vector as its first basis vector, projects away its span, and calls HKZ on the remainder (and lifts the result back to the original lattice). Hanrot–Stehlé established that, even for worst-case lattices, the tree size exponent for (quasi-)HKZ bases satisfies $c \leq 1/(2e)$ [26,28]. The same authors showed that this worst-case bound is tight [27] and conjectured that the same exponent applies to random lattices [28, Section 6]. The heuristic HKZ model of Hanrot–Pujol–Stehlé [25, Section 4.1] and its subsequent refinement by Haque–Pieprzyk [29] support this prediction, but rely on unproven heuristic assumptions about

[1] Note that for random lattices, with high probability $\lambda_1(\mathcal{L}) = (1+o(1)) \cdot \sqrt{n/(2\pi e)}$. As the leading constant c in the enumeration tree size does not change when multiplying the radius with $1+o(1)$, we can equivalently replace $\lambda_1(\mathcal{L})$ in the definition of $T(\mathbf{B})$ and c with the Gaussian heuristic.

the geometry of HKZ-reduced bases. Aono–Nguyen–Seito–Shikata [5] experimentally observed that HKZ bases may be suboptimal for minimizing the tree sizes. In general, it is not known whether HKZ bases are asymptotically optimal (up to lower order terms), or if bases exist achieving $c < 1/(2e) - \varepsilon$ for fixed $\varepsilon > 0$ and large n.

Enumeration tree size estimates of $n^{cn+o(n)}$ naturally translate to classical enumeration complexities for the shortest vector problem (SVP) of $n^{cn+o(n)}$ time (and $n^{O(1)}$ space), by searching the entire tree for a shortest non-zero vector $\boldsymbol{s}$ of norm $\|\boldsymbol{s}\| = \lambda_1(\mathcal{L})$. By suitably pruning the tree [6,24,5], in practice one can shave off valuable factors $2^{\Theta(n)}$ from the search costs, but known pruning techniques do not affect the asymptotics of the constant c in the time exponent. Recent work of Feng–Gao–Liu [19] demonstrated that by a collision search procedure for a pair $(\boldsymbol{v}, \boldsymbol{w})$ with $\boldsymbol{v} - \boldsymbol{w} = \boldsymbol{s}$, one can optimistically reduce the enumeration cost exponent in half; more details on this method will follow later. Building on quantum backtracking of Montanaro [43], Aono–Nguyen–Shen [6] further showed how to search a lattice enumeration tree in time proportional to the square root of its size, up to polynomial factors. Their method also applies to exact CVP and uses polynomial space.

When considering enumeration trees for other bases, past results have hinted that one may hope to do better. The best known heuristic lower bound of Nguyen for arbitrary bases [45, "A Heuristic Lower Bound"] says $c \geq 1/8$ for all bases, which leaves a clear gap with the heuristic estimate of $c = 1/(2e)$ for HKZ bases. In the idealized setting of a basis following the geometric series assumption [49], one would actually obtain $c = 1/8$ [28, Section 6, p. 27]. This reduced tree size for GSA bases was cleverly used by Albrecht et al. [3] to obtain a faster BKZ algorithm. However, for random lattices we generally do not expect bases with such a GSA-profile (with $c = 1/8$) to exist.

A concrete direction for finding bases with smaller enumeration trees is based on finding dense sublattices. This approach was already suggested by Gama–Howgrave-Graham–Koy–Nguyen and Hanrot–Stehlé [23,28]: search for a half-rank sublattice with small determinant, split the lattice basis in two halves, and reduce each part separately. The same relation between small enumeration trees and dense sublattices was observed by Li–Nguyen [40]. An open question is whether using dense sublattices can help find bases with strictly smaller c than HKZ. Note that finding dense sublattices seems to be incredibly hard; currently the best known algorithm for finding a densest rank-k sublattice for $k = \Theta(n)$, due to Dadush–Micciancio [14, Theorem 3.1], runs in time $n^{O(n^2)}$ and space $2^{O(n)}$.

For using dense sublattices in cryptanalytic algorithms, Ducas–Loyer [18, Sections 3–6] recently studied BKZ with low-rank densest-sublattice oracles, especially ranks 2 and 3, using Kim's sublattice statistics [35] to predict the output slope. Their heuristic estimates give only small gains, insufficient to compensate for the loss of dimensions for free, and their experiments show that these estimates are optimistic. Their experiments further reveal strong clustering in dense-sublattice counts, illustrating the difficulty of inferring typical existence from first moments.

For discrete Gaussian sampling, we use the rejection-corrected version of Klein's sequential sampler [36] analyzed by Brakerski–Langlois–Peikert–Regev–Stehlé [11, Section 5]. Similar to enumeration, its cost can be described as a function of the Gram–Schmidt lengths of the supplied basis. Using an HKZ-basis and bounds on the worst-case HKZ-shape from [28], this currently gives a polynomial-space discrete Gaussian sampler at any width running in time $n^{n/(2e)+o(n)}$.

1.2 Contributions

In this work, we answer several of the open questions above. Can we rigorously prove asymptotics on c for quasi-HKZ bases on random lattices? Are HKZ bases asymptotically optimal, or can we find better bases with lower asymptotics for c? And if so, how small can c be, and can we prove sharper lower bounds on c than the heuristic bound $c \geq 1/8$? We go through the main contributions in

the order they are described in the paper. (Section 2 concerns general preliminaries, and Section 3 establishes uniform estimates relating enumeration tree sizes to prefix determinants.)

§4: Random-lattice HKZ geometry. Our first main contribution confirms heuristic predictions about the HKZ geometry of Hanrot–Pujol–Stehlé and Haque–Pieprzyk [25,29], as well as the conjecture of Hanrot–Stehlé [28]: for random lattices, quasi-HKZ bases have trees of size $n^{n/(2e)+o(n)}$. More precisely: with high probability over the randomness of $\mathcal{L} \sim \mu_n$, with μ_n denoting the Haar probability measure over the space $\mathrm{SL}_n(\mathbb{R})/\mathrm{SL}_n(\mathbb{Z})$ of n-dimensional lattices of determinant one, *all* quasi-HKZ bases of $\mathcal{L}$ have this tree size scaling. To prove this, we show that the Gram–Schmidt profile of quasi-HKZ bases on random lattices essentially matches what heuristics predicted, from which the constant c follows. This matches the asymptotics of Hanrot–Stehlé for the sizes of enumeration trees for quasi-HKZ bases on worst-case lattices.

Theorem 1.1 (Tree size of all quasi-HKZ bases). *With probability* $1-o(1)$ *over* $\mathcal{L} \sim \mu_n$*, every quasi-HKZ basis* $\mathbf{B}_{\mathrm{HKZ}}$ *of* $\mathcal{L}$ *has enumeration tree size:*

$$T(\mathbf{B}_{\mathrm{HKZ}}) = n^{c_{\mathrm{HKZ}} \cdot n + o(n)}, \qquad c_{\mathrm{HKZ}} = \frac{1}{2e} = 0.18393\ldots. \tag{1.2}$$

§5: Beyond HKZ via dense sublattices. Whereas Theorem 1.1 proves what heuristics already suggested, to the best of our knowledge the following concrete asymptotic improvement has not been established before: (1) showing that useful dense sublattices exist with high probability, and (2) showing that with these dense sublattices, one can indeed go strictly below the $c_{\mathrm{HKZ}} = 1/(2e)$ exponent of quasi-HKZ bases, thereby showing quasi-HKZ bases are asymptotically suboptimal in the leading constant c. This idea was previously suggested by Hanrot–Stehlé [28], and here we rigorously establish the asymptotic improvement over quasi-HKZ bases. Using a single dense primitive sublattice of half the rank, and optimizing for the volume of this sublattice so that with high probability it exists, we obtain an explicit limiting Gram–Schmidt profile and asymptotic tree size constant c_2.

Theorem 1.2 (Tree size of a two-block split). *With probability* $1-o(1)$*,* $\mathcal{L} \sim \mu_n$ *has a primitive sublattice* $\mathcal{M}$ *of rank* $m = \lfloor (n-1)/2 \rfloor$ *with* $\det \mathcal{M} \leq e^{O(n)} n^{n/8}$*. Combining a reversed dual-HKZ basis of* $\mathcal{M}$ *with lifts of an HKZ basis of its orthogonal quotient gives a basis* $\mathbf{B}_2$ *of* $\mathcal{L}$ *with enumeration tree size:*

$$T(\mathbf{B}_2) = n^{c_2 \cdot n + o(n)}, \qquad c_2 = \frac{1}{4\sqrt{e}} = 0.15163\ldots. \tag{1.3}$$

Note that if we split a typical HKZ basis halfway, it has a sublattice $\mathcal{M}$ with determinant $\det \mathcal{M} = n^{((\ln 2)/4 + o(1))n}$. This is exactly the boundary where a denser sublattice would lead to smaller trees. As finding dense sublattices is difficult, a first goal for a faster SVP algorithm may be to find *any* half-rank sublattice with density higher than what we get from HKZ.

A natural question to ask then is: why stop at a two-way split? Can we recursively look for dense sublattices in each half, and further reduce the constant c? Or can we immediately go for multi-block splits, to obtain smaller trees? Using recursive binary splits, the constant can be lowered below 0.1474, with the limit numerically approaching $0.14739700\ldots$.

Theorem 1.3 (Tree size of recursive binary splits). *With probability* $1-o(1)$*,* $\mathcal{L} \sim \mu_n$ *contains a sequence of dense sublattices* $\mathcal{M}_0 \subseteq \mathcal{M}_1 \subseteq \cdots \subseteq \mathcal{L}$ *such that, combining (lifted) HKZ and reversed dual-HKZ bases of (orthogonal quotients of) these blocks, we obtain a basis* $\mathbf{B}_{\mathrm{rec}}$ *of* $\mathcal{L}$ *with enumeration tree size:*

$$T(\mathbf{B}_{\mathrm{rec}}) \leq n^{c_{\mathrm{rec}} \cdot n + o(n)}, \qquad c_{\mathrm{rec}} < 0.14740. \tag{1.4}$$

§6: Heuristic sublattice splits. For the previous upper bounds, it is important to note that these are based on "safe" splits: for every two-way split, we make sure that with high probability the desired dense sublattice exists. Better subdivisions may be possible if we allow splits which are exponentially rare, but where there are exponentially many potential sublattices to consider. However, analyzing subdivisions with rare sublattices rigorously appears challenging, as it requires controlling correlations between different sublattices. For instance, showing that with the subdivision strategy there are $n^{0.002n^2}$ potential sublattices, each having a chance of $n^{-0.001n^2}$ of being sufficiently dense, is not sufficient to show that one dense sublattice exists with high probability, as sublattices are not independent; a small fraction of Haar random lattices may contain all the dense sublattices we are looking for.

In general, it may be natural to assume that if the expected numbers of chains align when averaged over all random lattices, then we expect such towers of primitive sublattices to exist for most random lattices. Below we informally state this heuristic assumption; a precise definition follows later.

Assumption 1 (First-moment sufficiency, informal). *Consider nested chains of primitive sublattices, with prescribed ranks and determinant bounds. Suppose that, for some $\delta > 0$, the expected number of such chains in a Haar random lattice is at least $n^{\delta n^2 + o(n^2)}$, and the same bound holds for every non-empty subchain obtained by omitting intermediate sublattices. Then a Haar-random lattice contains such a complete chain with probability $1 - o(1)$.*

Ducas–Loyer [18, Section 6.2] observed strong clustering in rank-two dense-sublattice counts at determinant thresholds where the expected count is one. Assumption 1 concerns a different asymptotic regime: a fixed number of prescribed sublattice ranks, each proportional to n, with the expected count of every nonempty subchain at least $n^{\delta n^2}$ for some fixed $\delta > 0$ and all sufficiently large n. The assumption only predicts that the complete chain exists with high probability, without assuming Poisson statistics or concentration of the counts around their means. The observations of [18] do not contradict this statement, but emphasize that proving the existence of such chains in general may require justification beyond first moments.

Under this assumption, we show that we can slightly lower the previous upper bounds. Proving this result without the need for the heuristic assumption is left as an open problem.

Theorem 1.4 (Tree size of improved heuristic splits). *Suppose Assumption 1 holds, and let $\varepsilon > 0$ be fixed. Then with probability $1 - o(1)$ as $n \to \infty$, $\mathcal{L} \sim \mu_n$ contains a basis $\mathbf{B}$ satisfying:*

$$T(\mathbf{B}) \leq n^{(c_{\text{flat}}+\varepsilon)\cdot n + o(n)}, \qquad c_{\text{flat}} = \frac{2-\sqrt{2}}{4} = 0.14644\ldots. \tag{1.5}$$

§7: A sharp(er) tree size lower bound. To understand whether our improved upper bounds are (close to) optimal, we derived an improved asymptotic lower bound on c which with high probability over the randomness of the lattice holds simultaneously for all bases of this lattice. The previous best (heuristic) lower bound described by Nguyen [45] was of the form $c \geq 1/8 = 0.125$; we show this bound can be increased to $c \geq (2-\sqrt{2})/4 \approx 0.146$. This leaves a gap of slightly less than 0.001 in the constant c between our best proven upper and lower bounds, and this matches our best heuristic upper bound.

Theorem 1.5 (Tree size lower bound). *With probability $1 - o(1)$ over $\mathcal{L} \sim \mu_n$, every basis $\mathbf{B}$ of $\mathcal{L}$ satisfies:*

$$T(\mathbf{B}) \geq n^{c_{\text{flat}}\cdot n + o(n)}, \qquad c_{\text{flat}} = \frac{2-\sqrt{2}}{4} = 0.14644\ldots. \tag{1.6}$$

If the heuristic assumption holds, then this lower bound is asymptotically tight. Without this assumption, our upper and lower bounds leave a gap of less than 0.001 on the optimal c. The main results regarding the tree sizes are summarized and compared against related work in Table 1.

§8.1: Classical enumeration with preprocessing. Knowing that smaller trees exist, we can consider applications where smaller trees lead to faster algorithms. The smaller trees reduce the search cost once a suitable basis is available. Our uniform shifted counting bound extends this improvement to exact CVPP for arbitrary targets. With probability $1-o(1)$ over the lattice, polynomial-size classical advice supports deterministic queries in polynomial space and time

$$Q_{\mathrm{enum}} \leq n^{c_{\mathrm{rec}}n+o(n)}, \qquad c_{\mathrm{rec}} < 0.14740. \tag{1.7}$$

Under Assumption 1, the exponent can be replaced by $c_{\mathrm{flat}}+\varepsilon$ for every fixed $\varepsilon > 0$; see Corollary 8.1.

These bounds exclude the cost of basis construction. For proportional ranks, the exact densest-sublattice algorithm of Dadush–Micciancio [14, Theorem 3.1] takes time $n^{O(n^2)}$ and space $2^{O(n)}$. Finding the required sublattices efficiently enough to improve the total cost of SVP remains open. In CVPP, preprocessing depends only on the lattice and has unrestricted running time; only the resulting basis is retained as advice.

§8.2: Collision search beyond HKZ. In recent work, Feng–Gao–Liu [19] showed how to replace exhaustive enumeration over the entire enumeration tree by a polynomial-space collision search, to find a pair of vectors whose difference is a shortest vector. The relevant geometric factor in their analysis is $\sqrt{\mathcal{P}(\mathbf{B})}$, where $\mathcal{P}(\mathbf{B}) = \prod_i \max\{1, \|\boldsymbol{b}_i^*\|^{-1}\}$ for determinant-one lattices. Using quasi-HKZ bases, they then achieve an SVP cost of:

$$\sqrt{\mathcal{P}(\mathbf{B})} \leq n^{\tau_{\mathrm{HKZ}}\cdot n+o(n)}, \qquad \tau_{\mathrm{HKZ}} = \frac{1}{4e} = 0.09196\ldots. \tag{1.8}$$

We show that our lower bound carries over to their framework, with an additional factor 1/2. This bounds the geometric factor in the FGL analysis; a lower bound on the algorithm's actual running time would require a separate argument.

Theorem 1.6 (Collision search cost factor lower bound). *With probability* $1-o(1)$ *over* $\mathcal{L} \sim \mu_n$*, simultaneously for every basis* $\mathbf{B}$ *of* $\mathcal{L}$*,*

$$\sqrt{\mathcal{P}(\mathbf{B})} \geq n^{\tau_{\mathrm{low}}\cdot n+o(n)}, \qquad \tau_{\mathrm{low}} = \frac{2-\sqrt{2}}{8} = 0.07322\ldots. \tag{1.9}$$

To find the optimal basis minimizing the dominant cost, it is important to note that its asymptotic scale is controlled by the total variation of the normalized prefix profile, rather than its maximum height. Thus minimizing the largest enumeration level in the tree and minimizing the costs of this collision-search approach are different optimization problems. With the same binary split procedure, we do immediately obtain an improved cost estimate.

Theorem 1.7 (Collision search cost factor of a two-block split). *With probability* $1-o(1)$*,* $\mathcal{L} \sim \mu_n$ *contains a basis* $\mathbf{B}_2$ *with FGL costs:*

$$\sqrt{\mathcal{P}(\mathbf{B}_2)} \leq n^{\tau_2\cdot n+o(n)}, \qquad \tau_2 = \frac{1}{4\sqrt{e}} - \frac{1}{16} = 0.08913\ldots. \tag{1.10}$$

For the recursive and heuristic block profiles described above, we do not obtain improved cost estimates for collision search compared to a two-block split. These comparisons do not establish optimality among all bases or all configurations allowed by Assumption 1. Whether the lower bound on the FGL cost factor is attainable remains open.

We extend the collision procedure to exact CVPP for arbitrary targets, using the two-block basis as lattice-dependent advice.

Theorem 1.8 (Polynomial-space classical CVPP). *With probability* $1 - o(1)$ *over* $\mathcal{L} \sim \mu_n$*, there is a basis* $\mathbf{B}_2$ *of* $\mathcal{L}$ *such that, for every target* $\boldsymbol{t} \in \mathbb{R}^n$*, a randomized algorithm finds an exact closest lattice vector with probability at least* $2/3$*, using polynomial space and query time:*

$$Q_{\mathrm{coll}} \le n^{\tau_2 n + o(n)}. \tag{1.11}$$

§8.3: Quantum enumeration. Our smaller enumeration trees also give faster quantum CVPP algorithms. Applying quantum enumeration [6] halves the leading tree exponent. Thus recursive splits give a query exponent below 0.0737, and under Assumption 1 the exponent approaches $c_{\mathrm{flat}}/2$.

Theorem 1.9 (Polynomial-space quantum CVPP). *With probability* $1 - o(1)$ *over* $\mathcal{L} \sim \mu_n$*, there is polynomial-size classical advice depending only on* $\mathcal{L}$ *such that, for every target* $\boldsymbol{t} \in \mathbb{R}^n$*, a randomized quantum algorithm finds an exact closest lattice vector with probability at least* $2/3$*, using polynomial classical and quantum space and query time:*

$$Q_{\mathrm{quant}} \le n^{(c_{\mathrm{rec}}/2)n + o(n)} < n^{0.07370n + o(n)}. \tag{1.12}$$

Under Assumption 1, for every fixed $\varepsilon > 0$*, the query bound can be replaced by*

$$Q_{\mathrm{quant}} \le n^{(c_{\mathrm{flat}}/2 + \varepsilon)n + o(n)}, \qquad \frac{c_{\mathrm{flat}}}{2} = \frac{2 - \sqrt{2}}{8} = 0.0732233\ldots. \tag{1.13}$$

Note that the square-root improvements from classical collision search and quantum enumeration do not automatically compound. The FGL procedure [19] already succeeds with constant probability in time proportional to $\sqrt{\mathcal{P}(\mathbf{B})}$, so applying amplitude amplification to this procedure does not yield another square-root improvement. Quantum search over sample pairs likewise recovers this dependence from the overlap bounds in Section 8.2. The Brassard–Høyer–Tapp collision algorithm [12] achieves an additional speedup by storing a large table; restricting its memory to polynomial size removes this improvement in the leading exponent. Obtaining a smaller exponent by combining collision search with quantum computation would therefore require an additional algorithmic idea. We do not rule out such a future improvement, potentially exploiting the lattice structure.

The best known algorithmic complexities for CVPP with polynomial-sized advice are summarized and compared against related work in Table 2 (classical) and Table 3 (quantum). To the best of our knowledge, Theorems 1.8 and 1.9 currently provide the fastest (polynomial-space) classical and quantum CVPP complexities for random lattices, and may be of independent interest for cryptanalytic applications requiring solving large batches of CVP instances on random lattices [34,17,46].

§8.4: Discrete Gaussian sampling. The two-block basis also improves a bound for corrected Klein sampling that is uniform over all widths. For $s > 0$, let $D_{\mathcal{L},s}$ denote the centered discrete Gaussian on $\mathcal{L}$, assigning probability proportional to $e^{-\pi\|\boldsymbol{x}\|^2/s^2}$ to $\boldsymbol{x} \in \mathcal{L}$.

Theorem 1.10 (Polynomial-space discrete Gaussian sampling). *With probability* $1 - o(1)$ *over* $\mathcal{L} \sim \mu_n$*, there is a polynomial-size preprocessed basis, depending only on* $\mathcal{L}$*, such that for every width* $s > 0$ *one can sample within statistical distance at most* e^{-n^3} *of* $D_{\mathcal{L},s}$ *in time*

$$n^{\tau_{\mathrm{DGS}} n + o(n)} \operatorname{poly}(\ell), \qquad \tau_{\mathrm{DGS}} = 2\tau_2 = \frac{1}{2\sqrt{e}} - \frac{1}{8} = 0.17826\ldots, \tag{1.14}$$

and space polynomial in ℓ*, where* ℓ *is the combined bit length of the supplied basis and width.*

The basis is retained as advice, and its construction is excluded from the sampling bound. Section 8.4 proves the theorem using finite-precision corrected Klein sampling.

§9: Worst-case lattice outlook. Section 9 discusses the implications for worst-case lattices. Under an additional extension of the dense-sublattice heuristic to the conditioned Hanrot–Stehlé ensemble [27], we predict smaller enumeration trees for their examples after suitable splits. This motivates the following conjecture about the best basis of an arbitrary lattice.

Conjecture 1.11 (Improved worst-case enumeration bases). *There exist constants* $\delta > 0$ *and* $n_0 \geq 2$ *such that, for every* $n \geq n_0$ *and every unimodular rank-n lattice* $\mathcal{L}$*, there is a basis* $\mathbf{B}$ *of* $\mathcal{L}$ *satisfying* $T(\mathbf{B}) \leq n^{(c_{\mathrm{HKZ}} - \delta)n}$.

Visualizing the results. Figure 1 summarizes what enumeration trees look like for different bases, plotting the enumeration tree width against the prefix dimension (the reverse of the depth in the enumeration tree). The maximum of the curve determines the tree size exponent c, while the area under the curve is dictated by the area rule in Section 7.1, which essentially says the volume under the curve is at least $1/8$.

The lower bound has a geometric interpretation: after reducing the two outer blocks, a basis with tree coefficient at most c admits a prefix ceiling of area $c - c^2$. The area rule requires this area to be at least $1/8$, yielding $c \geq c_{\mathrm{flat}}$; Section 1.3 explains the argument.

Note that in Figure 1, recursive binary splits with a first split halfway have a fixed "dip" at $(1/2, 1/8)$, and consecutive splits do not alter this dip. Even with a large recursion depth, the gap between the binary splits and the plateau at $x = 1/2$ remains visible, which intuitively explains the gap between the lower bound and the best upper bound with recursive splits. The heuristic splits avoid this dip halfway, and make sure we approximate the plateau well at every value x.

An alternative visual representation of the different profiles is given in Figure 2, showing the Gram–Schmidt profiles at various (reversed) depths in the tree. This figure essentially shows the derivatives of the curves in Figure 1. The two-block profile jumps from $-1/4$ to $1/4$ at the split, corresponding to a factor $n^{1/2+o(1)}$ between adjacent Gram–Schmidt lengths. Note that good enumeration bases do not necessarily begin with a short vector, and better bases may even have superpolynomial first-vector lengths. The endpoint prediction and a rigorous first-vector lower bound are discussed in Section 5.2.

1.3 Techniques

Below we summarize some of the main technical tools to obtain the above main results. Our proofs combine uniform lattice-point estimates with averaging formulas for primitive vectors and sublattices.

Enumeration tree size bounds	**Lattice model**	c
HKZ bases [25,29]	Heuristic	0.184
Plateau profile (**Theorem 1.4**)	Heuristic	**0.146**
Rankin-based bound [45]	Heuristic (all bases)	≥ 0.125
Quasi-HKZ bases (**Theorem 1.1**)	Haar-random lattices	0.184
Two-block split bases (**Theorem 1.2**)	Haar-random lattices	0.152
Recursive binary split bases (**Theorem 1.3**)	Haar-random lattices	**0.147**
Arbitrary bases (**Theorem 1.5**)	Haar-random lattices (all bases)	$\geq$ **0.146**
Quasi-HKZ bases [30,33]	Worst-case lattices	0.500
Quasi-HKZ bases [28, Thm. 2.5]	Worst-case lattices	**0.184**
HKZ worst-case examples [27, Thm. 3]	Worst-case lattices (HKZ bases)	≥ 0.184

Table 1. Bounds on the constant c, rounded to three decimals, in the enumeration tree size $T(\mathbf{B}) = n^{cn+o(n)}$. Lower bounds are highlighted as $\geq c$, and explicitly state for which bases this lower bound holds. These bounds concern the sizes of full, unpruned enumeration trees, at a radius proportional to the shortest vector length.

Query times for classical poly-space CVPP	**Lattice model**	τ
Approximate Voronoi cells [16,38]	Heuristic	0.500
Large-tuple sieving [7]	Heuristic	0.500
Enumeration: Dual-HKZ bases [28]	Heuristic	0.184
Enumeration: Plateau profile (**Corollary 8.1**)	Heuristic	0.146
Collision search: Dual-HKZ bases [19]	Heuristic	**0.092**
Dual BDD with perturbations [15, Thm. 3.1]	Haar-random lattices	0.500
Enumeration: Recursive splits (**Corollary 8.1**)	Haar-random lattices	0.147
Collision search: Quasi-HKZ bases (**Theorem 1.1**)	Haar-random lattices	0.092
Collision search: Two-block split (**Theorem 1.8**)	Haar-random lattices	**0.089**
Collision search: Any basis (**Theorem 1.6**)	Haar-random lattices (FGL)	$\geq$ **0.073**
Compact Voronoi cells [31]	Worst-case lattices	2.000
Enumeration: Dual-HKZ bases [10,55]	Worst-case lattices	1.000
Kannan enumeration [33,28,22]	Worst-case lattices	**0.500**

Table 2. Bounds on the constant τ, rounded to three decimals, in the query time $Q(\mathbf{B}, \boldsymbol{t}) = n^{\tau n+o(n)}$ for classical CVPP algorithms using polynomial-size advice and query space. The lower bound is highlighted as $\geq \tau$, and concerns only the dominant FGL cost factor. The stated bounds include CVPP-adaptations of SVP and BDD algorithms, as well as heuristic extrapolations. Except for the first two rows, all rows allow arbitrary targets.

Query times for quantum poly-space CVPP	**Lattice model**	τ
Collision search: Dual-HKZ bases [19]	Heuristic	0.092
Quantum enumeration: HKZ bases [6]	Heuristic	0.092
Quantum enumeration: Plateau profile (**Theorem 1.9**)	Heuristic	**0.073**
Quantum enumeration: Two-block split (**Thm. 1.2; Lem. 8.2**)	Haar-random lattices	0.076
Quantum enumeration: Recursive splits (**Theorem 1.9**)	Haar-random lattices	**0.074**
Quantum enumeration: Kannan [28, Thm. 2.6]; [6, Thm. 6]	Worst-case lattices	**0.250**

Table 3. Bounds on the constant τ, rounded to three decimals, in the query time $Q(\mathbf{B}, \boldsymbol{t}) = n^{\tau n+o(n)}$ for quantum CVPP algorithms using polynomial-size advice and query space. All rows allow arbitrary targets; the Haar rows hold with probability $1 - o(1)$ over the lattice.

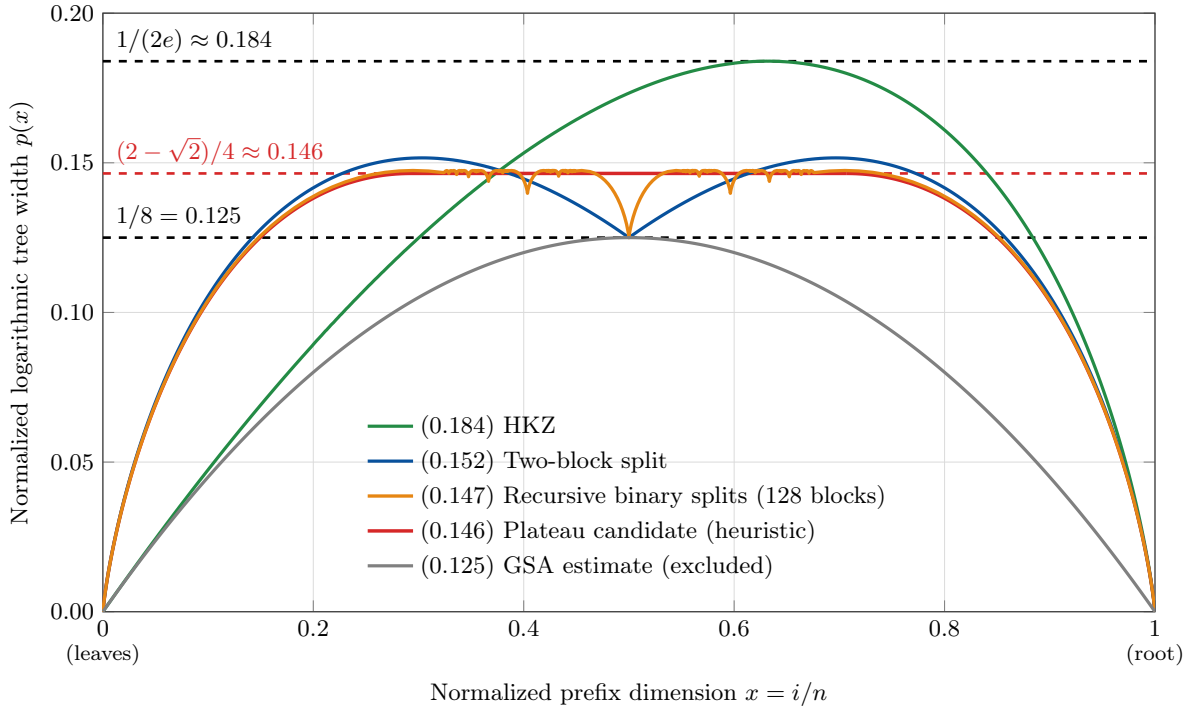


Fig. 1. Normalized tree width profiles. The solid red curve is the plateau candidate from (7.9), whose attainability is only heuristically proved. All non-gray profiles have area 1/8; the gray GSA estimate has area 1/12 and is excluded by the area rule in Section 7.1. The red curve is the limiting prefix scaling.

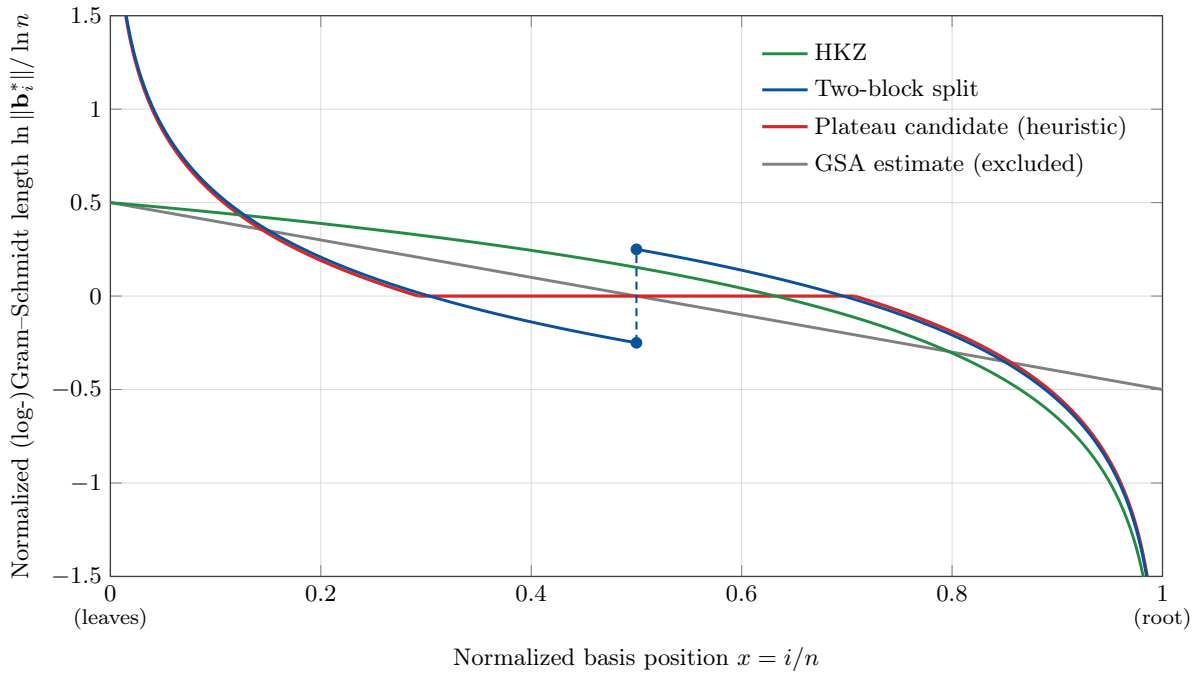


Fig. 2. Normalized logarithmic Gram–Schmidt profiles. Logarithmic singularities are clipped by the vertical plotting window. The red curve is the derivative of the limiting prefix scaling for the ideal profile; we do not prove that the actual Gram–Schmidt profiles converge to this curve. The GSA estimate is impossible for random lattices.

From prefix determinants to enumeration widths. Let P_k denote the logarithmic determinant of the sublattice generated by the first k basis vectors. We prove that, with high probability over the random lattice $\mathcal{L}$, and simultaneously for all bases $\mathbf{B}$ of $\mathcal{L}$,

$$\log T(\mathbf{B}) = \max_k P_k + O(n). \tag{1.15}$$

The lower estimate follows by averaging lattice-point counts over translated balls; the upper estimate uses Gaussian mass and Poisson summation. The key to uniformity is that, for a primitive prefix $\mathcal{M}$ and its orthogonal quotient $\mathcal{Q}$,

$$\mathcal{Q}^* = \mathcal{L}^* \cap \operatorname{span}(\mathcal{M})^\perp \subseteq \mathcal{L}^*. \tag{1.16}$$

Poisson summation therefore bounds every projected lattice-point count by $e^{P_k+\pi R^2}\rho(\mathcal{L}^*)$, where ρ denotes the Gaussian mass. Since $\mathbb{E}[\rho(\mathcal{L}^*)] = 2$ and $R^2 = O(n)$, a single high-probability bound on this mass controls every projection simultaneously.

Controlling adaptive HKZ projections. Successive projected lattices in HKZ reduction depend on previously selected shortest vectors. Their distribution therefore cannot simply be assumed to remain Haar. We control these dependencies through the following bound.

Lemma 1.12 (Moment bound; informal restatement of Lemma 4.2). *For a Haar-random lattice $\mathcal{L}$, consider the logarithmic shortfalls of projected shortest-vector lengths from twice their Gaussian-heuristic radii. Weight each shortfall by the remaining dimension, and let $Z_n(\mathcal{L})$ be the largest resulting sum over all HKZ bases. Then*

$$\mathbb{E}_{\mu_n}\left[e^{Z_n(\mathcal{L})/2}\right] \le e^{O(n^2)}. \tag{1.17}$$

In particular, with high probability this weighted sum is $O(n^2)$ simultaneously for all HKZ bases.

The key proof step replaces the maximum over possible shortest-vector choices by a sum over all primitive candidates. The primitive-vector averaging formula then allows us to integrate over the corresponding projected lattices and proceed inductively. This controls all HKZ choices without assuming that the selected quotients remain Haar. Together with deterministic HKZ bounds, the estimate yields the predicted bulk Gram–Schmidt geometry and a uniform prefix profile for all quasi-HKZ bases.

Dense sublattices with suitable blocks. To construct better bases, we need dense sublattices whose two resulting blocks also admit further reductions or splits. Our main existence result provides this additional control.

Lemma 1.13 (Existence of dense sublattices; informal restatement of Lemma 5.1). *Fix $0 < \alpha < 1$. With probability $1 - o(1)$, a Haar-random lattice $\mathcal{L}$ contains a primitive sublattice $\mathcal{M}$ of rank $m = \alpha n + O(1)$ with determinant satisfying:*

$$\det(\mathcal{M}) \le e^{O(n)} n^{m(n-m)/(2n)}. \tag{1.18}$$

Moreover, after rescaling $\mathcal{M}$ and its orthogonal quotient to determinant one, both blocks can satisfy any prescribed rotation-invariant measurable properties that hold with Haar probability $1 - o(1)$ in their respective dimensions.

First moments identify the relevant determinant scale, but overlapping candidates can make the second moment diverge. We remove candidates containing an unusually dense proper sublattice. The determinant cutoffs are chosen so that this removes only a vanishing fraction of the expected count, while controlling the contribution of intersecting pairs at every possible intersection rank. An averaging formula also bounds the expected number of candidates with unsuitable blocks. The resulting second-moment estimate proves that a suitable split exists with high probability. Prescribing that both blocks themselves admit further splits makes any fixed finite recursive construction possible.

From dense splits to smaller trees. To turn a dense sublattice into a small enumeration tree, we must also control the prefixes inside each resulting block. Lemma 4.1 provides a deterministic HKZ bound, valid for every lattice after normalization to determinant one, with profile $p_{\mathrm{H}}(t) = -(1-t)\ln(1-t)/2$. This profile can then also be used for smaller sublattices encountered from basis splits. Let p_n interpolate the normalized logarithmic prefixes $P_k/(n \ln n)$. A block of fixed relative width w, starting at x_0, whose boundary heights are at most u and v, satisfies after HKZ reduction

$$p_n(x_0 + wt) \leq (1-t)u + tv + wp_{\mathrm{H}}(t) + o(1). \tag{1.19}$$

Reversed dual-HKZ reduction replaces $p_{\mathrm{H}}(t)$ by $p_{\mathrm{H}}(1-t)$, reversing the profile. Thus block reduction adds a scaled HKZ curve to the linear interpolation between its boundary heights.

For a half-rank split with determinant $n^{n/8+o(n)}$, the boundary heights are 0, 1/8, and 0. Using reversed dual-HKZ on the left and HKZ on the right gives the upper profile

$$p_2(x) = \frac{v}{8} - \frac{v}{4}\ln v, \qquad v = 2\min\{x, 1-x\}. \tag{1.20}$$

Its maximum is $1/(4\sqrt{e})$. The prescribed typical block properties make this bound asymptotically tight. The same affine composition applies to recursively constructed child profiles, yielding the improved recursive upper bounds.

Joint constraints across ranks. Bounding sublattice determinants at each rank separately does not capture the requirement that the prefixes of a basis form a single nested chain. We exploit this additional structure through Kim's mean-value formula. For a fixed continuous function $h : [0,1] \to \mathbb{R}$, let $F_n(h; \mathcal{L})$ count complete chains of primitive sublattices whose rank-k member has determinant at most $n^{nh(k/n)}$. We obtain the following area rule (Proposition 7.1):

$$\mathbb{E}_{\mu_n}[F_n(h; \mathcal{L})] = n^{\left(2\int_0^1 h(x)\,dx - \frac{1}{4} + o(1)\right)n^2}. \tag{1.21}$$

If the area under h is less than 1/8, this expectation tends to zero. Markov's inequality then rules out every basis satisfying these prefix bounds, with high probability. This conclusion holds simultaneously for all bases, including bases chosen after observing the lattice.

The area rule alone gives a lower bound of 1/8 on the maximum normalized prefix determinant. To strengthen it, suppose a basis has all its normalized prefix determinants at most $c < 1/4$. We replace its first block, of rank approximately $2cn$, by a reversed dual-HKZ basis, and reduce its final quotient, of the same rank, by HKZ. The middle prefix determinants remain unchanged. Deterministic reduction bounds place the modified profile below a fixed curve g_c, up to a vanishing error. This curve has a plateau of height c on $[2c, 1-2c]$ and curved flanks near the endpoints. Compared with the rectangle of height c, the two flanks remove a total area of c^2.

Applying the area rule with a small fixed slack therefore requires

$$\int_0^1 g_c(x)\,dx = c - c^2 \geq \frac{1}{8}, \qquad \text{hence} \qquad c \geq \frac{2-\sqrt{2}}{4} = c_{\text{flat}}. \tag{1.22}$$

The uniform comparison between prefix determinants and enumeration tree size then gives the lower bound for every basis (Section 7.2).

The conditional plateau construction. For the heuristic upper bound, we prescribe a coarse chain with 2 outer blocks and $K-2$ middle blocks. We verify that every nonempty subchain has a strictly positive first-moment exponent. Under the stated existence assumption, deterministic block reductions fill the gaps between prescribed prefixes, introducing peaks of height $O(1/K)$ above the critical plateau. Taking sufficiently many blocks therefore approaches the universal lower bound. For these chains, the first-moment exponent of every nonempty subchain can be written as a sum of squared gaps between the selected ranks and the plateau endpoints, making the required positive slack explicit (Lemma 6.1).

Collision search and affine targets. The FGL cost factor depends on the total variation of the normalized prefix profile. We combine blockwise HKZ product bounds with the two-block construction to control this factor and the sampling cost.

For the piecewise-linear normalized prefix profile p_n, the FGL collision search cost factor satisfies

$$\frac{\ln \sqrt{\mathcal{P}(\mathbf{B})}}{n \ln n} = \frac{1}{4}\,\mathrm{TV}(p_n) \geq \frac{1}{2}\max_x p_n(x). \tag{1.23}$$

The inequality uses $p_n(0) = p_n(1) = 0$ and transfers the enumeration lower bound to this cost factor. Repeated rises and falls can increase total variation even when the maximum height decreases, explaining why further splits need not improve collision search. For the upper bound, we use blockwise HKZ product estimates: convergence of prefix profiles alone does not control total variation.

For CVPP, we sample separately from $\mathcal{L}$ and $\boldsymbol{t} + \mathcal{L}$. Translation by a closest-vector offset gives a uniform overlap bound between these distributions. Packing each coset separately controls the number of points per grid cell, allowing the collision analysis to extend to arbitrary targets. Both samplers can be evaluated without knowing the closest-vector offset used in the analysis.

Quantum speedups for enumeration. For extending the results to quantum algorithms for CVPP, the shifted lattice-point bound also controls the enumeration tree uniformly over all targets. Quantum backtracking searches this tree at a cost proportional to the square root of its size. Its resulting exponent again depends on the maximum prefix height, a quantity we optimized with the repeated splits and the heuristic plateau profile. The small peaks introduced by the heuristic block construction disappear from this exponent as the number of blocks increases, even though their total variation may not disappear from the FGL cost factor.

Corrected Klein sampling. Klein's sequential sampler assigns a vector its Gaussian weight divided by a product of one-dimensional theta masses. The rejection correction of [11, full version, Section 5.2] removes this denominator. For centered sampling, Poisson summation and the zero vector give $\rho_s(\mathcal{L}) \geq \max\{1, s^n\}$, which bounds the expected number of proposals by $2^n\mathcal{P}(\mathbf{B})$ uniformly over all widths. Thus the two-block product estimate also improves the sampling exponent. A fixed cap on the number of proposals and polynomial numerical precision give statistical distance at most e^{-n^3} while preserving polynomial space.

2 Preliminaries

We first introduce the lattice notation and averaging formulas used in the analysis. All logarithms are natural. In the algorithmic bounds, we suppress polynomial factors in the bit length of the bases, targets, and sampling widths. The Haar model describes the geometric distribution of the lattices.

2.1 Lattices, projections and reduction

A full-rank lattice is $\mathcal{L} = \mathbf{B}\mathbb{Z}^n \subset \mathbb{R}^n$, with covolume $\det \mathcal{L} = |\det \mathbf{B}|$ and minimum $\lambda_1(\mathcal{L}) = \min_{\boldsymbol{v}\in\mathcal{L}\setminus\{\mathbf{0}\}}\|\boldsymbol{v}\|$. A lattice is *unimodular* if $\det \mathcal{L} = 1$. For a sublattice $\mathcal{M}$, $\det \mathcal{M}$ is its covolume in its span. It is *primitive* if a basis of $\mathcal{M}$ can be extended to a basis of $\mathcal{L}$. Equivalently, it is primitive if $\mathcal{M} = \mathcal{L} \cap \operatorname{span}_{\mathbb{R}} \mathcal{M}$. A vector $\boldsymbol{v}$ is primitive if $\mathbb{Z}\boldsymbol{v}$ is primitive. The quotient $\mathcal{X}_n = \mathrm{SL}_n(\mathbb{R})/\mathrm{SL}_n(\mathbb{Z})$ is the space of unimodular lattices, and μ_n is its Haar probability measure. When referring to a Haar-random lattice, we mean $\mathcal{L} \sim \mu_n$. The quantities studied in this paper are invariant under orthogonal transformations, and all choices of bases and shortest vectors are understood to be measurable.

For an ordered column basis $\mathbf{B} = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$, let $\boldsymbol{b}_i^*$ denote its Gram–Schmidt vectors, and write

$$\boldsymbol{b}_j = \boldsymbol{b}_j^* + \sum_{i<j} \mu_{j,i}\boldsymbol{b}_i^*, \qquad a_i = \|\boldsymbol{b}_i^*\|, \qquad y_i = \ln a_i. \tag{2.1}$$

For unimodular lattices $\mathcal{L}$, $\sum_i y_i = 0$. For a subspace W, write π_W for orthogonal projection onto W. For a fixed basis, we also write π_i for projection onto $\operatorname{span}(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_{i-1})^\perp$. The remaining dimension is then $d = n - i + 1$, and the projected lattice satisfies:

$$\mathcal{L}_{\mathbf{B},d} = \pi_i(\mathcal{L}), \qquad \det\left(\mathcal{L}_{\mathbf{B},d}\right) = \prod_{j=i}^{n} a_j. \tag{2.2}$$

Projection removes the components along the first $i-1$ basis vectors: the images of $\boldsymbol{b}_i, \ldots, \boldsymbol{b}_n$ generate the remaining rank-d lattice. We identify this orthogonal complement with $\mathbb{R}^d$. The normalized basis depth is $x_i = (i-1)/n = 1 - d/n$. The figures place the ith Gram–Schmidt coordinate at i/n; this shift of $1/n$ has no effect on the limiting curves. Throughout the paper we use the following notation for the log-determinant of the sublattice generated by $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_k$:

$$P_k(\mathbf{B}) = \sum_{i=1}^{k} \ln a_i \qquad (0 \le k \le n). \tag{2.3}$$

For lattices with $\det \mathcal{L} = 1$, this gives $P_0(\mathbf{B}) = P_n(\mathbf{B}) = 0$ and $\ln \det \mathcal{L}_{\mathbf{B},d} = -P_{n-d}(\mathbf{B})$. For unimodular lattices, $P_k(\mathbf{B})$ coincides with the Rankin profile studied by Li–Nguyen [41]. This also relates to Neumaier's bit profile [44], defined using leading Gram determinants, which has coordinates $2P_k(\mathbf{B})/\ln 2$.

A basis is called *size-reduced* if $|\mu_{j,i}| \le 1/2$ for all $i < j$. It is *HKZ-reduced* if it is size-reduced and $a_i = \lambda_1(\pi_i(\mathcal{L}))$ for every i. For $n \ge 2$, a basis is *quasi-HKZ-reduced* [28, Definition 2.4] if it is size-reduced, the tail $(\pi_2(\boldsymbol{b}_2), \ldots, \pi_2(\boldsymbol{b}_n))$ is HKZ-reduced, and

$$\frac{3}{4}a_1^2 \le a_2^2 + \mu_{2,1}^2 a_1^2. \tag{2.4}$$

This is the usual LLL condition [39, (1.4)–(1.5)] on the first pair $(\boldsymbol{b}_1, \boldsymbol{b}_2)$. In particular, $a_1 \le \sqrt{2}a_2$. Every HKZ basis is quasi-HKZ: comparing $\boldsymbol{b}_1$ with $\boldsymbol{b}_2$ gives the stronger inequality $a_1^2 \le a_2^2 + \mu_{2,1}^2 a_1^2$.

2.2 Geometric bounds and averaging

Let $\mathbb{B}_d(R)$ denote the Euclidean ball of radius R, and let Γ and ζ denote the gamma and Riemann zeta functions. We will use the following constants:

$$V_d = \frac{\pi^{d/2}}{\Gamma(1+d/2)}, \qquad r_d = V_d^{-1/d} \quad (d \geq 1), \qquad \kappa_d = \frac{V_d}{2\zeta(d)} \quad (d \geq 2). \tag{2.5}$$

Here V_d is the volume of the unit ball in d dimensions, and r_d is the radius of a ball of volume one.

Lemma 2.1 (Geometric bounds). *For every rank-d lattice $\mathcal{L}$, the following inequalities hold, with the asymptotic estimate taken as $d \to \infty$:*

$$\frac{\lambda_1(\mathcal{L})}{(\det \mathcal{L})^{1/d}} \leq 2r_d \leq \sqrt{d}, \qquad \ln r_d = \frac{1}{2}\ln\frac{d}{2\pi e} + o(1). \tag{2.6}$$

Proof. Minkowski's convex body theorem [45, Theorem 4] gives the first inequality; the cube of side $2/\sqrt{d}$ inside the unit ball gives the second. Stirling's formula gives the asymptotic and $\ln V_d = -d\ln d/2 + O(d)$. □

For Haar-random unimodular lattices, the sharper estimate $\lambda_1(\mathcal{L}) = (1+o(1))r_d$ holds with probability $1-o(1)$ [52, Theorem 1]. We will also need the deterministic bound: the projections arising later in the analysis depend on the chosen basis, and need not be Haar-random. A standard consequence of Siegel's primitive-vector mean [51] is the explicit lower-tail estimate

$$\mathbb{P}_{\mu_d}\{\lambda_1(\mathcal{L}) < r_d/2\} \leq \kappa_d (r_d/2)^d = \frac{2^{-d}}{2\zeta(d)}. \tag{2.7}$$

Indeed, the expected number of primitive pairs $\{\pm\boldsymbol{v}\}$ in that ball is its volume divided by $2\zeta(d)$, and every shortest vector is primitive.

For a primitive pair $[\boldsymbol{v}] = \{\pm\boldsymbol{v}\}$ in a unimodular rank-d lattice, we obtain a smaller unimodular lattice by projecting away the line spanned by $\boldsymbol{v}$ and rescaling the result to determinant one:

$$\mathcal{Q}_{\boldsymbol{v}} = \|\boldsymbol{v}\|^{1/(d-1)}\pi_{\boldsymbol{v}^\perp}(\mathcal{L}) \in \mathcal{X}_{d-1}. \tag{2.8}$$

The projected lattice consists of the components of lattice vectors perpendicular to $\boldsymbol{v}$. Its determinant is $1/\|\boldsymbol{v}\|$, which explains the scaling factor. This is exactly the smaller lattice encountered after one HKZ step, normalized to unit determinant.

Proposition 2.2 (Averaging over primitive vectors). *Let $d \geq 2$ and let $f(s, \mathcal{Q})$ be nonnegative, measurable and rotation-invariant in $\mathcal{Q}$. Then*

$$\mathbb{E}_{\mu_d}\left[\sum_{[\boldsymbol{v}]\in\mathcal{L}_{\mathrm{prim}}/\{\pm1\}} f(\|\boldsymbol{v}\|, \mathcal{Q}_{\boldsymbol{v}})\right] = d\kappa_d \int_0^\infty s^{d-1}\mathbb{E}_{\mu_{d-1}}[f(s,\mathcal{Q})]\,\mathrm{d}s. \tag{2.9}$$

Here $\mathcal{X}_1 = \{\mathbb{Z}\}$ and μ_1 is its point mass.

Proof. We apply the primitive component ($k=1$) of Strömbergsson's averaging formula [53, Proposition 5.1 and (5.1)–(5.4)]. For a specified primitive vector of length s, the normalized quotient carries Haar measure μ_{d-1}, while the choices of lifts have total measure one. Counting $\boldsymbol{v}$ and $-\boldsymbol{v}$ separately gives primitive-vector intensity $\zeta(d)^{-1}\,\mathrm{d}\boldsymbol{v}$. Passing to polar coordinates and dividing by two to count primitive pairs gives the factor $d\kappa_d$. □

In Section 4.2, we use this with $f(s,\mathcal{Q}) = \mathbf{1}_{s\leq 2r_d}(2r_d/s)^{d/2}\exp(Z_{d-1}(\mathcal{Q})/2)$, where Z_{d-1} measures cumulative HKZ length deficits. We also use the same quotient weight with the simpler cutoff $\mathbf{1}_{s\leq 4r_d}$ for quasi-HKZ bases. Summing over all eligible first vectors will allow us to obtain a bound that holds for every possible first HKZ step.

2.3 Duality, Gaussian mass, and sublattice averages

The dual of a lattice $\mathcal{L}$ in its span is $\mathcal{L}^* = \{\boldsymbol{y} : \langle \boldsymbol{y}, \mathcal{L} \rangle \subset \mathbb{Z}\}$. Haar measure is invariant under duality. The Gaussian mass and its basic identities are

$$\rho(\mathcal{L}) = \sum_{\boldsymbol{v} \in \mathcal{L}} e^{-\pi \|\boldsymbol{v}\|^2}, \qquad \rho(\mathcal{L}) = \frac{\rho(\mathcal{L}^*)}{\det \mathcal{L}}, \qquad \mathbb{E}_{\mu_n}[\rho(\mathcal{L})] = 2 \quad (n \geq 2). \tag{2.10}$$

The middle identity is Poisson summation; the last is Siegel's mean-value formula [51], including the zero vector.

Following Kim, write $\mathrm{Gr}(\mathcal{L}, r)$ for the primitive rank-r sublattices, counted without orientations or choices of bases. The notation comes from the Grassmannian: each such sublattice is determined by its span through $\mathcal{M} = \mathcal{L} \cap \operatorname{span}_{\mathbb{R}} \mathcal{M}$. The coefficient $a(n, r)$ is the expected number of such sublattices with determinant at most one:

$$a(n, r) = \frac{1}{n} \binom{n}{r} \prod_{j=1}^{r} \frac{V_{n-j+1} \zeta(j)}{V_j \zeta(n-j+1)}, \qquad \zeta(1) := 1. \tag{2.11}$$

Kim's mean-value formula [35, Theorem 3, $k = 1$] states that

$$N_{r,H}(\mathcal{L}) = \#\{\mathcal{M} \in \mathrm{Gr}(\mathcal{L}, r) : \det \mathcal{M} \leq H\}, \qquad \mathbb{E}_{\mu_n}[N_{r,H}(\mathcal{L})] = a(n, r) H^n. \tag{2.12}$$

The mean-value formula for a fixed rank also follows from the classical averaging framework of Thunder [54]. The $k = 2$ case of the same theorem gives the mean number of ordered pairs with ranks r, t, determinant bounds H, J, and zero intersection: it equals $a(n, r)a(n, t)H^n J^n$ when $r + t \leq n - 1$. This identity concerns the average number of such pairs. In particular, it does not imply that sublattices chosen by an algorithm are independent.

The next lemma also keeps track of the two smaller lattices obtained by splitting along $\mathcal{M}$: the sublattice itself and its orthogonal quotient. After normalizing their determinants, their averages separate into two Haar integrals.

Lemma 2.3 (Averaging the two blocks of a split). *For $\mathcal{M} \in \mathrm{Gr}(\mathcal{L}, r)$, let $D = \det \mathcal{M}$ and $\mathcal{Q} = \pi_{\mathcal{M}^\perp}(\mathcal{L})$, so that $\det \mathcal{Q} = D^{-1}$. For every nonnegative measurable function F that is rotation-invariant in both lattice arguments:*

$$\mathbb{E}_{\mu_n}\left[\sum_{\mathcal{M} \in \mathrm{Gr}(\mathcal{L}, r)} F(D, D^{-1/r}\mathcal{M}, D^{1/(n-r)}\mathcal{Q})\right] = n a(n, r) \int_0^\infty D^{n-1} \mathbb{E}_{\substack{\mathcal{L}_1 \sim \mu_r \\ \mathcal{L}_2 \sim \mu_{n-r}}} [F(D, \mathcal{L}_1, \mathcal{L}_2)] \, \mathrm{d}D, \tag{2.13}$$

where $\mathcal{L}_1$ and $\mathcal{L}_2$ on the right are independent.

Proof. Kim's decomposition of Haar measure [35, Lemma 10 and Sections 4–5] separates the determinant, the two normalized lattices, rotations, and choices of lifts. The determinant density is proportional to $D^{n-1} \, \mathrm{d}D$, and the two lattice factors carry Haar probability measure. Integrating out the rotations and lifts leaves a constant; taking $F = \mathbf{1}_{D \leq H}$ and using (2.12) identifies it as $n a(n, r)$. □

For example, if the two blocks must belong to prescribed sets of Haar probabilities p_1, p_2, their expected count up to determinant H is $a(n, r)H^n p_1 p_2$. This is an average over all candidate splits. A sublattice selected by an algorithm need not itself have a Haar distribution. Section 5.1 uses the identity to exclude candidates with unsuitable blocks before choosing a split.

3 Enumeration trees and prefix determinants

To compare different bases, we first relate their enumeration tree sizes to their prefix determinants. We show that, for random lattices, the largest prefix determinant determines the tree size up to a factor $e^{O(n)}$. This estimate holds simultaneously for all bases, and will therefore be used for both the upper and lower bounds. The connection between prefix determinants and enumeration costs also appears in the statistical study of Yu–Ducas [56, Section 5], who analyzed the half-volume of BKZ-reduced bases as an indicator of enumeration costs. Here we use the prefix-determinant profile to control the full enumeration tree size.

3.1 Levels, projections, and prefixes

At depth d, full enumeration has chosen the coefficients of the last d basis vectors $\boldsymbol{b}_n, \boldsymbol{b}_{n-1}, \ldots, \boldsymbol{b}_{n-d+1}$. The root corresponds to $d = 0$, the leaves correspond to $d = n$, and the complementary prefix has rank $k = n - d$. The number of nodes at depth d and the total tree size are given by:

$$N_d(\mathbf{B}, R) = \#\big(\mathcal{L}_{\mathbf{B},d} \cap \mathbb{B}_d(R)\big), \qquad T(\mathbf{B}, R) = 1 + \sum_{d=1}^{n} N_d(\mathbf{B}, R). \tag{3.1}$$

Each projected vector corresponds to one integer coefficient tuple. The root contributes one; sign symmetry changes the count by at most a constant factor. The fixed-radius full tree is independent of traversal order, with polynomial arithmetic overhead per node [20,48].

For $\mathcal{M}_k = \mathbb{Z}\boldsymbol{b}_1 + \cdots + \mathbb{Z}\boldsymbol{b}_k$ and $\mathcal{Q}_k = \pi_{\mathcal{M}_k^\perp}(\mathcal{L})$, primitivity of basis prefixes gives

$$\det \mathcal{M}_k = e^{P_k(\mathbf{B})}, \qquad \det \mathcal{Q}_k = e^{-P_k(\mathbf{B})}, \qquad \mathcal{Q}_k^* = \mathcal{L}^* \cap \mathcal{M}_k^\perp. \tag{3.2}$$

A large prefix determinant therefore makes the corresponding projected lattice dense, which suggests that the enumeration tree will be wide at this depth. We make this relation precise below. In Section 7, we will also use that the prefixes form a nested chain of primitive sublattices.

3.2 Uniform lower and upper estimates

Proposition 3.1 (Tree size from the prefix profile). *With probability $1 - e^{-\Omega(n)}$ over $\mathcal{L} \sim \mu_n$, simultaneously for all bases at radius $R = \lambda_1(\mathcal{L})$,*

$$\ln T(\mathbf{B}) = \max_{0 \le k \le n} P_k(\mathbf{B}) + O(n), \qquad \ln N_{n-k}(\mathbf{B}, R) = P_k(\mathbf{B}) + O(n) \quad (k < n). \tag{3.3}$$

Proof. We first prove the lower bound, with $d = n - k$. Averaging over a fundamental domain of $\mathcal{Q}_k$ shows that some translate of $\mathbb{B}_d(R/2)$ contains at least $V_d(R/2)^d / \det \mathcal{Q}_k$ lattice points. Subtracting one of these points from all of them gives distinct points in $\mathcal{Q}_k \cap \mathbb{B}_d(R)$, because distances within a ball of radius $R/2$ are at most R.

For the upper bound, every point with $\|\boldsymbol{v}\| \le R$ contributes at least one to $e^{\pi R^2}\rho(\mathcal{Q}_k)$. Equivalently, the Gaussian function $e^{\pi(R^2 - \|\boldsymbol{v}\|^2)}$ bounds the ball's indicator from above. Applying Poisson summation and then $\mathcal{Q}_k^* \subseteq \mathcal{L}^*$ gives

$$e^{P_k} V_d(R/2)^d \le N_d(\mathbf{B}, R) \le e^{\pi R^2}\rho(\mathcal{Q}_k) = e^{\pi R^2 + P_k}\rho(\mathcal{Q}_k^*) \le e^{\pi R^2 + P_k}\rho(\mathcal{L}^*). \tag{3.4}$$

By (2.7) and Stirling's formula, with probability $1 - e^{-\Omega(n)}$ we have $R \ge a\sqrt{n}$ for an absolute $a > 0$. Markov's inequality and (2.10) also give $\rho(\mathcal{L}^*) \le e^n$ with probability at least $1 - 2e^{-n}$. Both bounds

therefore hold with probability $1 - e^{-\Omega(n)}$. Using $R \leq \sqrt{n}$ from Lemma 2.1 and the cube inside the unit ball,

$$V_d(R/2)^d \geq a^d (n/d)^{d/2} \geq e^{-O(n)}, \qquad e^{\pi R^2} \rho(\mathcal{L}^*) \leq e^{(\pi+1)n}. \tag{3.5}$$

Taking logarithms gives the claimed estimate for each level of the tree. As there are only $n+1$ levels, the logarithm of the total tree size differs from that of the widest level by at most $\ln(n+1)$. This proves the proposition. □

Shifted enumeration and CVPP. For a target $\boldsymbol{t}$, write $\boldsymbol{t}_k = \pi_{\mathcal{M}_k^\perp}(\boldsymbol{t})$. In Poisson summation a shift only introduces complex factors of absolute value one, so the same upper bound holds uniformly over all targets:

$$\#\big((\mathcal{Q}_k - \boldsymbol{t}_k) \cap \mathbb{B}_{n-k}(R)\big) \leq e^{\pi R^2 + P_k(\mathbf{B})} \rho(\mathcal{L}^*). \tag{3.6}$$

The covering-radius estimate and the resulting classical CVPP bounds are given in Section 8.1; Sections 8.2 and 8.3 give collision-search and quantum query bounds.

In the rest of the paper, we will express prefix determinants through limiting profiles. If $P_{\lfloor xn \rfloor}/(n \ln n) = p(x) + o(1)$ uniformly in x, the enumeration tree coefficient is $\max_x p(x)$. The figures use prefix dimension $x = k/n$, so enumeration runs from the root at $x = 1$ to the leaves at $x = 0$.

4 HKZ profiles on random lattices

We next analyze the enumeration trees of quasi-HKZ bases. We first derive a deterministic bound on HKZ prefix determinants, which will also be useful for the constructions in Section 5. To show that this bound is tight on random lattices, we then bound how far the successive HKZ lengths can fall below the Gaussian heuristic. This gives the limiting prefix profile and the enumeration exponent $1/(2e)$.

4.1 HKZ prefix determinants

Consider an HKZ basis $\mathbf{B}$ of a unimodular rank-m lattice. At remaining dimension j, HKZ chooses a shortest vector in $\mathcal{L}_{\mathbf{B},j}$; its length is the Gram–Schmidt length a_{m-j+1}. We divide by $(\det \mathcal{L}_{\mathbf{B},j})^{1/j}$ to compare different scales, and let h_j be the logarithm of this normalized length:

$$h_j = \ln \frac{\lambda_1(\mathcal{L}_{\mathbf{B},j})}{(\det \mathcal{L}_{\mathbf{B},j})^{1/j}} = y_{m-j+1} + \frac{P_{m-j}(\mathbf{B})}{j}. \tag{4.1}$$

The recursive Gaussian heuristic (GH) predicts $h_j \approx \ln r_j$. Substituting this prediction at successive ranks yields the heuristic HKZ Gram–Schmidt profile. The sandpile model of Hanrot–Pujol–Stehlé [25, Section 4.1] already yields the same leading profile; Haque–Pieprzyk [29, Sections 3.2 and 4] analyze a Gaussian-heuristic variant and derive more precise asymptotic estimates. Our aim is to justify the leading profile for random lattices despite the dependence between successive HKZ choices. Starting from $P_0 = 0$, the relation $P_{m-j+1}/(j-1) = P_{m-j}/j + h_j/(j-1)$ for $j \geq 2$ gives the exact identity

$$P_{m-d}(\mathbf{B}) = d \sum_{j=d+1}^{m} \frac{h_j}{j-1}, \qquad 1 \leq d \leq m. \tag{4.2}$$

Lemma 4.1 (Deterministic HKZ prefix bound). *For every HKZ basis* $\mathbf{B}_{\mathrm{H}}$ *of a unimodular rank-d lattice, uniformly for* $0 \le t \le 1$,

$$\frac{P_{\lfloor td \rfloor}(\mathbf{B}_{\mathrm{H}})}{d \ln d} \le p_{\mathrm{H}}(t) + o(1), \qquad p_{\mathrm{H}}(t) = -\frac{1-t}{2}\ln(1-t). \tag{4.3}$$

Proof. Minkowski's bound gives $h_j \le \frac{1}{2}\ln j$. For $r = d-k \ge 1$, (4.2) implies

$$P_{d-r} \le \frac{r}{2}\ln d \sum_{j=r+1}^{d} \frac{1}{j-1} \le \frac{r}{2}\ln d\left(\ln(d/r) + \frac{1}{r}\right). \tag{4.4}$$

After dividing by $d\ln d$, the rounding error is uniform in k. At the remaining endpoint $k = d$, the claim follows from $P_d = 0$. □

In the split constructions, we will also use dual HKZ bases. Reversing and dualizing an HKZ basis of the dual lattice changes the prefix determinant P_k to P_{d-k}. The corresponding upper bound is therefore the reflected profile:

$$p_{\mathrm{D}}(t) = p_{\mathrm{H}}(1-t) = -\frac{t}{2}\ln t. \tag{4.5}$$

4.2 A one-sided moment bound

To compare the actual HKZ lengths with the recursive GH prediction, we measure their shortfall from $\ln(2r_j)$. We use this slightly larger reference value so that the deficits $e_j = \ln(2r_j) - h_j$ are nonnegative for every lattice. The additional $\ln 2$ does not change the leading profile. For a unimodular rank-m lattice $\mathcal{L}$, let $Z_m(\mathcal{L})$ be the maximum of $\sum_{j=2}^{m} j e_j$ over its HKZ bases, and put $Z_1 = 0$. Thus Z_m measures the total shortfall of the normalized logarithmic GS lengths from the recursive GH reference, weighted by the remaining dimension. The maximum is finite because each HKZ step has finitely many shortest-vector choices.

Lemma 4.2 (A moment bound for the cumulative deficit). *For* $\mathcal{L} \sim \mu_n$,

$$\mathbb{E}_{\mu_n}[e^{Z_n(\mathcal{L})/2}] = e^{O(n^2)}. \tag{4.6}$$

Proof. We prove the claim by induction on the dimension, writing $M_d = \mathbb{E}_{\mu_d}[e^{Z_d/2}]$ and starting from $M_1 = 1$. For a fixed unimodular rank-d lattice, let $\mathcal{S}(\mathcal{L})$ be its finite nonempty set of shortest-vector pairs. Once a first vector is chosen, all later HKZ choices are choices in its projected quotient. Multiplying every lattice vector by the same factor changes its length and $(\det\mathcal{L})^{1/j}$ by that factor, so h_j and e_j stay unchanged. We may therefore rescale the quotient to determinant one without changing the remaining deficits:

$$Z_d(\mathcal{L}) = \max_{[\boldsymbol{v}]\in\mathcal{S}(\mathcal{L})}\left(d\ln\frac{2r_d}{\|\boldsymbol{v}\|} + Z_{d-1}(\mathcal{Q}_{\boldsymbol{v}})\right). \tag{4.7}$$

Every shortest vector is primitive and has length at most $2r_d$. After exponentiating, we can bound the maximum by the sum over all shortest-vector pairs, and then enlarge this sum to all primitive pairs in the ball of radius $2r_d$. All summands are nonnegative, so sums and integrals may be interchanged without first proving integrability (Tonelli's theorem [21, Theorem 2.37]). Proposition 2.2 then gives:

$$M_d \le d\kappa_d M_{d-1}\int_0^{2r_d} s^{d-1}\left(\frac{2r_d}{s}\right)^{d/2}\mathrm{d}s = \frac{2^d}{\zeta(d)}M_{d-1} \le 2^d M_{d-1}. \tag{4.8}$$

Here $V_d r_d^d = 1$. The factor $1/2$ in the moment ensures convergence at $s = 0$; removing it would give the divergent integral $\int_0^{2r_d} s^{-1}\,\mathrm{d}s$. Iteration gives $\ln M_n \le (\ln 2)\sum_{d=2}^{n} d = O(n^2)$. Because we average over all primitive vectors, the argument does not require the quotients selected by HKZ to be Haar-random or independent. □

4.3 Quasi-HKZ profiles

Every quasi-HKZ basis satisfies $\lambda_1(\mathcal{L}) \le a_1 \le \sqrt{2}\lambda_1(\mathcal{L}) < 4r_n$. Indeed, if the first vector is not shortest, a shortest vector has nonzero projection onto the HKZ tail, so $a_2 \le \lambda_1(\mathcal{L})$ and (2.4) gives the claim. The last inequality follows from Minkowski's bound.

We now use the moment bound to justify the leading HKZ profile predicted by the heuristic models of Hanrot–Pujol–Stehlé and Haque–Pieprzyk [25,29]. The next proposition shows that the actual logarithmic Gram–Schmidt lengths y_i differ from $\frac{1}{2}\ln n\,(1+\ln(1-x_i))$ by $o(\ln n)$ on average, away from the final ranks. Summing these coordinates gives the uniform prefix profile p_{H}, whose maximum determines the tree exponent.

Proposition 4.3 (Quasi-HKZ profiles on Haar-random lattices). *With probability $1-e^{-\Omega(n)}$ over $\mathcal{L} \sim \mu_n$, simultaneously for all quasi-HKZ bases,*

$$\max_{0\le k\le n}\left|\frac{P_k(\mathbf{B})}{n\ln n} - p_{\mathrm{H}}(k/n)\right| = o(1). \tag{4.9}$$

With the same probability, for every fixed $0 < \varepsilon < 1$,

$$\frac{1}{n}\sum_{i:\,x_i\le 1-\varepsilon}\left|\frac{y_i}{\ln n} - \frac{1}{2}\big(1+\ln(1-x_i)\big)\right| = o(1). \tag{4.10}$$

Proof. Let $m = n-1$, and let W_n be the maximum of $Z_m(\mathcal{Q}_{\boldsymbol{v}})$ over primitive pairs with $\|\boldsymbol{v}\| \le 4r_n$. Proposition 2.2 and Lemma 4.2 give

$$\mathbb{E}_{\mu_n}[e^{W_n/2}] \le \kappa_n(4r_n)^n\mathbb{E}_{\mu_m}[e^{Z_m/2}] = e^{O(n^2)}. \tag{4.11}$$

Hence $W_n = O(n^2)$ with probability $1-e^{-\Omega(n^2)}$. Together with (2.7) and Proposition 3.1, this gives an event of probability $1-e^{-\Omega(n)}$ on which all three bounds hold. In particular, every quasi-HKZ basis then has $|\ln a_1| = O(\ln n)$.

For its normalized HKZ tail $\mathbf{B}'$, put $g_j = \ln(2r_j)$ for $j \ge 2$, with $g_1 = e_1 = 0$. The prefix identity gives

$$y'_{m-d+1} = h_d - \sum_{j=d+1}^{m}\frac{h_j}{j-1}, \qquad R_{m,d} = g_d - \sum_{j=d+1}^{m}\frac{g_j}{j-1}. \tag{4.12}$$

Using $h_j = g_j - e_j$ and the triangle inequality,

$$\sum_{d=1}^{m} d|y'_{m-d+1} - R_{m,d}| \le \sum_{j=2}^{m}\left(j + \sum_{d<j}\frac{d}{j-1}\right)e_j \le \frac{3}{2}W_n = O(n^2). \tag{4.13}$$

For fixed $\varepsilon > 0$ and $d \ge \varepsilon m$, Stirling's formula and harmonic sums give $R_{m,d} = \frac{1}{2}\ln m\,(1+\ln(d/m)) + O(1)$, where the constant may depend on ε. The weighted bound above therefore gives an average

coordinate error $O(1)$ in this range. Rescaling by $y_i = y'_{i-1} - \ln a_1/m$ for $i \geq 2$ and dividing by $\ln n$ proves (4.10).

Summing the coordinate errors and using $\frac{1}{2}\int_0^x (1+\ln(1-t))\,\mathrm{d}t = p_{\mathrm{H}}(x)$ proves (4.9) for $k/n \leq 1-\varepsilon$.

For the remaining prefixes, Lemma 4.1 applied to the tail gives $P_k/(n\ln n) \leq p_{\mathrm{H}}(k/n) + o(1)$, while Proposition 3.1 and $N_{n-k} \geq 1$ give $P_k \geq -O(n)$. Since $P_n = 0$ and $p_{\mathrm{H}}(x) \to 0$ as $x \to 1$, letting $\varepsilon \downarrow 0$ proves the uniform claim. □

4.4 Matching enumeration bounds

Proof (Theorem 1.1). With probability $1 - e^{-\Omega(n)}$, Propositions 3.1 and 4.3 both hold, giving uniformly over all quasi-HKZ bases:

$$\frac{\ln T(\mathbf{B})}{n\ln n} = \max_{0\leq x\leq 1} p_{\mathrm{H}}(x) + o(1) = \frac{1}{2e} + o(1). \tag{4.14}$$

The maximum occurs at $x = 1 - 1/e$. □

5 Splitting bases with dense sublattices

To improve on HKZ, we construct a basis by first finding a dense primitive sublattice and then reducing this sublattice and its quotient separately. We show that, with high probability, a suitable split of a Haar random lattice exists and that the same approach can be applied recursively. For the two-block construction, Proposition 4.3 gives the exact leading tree exponent. For the recursive construction, the deterministic upper bound from Lemma 4.1 suffices.

5.1 Dense sublattices with suitable blocks

A primitive rank-m sublattice $\mathcal{M}$ splits $\mathcal{L}$ into $\mathcal{M}$ and the rank-$(n-m)$ quotient $\mathcal{Q} = \pi_{\mathcal{M}^\perp}(\mathcal{L})$. Their determinants are reciprocal. We need a split with small $\det\mathcal{M}$ whose two normalized blocks also have the usual high-probability properties of random lattices, such as the HKZ profile or the existence of further splits.

A lower bound at the relevant determinant scale follows from Shapira–Weiss [50, Proposition 3]: for some absolute constant $C > 0$, with probability $1-o(1)$, every rank-k sublattice $\mathcal{M} \subseteq \mathcal{L}$ satisfies

$$\det\mathcal{M} \geq (n/C)^{k(n-k)/(2n)} \qquad (1 \leq k < n). \tag{5.1}$$

In particular, at half rank this gives the lower scale $e^{-O(n)}n^{n/8}$. The following lemma provides a matching upper scale, up to $e^{O(n)}$ factors, while also ensuring suitable geometry of both normalized blocks.

Lemma 5.1 (Dense sublattices with suitable blocks). *Fix $0 < \alpha < 1$. Take $m = \lfloor \alpha n \rfloor$ for $\alpha \neq 1/2$, and $m = \lfloor (n-1)/2 \rfloor$ for $\alpha = 1/2$. There is a threshold*

$$H = e^{O(n)} n^{m(n-m)/(2n)} \tag{5.2}$$

with the following property. Prescribe any rotation-invariant measurable properties that hold with Haar probability $1-o(1)$ in dimensions m and $n-m$. With probability $1-o(1)$, $\mathcal{L} \sim \mu_n$ contains a primitive rank-m sublattice $\mathcal{M}$ with $H/2 \leq D = \det\mathcal{M} \leq H$ such that $D^{-1/m}\mathcal{M}$ and $D^{1/(n-m)}\mathcal{Q}$ satisfy the respective properties. Either property may instead concern the dual of that normalized block.

To prove the lemma, we count candidate splits and remove those containing an unusually dense proper sublattice. This removes only a small fraction of the candidates in expectation, while allowing us to control the correlations between intersecting candidates. The next two lemmas establish the required first- and second-moment bounds. We then use Lemma 2.3 to remove candidates with unsuitable blocks, and show that at least one suitable split remains with high probability.

We first consider the case $2m \leq n-1$ and $m/n \to \alpha > 0$. Let $N_H = N_{m,H}$ be the number of candidates and let $\mu = \mathbb{E}[N_H] = a(n,m)H^n$. We retain a candidate $\mathcal{M}$ only if every proper primitive sublattice $\mathcal{C} \subset \mathcal{M}$ satisfies:

$$\det \mathcal{C} \geq U_r := n^{r(n-r)/(2n)}, \qquad r = \operatorname{rank} \mathcal{C} \in \{1, \dots, m-1\}. \tag{5.3}$$

Let Z be the number of retained candidates. In the next two lemmas, we choose the positive constant implicit in the factor $e^{O(n)}$ in H sufficiently large.

Lemma 5.2 (Few candidates are removed). *The reduced count Z satisfies $\mathbb{E}[Z] = (1-e^{-\Omega(n)})\mu$.*

Proof. For a candidate of determinant D, rescaling to determinant one changes the cutoff to $U_r D^{-r/m}$. By Lemma 2.3, the mean number of violations at rank r in this normalized lattice is $a(m,r)U_r^m D^{-r}$. Summing over the ranks and integrating over D, we obtain:

$$\frac{\mathbb{E}[N_H - Z]}{\mu} \leq \frac{na(n,m)}{\mu} \sum_{r=1}^{m-1} a(m,r) U_r^m \int_0^H D^{n-r-1} \,\mathrm{d}D = \sum_{r=1}^{m-1} \frac{n}{n-r} a(m,r) U_r^m H^{-r}. \tag{5.4}$$

To estimate these terms uniformly in r, we write $b_j = \ln(V_j/\zeta(j))$, with $\zeta(1) = 1$. Stirling's formula gives $b_j = -j\ln n/2 + O(n)$ uniformly for $1 \leq j \leq n$, since $j\ln(n/j) = O(n)$. The product for $a(m,r)$ has r numerator and r denominator terms, so

$$\ln a(m,r) = -\frac{r(m-r)}{2} \ln n + O(rn). \tag{5.5}$$

Consequently, the logarithm of the rth summand in (5.4) is

$$-\frac{r(m-r)(n-m)}{2n} \ln n - r \ln \frac{H}{n^{m(n-m)/(2n)}} + O(rn). \tag{5.6}$$

Choosing the constant in H sufficiently large makes this at most $-\kappa nr$ for some $\kappa > 0$. Summing over r proves the claim. □

Lemma 5.3 (The reduced count has small variance). *For the same choice of H, $\mu \geq e^{\Omega(n^2)}$ and $\mathbb{E}[Z^2] \leq (1 + e^{-\Omega(n)})\mu^2$.*

Proof. We distinguish pairs according to the rank of their intersection, and write $s = n - 2m \geq 1$. Pairs with zero intersection have mean μ^2 before removing candidates, by Kim's pair formula in Section 2.3. The diagonal contributes at most μ. It remains to count distinct candidates intersecting in a rank-r lattice $\mathcal{C}$, where $1 \leq r < m$.

The intersection is primitive and has $D := \det \mathcal{C} \geq U_r$. Projecting away $\mathcal{C}$ gives two primitive rank-$(m-r)$ sublattices with zero intersection in a rank-$(n-r)$ quotient of determinant D^{-1}. Each projected sublattice determines its original candidate as its full inverse image, so there is no additional multiplicity. Normalizing the quotient changes the two determinant ceilings to $HD^{-(n-m)/(n-r)}$.

Kim's pair formula applies because $2(m-r) \le n-r-1$. Integrating its value against the intersection density $na(n,r)D^{n-1}\,\mathrm{d}D$ bounds the number of these pairs by

$$na(n,r)a(n-r,m-r)^2H^{2(n-r)}\int_{U_r}^{\infty} D^{2m-n-1}\,\mathrm{d}D = \frac{n}{s}a(n,r)a(n-r,m-r)^2H^{2(n-r)}U_r^{-s}. \quad (5.7)$$

The cutoff removes the divergence at zero, while $s \ge 1$ ensures convergence at infinity. This is why, for $\alpha = 1/2$, we chose $m = \lfloor (n-1)/2 \rfloor$ rather than $\lfloor n/2 \rfloor$.

Combining the diagonal, the pairs with zero intersection, and the remaining intersecting pairs gives:

$$\frac{\mathbb{E}[Z^2]}{\mu^2} \le 1 + \frac{1}{\mu} + \frac{n}{s}\sum_{r=1}^{m-1}\frac{a(n,r)a(n-r,m-r)^2}{a(n,m)^2}H^{-2r}U_r^{-s}. \quad (5.8)$$

Using the same b_j as above, cancellation in the products for $a(n,r)$ gives

$$\ln\frac{a(n,r)a(n-r,m-r)^2}{a(n,m)^2} = 2\sum_{j=m-r+1}^{m} b_j - \sum_{j=n-r+1}^{n} b_j - \sum_{j=1}^{r} b_j + O(r\ln n) \quad (5.9)$$

$$= \frac{r(s+r)}{2}\ln n + O(rn). \quad (5.10)$$

The logarithm of the rth contribution to (5.8) is therefore

$$-\frac{r(m-r)(n-m)}{n}\ln n - 2r\ln\frac{H}{n^{m(n-m)/(2n)}} + O(rn). \quad (5.11)$$

The error term also absorbs the factor $n/s \le n$. As in Lemma 5.2, increasing the constant in H makes the sum $e^{-\Omega(n)}$. Finally, $\ln a(n,m) = -m(n-m)\ln n/2 + O(n^2)$, so the same choice makes $\ln\mu = \Omega(n^2)$. □

Proof (Lemma 5.1). We first assume $2m \le n-1$. Restricting to candidates with determinant at least $H/2$ removes at most $2^{-n}\mu$ candidates in expectation. If the two prescribed properties have failure probabilities ε_m and ε_{n-m}, Lemma 2.3 bounds the expected number of candidates with an unsuitable block by $(\varepsilon_m + \varepsilon_{n-m})\mu = o(\mu)$. These estimates also apply after imposing the internal cutoff, since doing so can only decrease the number of unsuitable candidates. Let Y count the candidates satisfying the cutoff, the determinant window and both block properties. The two preceding lemmas give

$$\mathbb{E}[Y] = (1-o(1))\mu, \qquad \mathbb{P}\{Y>0\} \ge \frac{\mathbb{E}[Y]^2}{\mathbb{E}[Y^2]} \ge \frac{\mathbb{E}[Y]^2}{\mathbb{E}[Z^2]} = 1-o(1). \quad (5.12)$$

Here the probability estimate follows by applying Cauchy–Schwarz to $Y\mathbf{1}_{Y>0}$.

To cover the case $\alpha > 1/2$, we apply the result at rank $n-m$ in the dual lattice $\mathcal{L}^*$. If $\mathcal{N} \subset \mathcal{L}^*$ is the resulting primitive sublattice, take $\mathcal{M} = \mathcal{L} \cap \mathcal{N}^\perp$. Then $\det\mathcal{M} = \det\mathcal{N}$, while $\pi_{\mathcal{M}^\perp}(\mathcal{L}) = \mathcal{N}^*$ and $\pi_{\mathcal{N}^\perp}(\mathcal{L}^*) = \mathcal{M}^*$. The two normalized blocks are therefore swapped and dualized. Haar invariance under duality transfers the prescribed properties, proving the remaining rank fractions and the dual variants. □

The resulting determinant exponent is optimal to leading order. Indeed, for every fixed $\varepsilon > 0$, the first-moment bound gives:

$$\mathbb{P}\left\{\exists\mathcal{M} \in \mathrm{Gr}(\mathcal{L},m) : \det\mathcal{M} \le n^{m(n-m)/(2n)-\varepsilon n}\right\} \le n^{-\varepsilon n^2 + O(n^2/\ln n)} = o(1). \quad (5.13)$$

The exponent $m(n-m)/(2n)$ is maximized at $m = n/2$, where it is $n/8$. Thus the half-rank construction has the scale $n^{n/8+o(n)}$ used in the introduction.

5.2 The explicit two-block split

We now apply the lemma to a sublattice of approximately half the ambient rank, fixing the middle prefix determinant at $n^{n/8+o(n)}$. We complete the left block with a reversed dual-HKZ basis and the right block with an HKZ basis. Requiring the normalized blocks to have the usual Haar profiles allows us to determine the entire limiting prefix profile. Its two maxima occur inside the blocks, on either side of the split.

Proof (Theorem 1.2). Set $m = \lfloor (n-1)/2 \rfloor$. Lemma 5.1 supplies a primitive rank-m sublattice $\mathcal{M}$ with $H/2 \le D = \det \mathcal{M} \le H$, where $H = e^{O(n)} n^{m(n-m)/(2n)}$. In particular, $\ln D/(n \ln n) = 1/8 + o(1)$. Let $\mathcal{Q} = \pi_{\mathcal{M}^\perp}(\mathcal{L})$, so that $\det \mathcal{Q} = D^{-1}$.

When applying Lemma 5.1, we require the profile conclusions of Proposition 4.3 for the normalized lattices $D^{1/m}\mathcal{M}^*$ and $D^{1/(n-m)}\mathcal{Q}$. For the coordinate and first-vector claims below, we also require $r_d/2 \le \lambda_1 \le 2r_d$ in these two lattices and in $D^{-1/m}\mathcal{M}$. These properties hold with probability $1 - o(1)$ by (2.7) and Haar invariance under duality.

Let $\mathbf{C}$ and $\mathbf{E}$ be HKZ bases of these two normalized lattices, respectively. We undo both scalings, dualize and reverse the first basis, and append arbitrary lifts of the second basis. Since $\mathcal{M}$ is primitive, the resulting concatenation is a basis $\mathbf{B}$ of $\mathcal{L}$. We may size-reduce this basis without changing its Gram–Schmidt lengths. The prefix determinants obey the exact identities

$$P_k(\mathbf{B}) = \frac{k}{m} \ln D + P_{m-k}(\mathbf{C}), \qquad 0 \le k \le m, \tag{5.14}$$

$$P_{m+j}(\mathbf{B}) = \left(1 - \frac{j}{n-m}\right) \ln D + P_j(\mathbf{E}), \qquad 0 \le j \le n-m. \tag{5.15}$$

We now apply the prescribed profile estimates to $\mathbf{C}$ and $\mathbf{E}$. As $m/n \to 1/2$, the logarithms of both block dimensions are $\ln n + O(1)$. Together with $\ln D/(n \ln n) = 1/8 + o(1)$, this gives, uniformly for $0 \le x \le 1$:

$$\frac{P_{\lfloor xn \rfloor}(\mathbf{B})}{n \ln n} = p_2(x) + o(1), \qquad p_2(x) = \frac{v}{8} - \frac{v}{4} \ln v, \quad v = 2 \min\{x, 1-x\}. \tag{5.16}$$

The two maxima occur at $x = 1/(2\sqrt{e})$ and $x = 1 - 1/(2\sqrt{e})$, where $p_2(x) = 1/(4\sqrt{e})$. With probability $1 - o(1)$, this basis exists and Proposition 3.1 applies to every basis. Its normalized tree exponent is therefore $\max_x p_2(x) + o(1)$, proving (1.3). □

The Gram–Schmidt profile and the first vector. Proposition 4.3, with reflection and sign reversal in the first block, gives the curve in Figure 2:

$$g_2(x) = p_2'(x) = \begin{cases} -\frac{1}{4} - \frac{1}{2} \ln(2x), & 0 < x < \frac{1}{2}, \\ \frac{1}{4} + \frac{1}{2} \ln(2(1-x)), & \frac{1}{2} < x < 1. \end{cases} \tag{5.17}$$

The normalized GS lengths converge to this curve in averaged absolute error on every closed subinterval contained in either $(0, 1/2)$ or $(1/2, 1)$. At the split, the first lengths $c_1, e_1 = n^{1/2+o(1)}$ of $\mathbf{C}, \mathbf{E}$ give

$$a_m = \frac{D^{1/m}}{c_1} = n^{-1/4+o(1)}, \qquad a_{m+1} = D^{-1/(n-m)} e_1 = n^{1/4+o(1)}. \tag{5.18}$$

The ratio of the two Gram–Schmidt lengths at the split is therefore $n^{1/2+o(1)}$. For the first basis vector, $\lambda_1(\mathcal{L}) \le \sqrt{n}$ and the bound on $\lambda_1(D^{-1/m}\mathcal{M})$ give

$$\|\boldsymbol{b}_1\| \ge \lambda_1(\mathcal{M}) \ge D^{1/m} \frac{r_m}{2} = n^{3/4+o(1)}, \qquad \frac{\|\boldsymbol{b}_1\|}{\lambda_1(\mathcal{L})} \ge n^{1/4-o(1)}. \tag{5.19}$$

A stronger prediction for the first vector follows if we assume recursive GH, in the form $h_j = \frac{1}{2}\ln j + O(1)$ at every rank of $\mathbf{C}$. Under this additional assumption, (4.12) gives:

$$\ln \|\boldsymbol{b}_1\| = \frac{\ln D}{m} + \sum_{j=2}^{m} \frac{h_j}{j-1} = \frac{1}{4}(\ln n)^2 + O(\ln n). \tag{5.20}$$

The lower bound in (5.19) does not require this heuristic assumption.

5.3 Recursive splitting

The approach of finding a dense sublattice, splitting into blocks, and reducing each block separately, can be applied recursively. This recursive idea was already described by Hanrot–Stehlé [28], and recently studied by [1]. For any fixed finite binary tree with split fractions in $(0, 1)$, Lemma 5.1 gives the required splits with probability $1 - o(1)$. To see this, start with HKZ or reversed dual-HKZ bases at the leaves, and inductively require each normalized child lattice to admit the construction specified by its subtree.

For a split at fraction α, let $q_\alpha = \alpha(1-\alpha)/2$. The upper profiles of the two child lattices then combine as follows:

$$p(\alpha t) = t q_\alpha + \alpha p_{\mathrm{L}}(t), \tag{5.21}$$

$$p(\alpha + (1-\alpha)t) = (1-t)q_\alpha + (1-\alpha)p_{\mathrm{R}}(t), \qquad 0 \le t \le 1. \tag{5.22}$$

For a fixed tree, rounding and determinant slack contribute $o(n \ln n)$, so Proposition 3.1 gives $T(\mathbf{B}) \le n^{(\max p)n + o(n)}$.

To find trees with a small maximum prefix determinant, the accompanying script `recursive_binary_splits.py` performs the following search:

1. At depth 0, we tabulate the maxima of $p_{\mathrm{H}}(t) + st$ and $p_{\mathrm{D}}(t) + st$.
2. At each further depth, we balance the two child peaks using (5.21), with inherited tilts $s + (1-\alpha)/2$ and $s - \alpha/2$. We retain an earlier candidate if it gives a smaller maximum.
3. We reconstruct the finite tree and evaluate the maximum in each terminal block separately with high-precision arithmetic.

Proof (Theorem 1.3, with numerical evaluation). The numerical search first gives a coefficient below 0.1474 at depth 13. To leave a larger margin, we use its depth-14 candidate with 16 384 blocks. Evaluating the maximum in each terminal block separately gives:

$$\max_x p(x) = 0.147399340499\ldots < 0.14739935 < 0.1474. \tag{5.23}$$

The finite-tree construction above then proves the claimed bound as $n \to \infty$. The script and recorded trees accompany the article. The proof only uses this specific candidate, and does not require the numerical search to find an optimal tree. The 128-block curve in Figure 1 is the depth-7 example, and Figure 3 shows depths 0 through 10. □

6 Towards optimal bases under suitable heuristics

To approach the lower bound $(2-\sqrt{2})/4$, we use a different arrangement of blocks. The two outer blocks have fixed normalized width, while the remaining interval is divided into $K-2$ equally

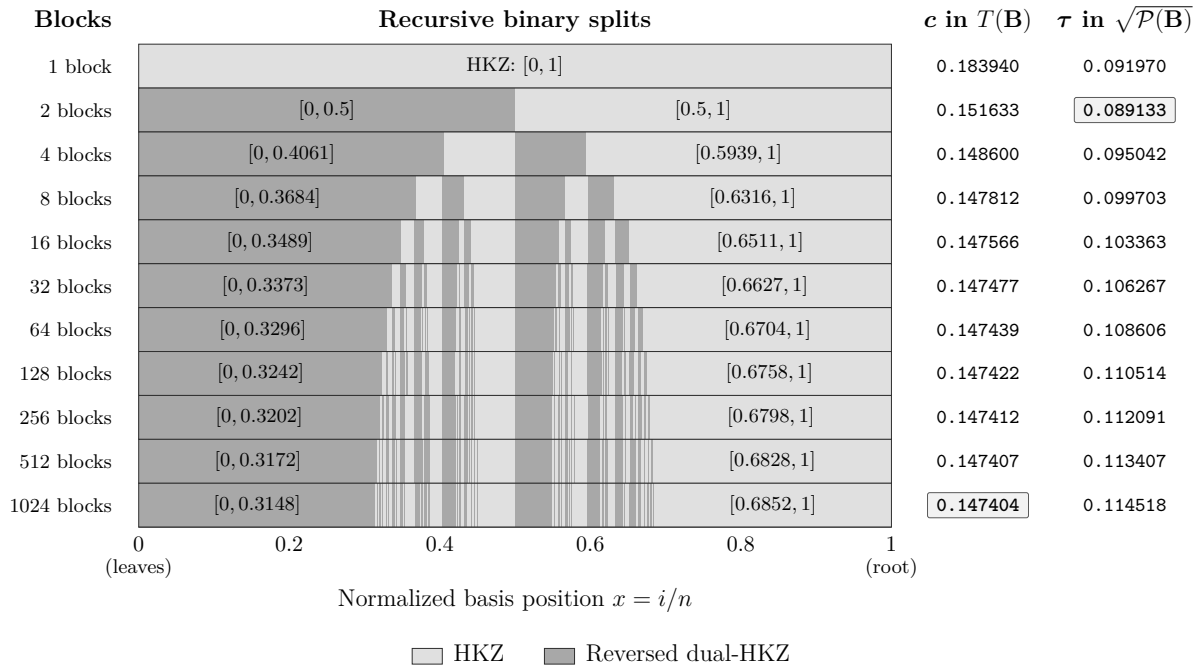


Fig. 3. Recursive binary splits for various recursion depths. As the depth increases, the computed upper bound on c stabilizes around $0.147397\ldots < 0.1474$. The same construction does not lower the FGL collision search cost constant τ beyond the binary split.

sized middle blocks. As K increases, the middle blocks become smaller. Under Assumption 1, this construction gives:

$$c_K \leq \frac{2-\sqrt{2}}{4}\left(1+\frac{\sqrt{2}}{e(K-2)}\right), \qquad K \geq 3. \tag{6.1}$$

The proof of Theorem 1.4 has two steps. We first verify that the prescribed chain satisfies the first-moment conditions in the assumption. We then complete this chain to a basis using HKZ and reversed dual-HKZ reduction inside the blocks.

Precise form of Assumption 1. Fix $r \geq 2$, rank fractions $0 = x_0 < \cdots < x_r = 1$, and determinant exponents $h_0 = h_r = 0, h_1, \ldots, h_{r-1}$, independently of n. Let $C_n(\mathcal{L})$ count chains of primitive sublattices satisfying

$$\{0\} = \mathcal{M}_0 \subset \cdots \subset \mathcal{M}_r = \mathcal{L}, \qquad \operatorname{rank} \mathcal{M}_i = \lfloor x_i n \rfloor, \qquad \det \mathcal{M}_i \leq n^{h_i n} \quad (1 \leq i < r). \tag{6.2}$$

For nonempty $J \subseteq \{1, \ldots, r-1\}$, let $C_{J,n}(\mathcal{L})$ count chains satisfying just the ranks and bounds indexed by J, whether or not they extend to a full chain. The assumption says that, if some fixed $\delta > 0$ satisfies

$$\min_{\varnothing \neq J \subseteq \{1,\ldots,r-1\}} \mathbb{E}_{\mu_n}[C_{J,n}(\mathcal{L})] \geq n^{\delta n^2} \tag{6.3}$$

for all sufficiently large n, then $\mathbb{P}\{C_n(\mathcal{L}) > 0\} = 1 - o(1)$. The condition is imposed on every subchain because a large full-chain average alone may hide a rank at which suitable partial chains are unlikely to exist.

For our construction, let $c = (2-\sqrt{2})/4$ and $a = 2c = 1 - 1/\sqrt{2}$. The two outer blocks each have width a, independently of K. We split the remaining interval of width $1 - 2a = \sqrt{2} - 1$ into $K - 2$ middle blocks, and impose the determinant bound n^{cn} at every internal boundary:

$$w = \frac{1-2a}{K-2}, \qquad x_i = a + (i-1)w \quad (1 \le i \le K-1), \qquad \det \mathcal{M}_i \le n^{cn}. \tag{6.4}$$

Lemma 6.1 (Positive first moments for every subchain). *For fixed $K \ge 3$, the ranks and bounds in* (6.4) *satisfy* (6.3) *with $\delta = (1-2a)^2/(8(K-2))$.*

Proof. Consider any nonempty set of boundaries $t_1 < \cdots < t_s$. Kim's chain formula [35, Theorem 5] and Stirling's formula give $\mathbb{E}[C_{J,n}] = n^{(\Gamma_J + o(1))n^2}$. Using $c = a/2$ and $a - a^2/2 = 1/4$, the exponent becomes a sum of squares:

$$\Gamma_J = c(1 + t_s - t_1) - \frac{1 - t_1^2 - (1-t_s)^2 - \sum_{\nu=1}^{s-1}(t_{\nu+1} - t_\nu)^2}{4} \tag{6.5}$$

$$= \frac{(t_1 - a)^2 + (1 - a - t_s)^2 + \sum_{\nu=1}^{s-1}(t_{\nu+1} - t_\nu)^2}{4}. \tag{6.6}$$

The distances in the last line are nonnegative integer multiples of w whose sum is $(K-2)w$. Their squared sum is at least $(K-2)w^2$, since $j^2 \ge j$ for every nonnegative integer j. Thus $\Gamma_J \ge (1-2a)^2/(4(K-2)) = 2\delta$. There are finitely many subchains for fixed K, so the same bound applies to all of them for sufficiently large n. □

Lemma 6.2 (Completing the blocks). *Any primitive chain with the ranks and determinant bounds in* (6.4) *extends to a basis with*

$$\max_j \frac{P_j(\mathbf{B})}{n \ln n} \le c + \frac{w}{2e} + o(1). \tag{6.7}$$

Proof. We reduce the first block by reversed dual-HKZ, the last block by HKZ, and each middle block by either orientation. Lifting and concatenating the resulting bases gives a basis of the full lattice. Within a block, the prefix profile is bounded by the linear interpolation of its two boundary heights plus its width times p_{H} or p_{D}, by Lemma 4.1. Smaller boundary determinants only lower this interpolation.

The left outer block therefore has upper profile $x(1 - \ln(x/a))/2$ for $0 \le x \le a$, with maximum $a/2 = c$; the right block is its reflection. Each middle block has both boundary heights at most c and width w, so its peak is at most $c + w \max p_{\mathrm{H}} = c + w/(2e)$, up to $o(1)$. This proves the bound. □

Proof (Theorem 1.4). Lemma 6.1 and Assumption 1 supply the chain with probability $1 - o(1)$. Lemma 6.2 and Proposition 3.1 then give (6.1), because

$$c + \frac{w}{2e} = \frac{2-\sqrt{2}}{4}\left(1 + \frac{\sqrt{2}}{e(K-2)}\right). \tag{6.8}$$

For any fixed $\varepsilon > 0$, we can choose K independently of n so that this bound is below $c + \varepsilon$. The resulting tree size is at most $n^{(c+\varepsilon)n + o(n)}$, proving (1.5).

□

Figure 4 shows this construction for total block counts $K = 4, 8, 16, \ldots$, alternating the orientations of the middle blocks. For comparison, the one- and two-block rows show the HKZ and two-block constructions. As K increases, the upper bound on the tree coefficient approaches the universal lower bound.

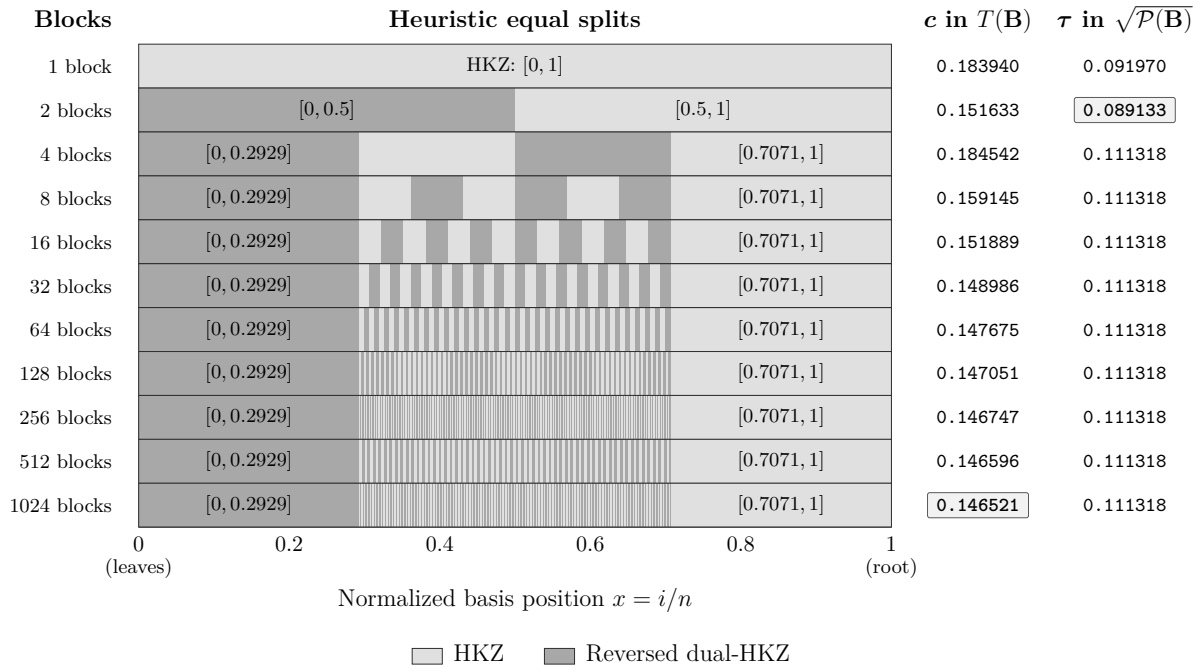


Fig. 4. Heuristic splits with equally sized middle blocks, using the same block counts as the recursive binary splits. The upper bound approaches the lower bound as $c_K = c_{\text{flat}} + O(1/K)$. The same construction does not lower the FGL collision search cost constant τ beyond the binary split.

7 Lower bounds on the tree size

To obtain a lower bound for all bases, we use that their primitive prefixes form nested chains. Kim's mean-value formula gives a lower bound on the area under a feasible prefix profile. By first reducing the two outer blocks, we turn this area constraint into the stronger lower bound on the maximum prefix determinant stated in Theorem 1.5.

7.1 The area rule

A basis determines one primitive sublattice at every prefix rank. To rule out bases whose prefix profiles all lie below a fixed curve, we count these complete nested chains. For a continuous ceiling $h : [0, 1] \to \mathbb{R}$, set $A(h) = \int_0^1 h(x)\,\mathrm{d}x$ and

$$F_n(h; \mathcal{L}) = \#\left\{\mathcal{M}_1 \subset \cdots \subset \mathcal{M}_{n-1} \subset \mathcal{L} : \begin{array}{l} \mathcal{M}_k \text{ primitive, } \operatorname{rank} \mathcal{M}_k = k, \\ \det \mathcal{M}_k \le n^{nh(k/n)} \quad (1 \le k < n) \end{array}\right\}. \tag{7.1}$$

We count chains of sublattices, so different bases representing the same chain contribute only once. The expectation of F_n is the average number of admissible chains in a Haar-random lattice. The next proposition shows that this average tends to zero if the area under h is below $1/8$. In that case, the probability that any such chain exists also tends to zero.

Proposition 7.1 (Area rule). *For every fixed continuous h,*

$$\mathbb{E}_{\mu_n}[F_n(h; \mathcal{L})] = n^{(2A(h) - 1/4 + o(1))n^2}. \tag{7.2}$$

If $A(h) < 1/8$, the probability that any basis satisfies all these prefix bounds is at most $n^{-(1/4-2A(h)+o(1))n^2}$.

Proof. For the complete rank sequence $(1, \dots, n-1)$, Kim's formula [35, Theorem 5] has determinant powers all equal to two and gives

$$\mathbb{E}_{\mu_n}[F_n(h;\mathcal{L})] = \frac{n!}{2^{n-1}} \left(\prod_{d=2}^{n} \kappa_d \right) \exp\left(2n \ln n \sum_{k=1}^{n-1} h(k/n) \right). \tag{7.3}$$

Stirling's formula gives logarithm $-n^2 \ln n/4 + O(n^2)$ for the prefactor. A Riemann sum proves (7.2), and Markov's inequality gives the probability bound. □

Applying the proposition to $h = p + \varepsilon$ shows that any fixed continuous limiting profile occurring with probability bounded away from zero must have area at least $1/8$. The area also has a determinant-product interpretation. If the normalized prefix profiles converge uniformly to a continuous function p, then

$$\int_0^1 p(x)\,dx = \lim_{n\to\infty} \frac{1}{n^2 \ln n} \ln\left(\prod_{k=1}^{n-1} \det M_k \right). \tag{7.4}$$

Products of covolumes along nested chains also occur in the theory of flag heights and generalized Hermite constants [42].

7.2 The universal lower bound

The area rule alone gives a lower bound of $1/8$ on the maximum prefix determinant, after normalization by $n \ln n$. To obtain the stronger constant $(2-\sqrt{2})/4$, we first replace the two outer parts of any hypothetical basis by reduced blocks. The reduced blocks admit curved upper bounds near the endpoints, while the middle prefix determinants remain unchanged.

Proof (Theorem 1.5). Fix $0 < c < (2-\sqrt{2})/4$, and suppose a basis satisfies $\max_k P_k \le cn \ln n$. We keep its prefix chain between ranks $\lfloor 2cn \rfloor$ and $\lfloor (1-2c)n \rfloor$, reduce the first block by reversed dual-HKZ, and reduce the final quotient by HKZ. The middle prefix determinants stay unchanged. Lemma 4.1 bounds the new left flank by $x(1 - \ln(x/(2c)))/2 + o(1)$, and the right flank by its reflection. Thus the entire new profile lies below $g_c + o(1)$, where

$$u_c(x) = \min\{x, 1-x, 2c\}, \qquad g_c(x) = \frac{u_c(x)}{2}\left(1 - \ln\frac{u_c(x)}{2c}\right). \tag{7.5}$$

The plateau has height c on $[2c, 1-2c]$. Each flank has area $3(2c)^2/8$, so $A(g_c) = c(1-4c) + 3c^2 = c - c^2 < 1/8$. We can therefore choose a fixed $\varepsilon > 0$ with $A(g_c + \varepsilon) < 1/8$. For sufficiently large n, the modified basis satisfies this ceiling, which the area rule excludes with probability $1 - o(1)$. Hence, with probability $1 - o(1)$, no original basis has $\max_k P_k \le cn \ln n$. Letting c approach $(2-\sqrt{2})/4$ sufficiently slowly as n grows, we obtain, with probability $1 - o(1)$:

$$\inf_{\mathbf{B} \text{ basis of } \mathcal{L}} \max_{0 \le k \le n} P_k(\mathbf{B}) \ge \left(\frac{2-\sqrt{2}}{4} - o(1) \right) n \ln n. \tag{7.6}$$

The error is independent of the basis. Proposition 3.1 then proves Theorem 1.5. □

7.3 Profile areas and the plateau

The HKZ/dual-HKZ profiles $p_{\mathrm{H}}, p_{\mathrm{D}}$ both have area $1/8$. Integrating (5.21) gives

$$A(p) = \frac{q_\alpha}{2} + \alpha^2 A(p_{\mathrm{L}}) + (1-\alpha)^2 A(p_{\mathrm{R}}) = \frac{1}{8} \tag{7.7}$$

whenever both child areas are $1/8$. It follows that every upper profile obtained from a fixed finite binary construction, including the orange 128-block example in Figure 1, has this area. A first split at $1/2$ fixes $p(1/2) = 1/8$, explaining the central dip retained by later refinements.

For comparison, the gray GSA profile gives:

$$p_{\mathrm{GSA}}(x) = \frac{x(1-x)}{2}, \qquad A(p_{\mathrm{GSA}}) = \frac{1}{12}, \qquad \max_x p_{\mathrm{GSA}}(x) = \frac{1}{8}. \tag{7.8}$$

The area rule therefore excludes this profile on Haar-random lattices with probability $1 - o(1)$. Its derivative $1/2 - x$ is the gray line in Figure 2.

Finally, consider the critical plateau profile, with $a = 1 - 1/\sqrt{2}$. The red curve satisfies:

$$p_{\mathrm{flat}} = g_{a/2}, \qquad A(p_{\mathrm{flat}}) = \frac{1}{8}, \qquad \max_x p_{\mathrm{flat}}(x) = \frac{a}{2} = \frac{2-\sqrt{2}}{4}. \tag{7.9}$$

Its derivative, shown in Figure 2, is

$$p'_{\mathrm{flat}}(x) = \begin{cases} \frac{1}{2}\ln(a/x), & 0 < x < a, \\ 0, & a \le x \le 1-a, \\ -\frac{1}{2}\ln(a/(1-x)), & 1-a < x < 1. \end{cases} \tag{7.10}$$

Under Assumption 1, the constructions in Section 6 give upper bounds on the prefix profile that approach p_{flat}, and attain its leading tree coefficient. This controls the prefix determinants, but does not imply that the individual Gram–Schmidt lengths converge to the derivative shown above. Whether the critical coefficient can be attained without a heuristic assumption remains open.

8 Applications

We now turn the geometric results into search and sampling bounds. In all preprocessing statements, the retained advice is a classical basis of polynomial encoding length, chosen independently of the target or sampling width. Preprocessing time is unrestricted. The space bounds concern the query or sampling phase, and polynomial factors in the encoded input lengths are suppressed when writing the leading exponent in n.

The four applications optimize different geometric quantities. Classical enumeration and quantum enumeration depend on the largest prefix determinant. Collision search and the width-uniform bound for corrected Klein sampling depend on $\mathcal{P}(\mathbf{B}) = \prod_i \max\{1, \|\boldsymbol{b}_i^*\|^{-1}\}$. The recursive construction is therefore used for enumeration, whereas the two-block construction supplies the best bounds established here for collision search and corrected Klein sampling.

8.1 Classical enumeration

We first turn the prefix bounds into a query algorithm by traversing the full enumeration tree. On the common high-probability event from Section 3.2, the shifted counting bound (3.6) holds for

every basis and every target. Banaszczyk's covering-radius bound [8], together with the shortest-vector estimate for the dual lattice, also gives a target-independent radius $R = 10\sqrt{n}$ containing a closest vector to every target. Summing over all levels therefore gives

$$Q_{\mathrm{enum}}(\mathbf{B}, \boldsymbol{t}) \le \mathrm{poly}(\text{input length}) \exp\left(\max_{0 \le k \le n} P_k(\mathbf{B}) + O(n) \right). \tag{8.1}$$

Depth-first traversal retains a closest full-depth leaf and stores only the basis, the target, one path, and the best candidate found so far.

Preprocessing and advice. All preprocessing is independent of the target. Its running time is unrestricted, and only the selected basis is retained as advice. The size of this advice also needs to be controlled. On the event used in Proposition 3.1, (3.4) and $N_{n-k} \ge 1$ imply $P_k \ge -O(n)$ for every basis prefix. If $\max_k P_k = O(n \ln n)$, then $|y_i| = O(n \ln n)$, and size reduction gives $\|\boldsymbol{b}_i\| \le \exp(O(n \ln n))$. Relative to an encoded input basis, the integer change-of-basis matrix has polynomial bit length. Preprocessing can thus restrict its search to finitely many changes of basis of polynomial bit length; the existence results guarantee a successful choice on the stated event. Query time and space include polynomial factors in the encoded input lengths, as in Section 2.

Corollary 8.1 (Classical CVPP from split bases). *With probability $1 - o(1)$ over $\mathcal{L} \sim \mu_n$, there is polynomial-size advice depending only on $\mathcal{L}$ such that, for every target $\boldsymbol{t}$, a deterministic algorithm finds an exact closest lattice vector using polynomial query space and query time*

$$Q_{\mathrm{enum}} \le n^{c_{\mathrm{rec}} n + o(n)}, \qquad c_{\mathrm{rec}} < 0.14740. \tag{8.2}$$

Using only the two-block construction gives the bound $n^{c_2 n + o(n)}$, where $c_2 = 1/(4\sqrt{e})$. Under Assumption 1, for every fixed $\varepsilon > 0$ there is polynomial-size advice giving the query bound

$$Q_{\mathrm{enum}} \le n^{(c_{\mathrm{flat}} + \varepsilon) n + o(n)}. \tag{8.3}$$

Proof. Apply (8.1) to the prefix ceilings of Theorems 1.2, 1.3, and 1.4. In the conditional case, choose a fixed number of blocks as a function of ε before letting n grow. The common geometric event is uniform over targets and bases, so preprocessing selects the advice before the target is known.

SVP on a supplied basis. For SVP, take the target to be zero, use $R = \sqrt{n}$, and discard the zero vector. Minkowski's bound ensures that the ball contains a shortest nonzero vector. The same time exponents apply to the search on the supplied basis. These are not end-to-end SVP bounds including the cost of constructing the basis.

8.2 Collision-search enumeration

Using the search radius $R = \sqrt{n}$ in the Feng–Gao–Liu (FGL) sampling and collision procedure [19] gives the geometric product

$$\mathcal{P}(\mathbf{B}) = \prod_{i=1}^{n} \max\{1, a_i^{-1}\}. \tag{8.4}$$

The dominant geometric factor in the FGL running-time bound is $\sqrt{\mathcal{P}(\mathbf{B})}$. We first bound this factor from below for all bases, and then determine it for the two-block construction. To obtain an algorithmic upper bound, we also control the cost of sampling from the chosen basis. Theorem 1.6

concerns the geometric factor in the analysis; it does not give a lower bound on the actual running time of the procedure.

The key difference with enumeration is that collision search depends on the total variation of the prefix profile, whereas enumeration depends on its maximum. To make this precise, write $[u]_+ = \max\{u, 0\}$, let TV denote total variation on $[0, 1]$, and let p_n linearly interpolate $p_n(k/n) = P_k/(n \ln n)$.

Proof (Theorem 1.6). Since $\sum_i y_i = 0$,

$$\frac{\ln \sqrt{\mathcal{P}(\mathbf{B})}}{n \ln n} = \frac{\sum_i [-y_i]_+}{2n \ln n} = \frac{1}{4} \operatorname{TV}(p_n) \geq \frac{\max_k P_k}{2n \ln n}. \tag{8.5}$$

The universal prefix lower bound (7.6) gives $\max_k P_k \geq ((2 - \sqrt{2})/4 - o(1))n \ln n$, simultaneously for all bases, with probability $1 - o(1)$. Because both the error term and the inequality above are uniform over the bases, the resulting lower bound holds simultaneously for all bases, with the leading constant divided by two. Multiplying the search radius by any fixed positive constant changes the logarithm of the cost factor by at most $O(n)$, uniformly in the basis. The same leading lower bound therefore applies to the radius used for CVPP in Section 8.2. □

Proof (Theorem 1.7). The normalized HKZ blocks used in Section 5.2 have dimensions $d = n/2 + O(1)$ and shortest-vector lengths $\rho_1 \asymp \sqrt{d}$, by the shortest-vector properties imposed in the construction. Let ρ_i denote the Gram–Schmidt lengths of either normalized block. The HKZ product bound [19, Lemma 1] gives, for every subset I of size k,

$$-\sum_{i \in I} \ln \rho_i \leq \frac{k}{2} \ln(d/k) \ln n + O(n). \tag{8.6}$$

Here and below the error is uniform in I, and $0 \ln(d/0) = 0$. Consequently,

$$\sum_{i=1}^{d} \left[\frac{1}{4} \ln n - \ln \rho_i \right]_+ \leq \ln n \max_{0 \leq k \leq d} \left\{ \frac{k}{4} + \frac{k}{2} \ln(d/k) \right\} + O(n) = \left(\frac{1}{2\sqrt{e}} + o(1) \right) d \ln n. \tag{8.7}$$

Rescaling the second block multiplies its lengths by $n^{-1/4} e^{O(1)}$, so the estimate above bounds its contribution to $\sum_i [-y_i]_+$. The first block is reversed and dualized, then scaled by $n^{1/4} e^{O(1)}$. Since $\sum_i \ln \rho_i = 0$, its contribution is bounded by $(1/(2\sqrt{e}) - 1/4 + o(1)) d \ln n$. Adding both contributions and dividing by $2n \ln n$ proves the upper bound in (1.10). In fact, these estimates give $\ln \mathcal{P}(\mathbf{B}_2) \leq 2\tau_2 n \ln n + O(n)$.

For the reverse inequality, Section 5.2 gives uniform convergence $p_n \to p_2$. The two peaks of p_2 have height $1/(4\sqrt{e})$, the intervening dip has height $1/8$, and its endpoints are zero. Evaluating the variation along these five points gives $\operatorname{TV}(p_n) \geq 1/\sqrt{e} - 1/4 - o(1)$. Equation (8.5) therefore gives a matching lower bound, showing that the constructed basis satisfies $\sqrt{\mathcal{P}(\mathbf{B}_2)} = n^{\tau_2 n + o(n)}$. □

Search on the preprocessed basis. To turn the product estimate into a search bound, we must control the sampling cost and choose an appropriate radius. Applying the same HKZ product bound with $k = 1$ and $k = d - 1$, and using $\prod_i \rho_i = 1$, gives $\max_i \max\{1, a_i^{-1}\} \leq e^{O(n)} = n^{o(n)}$ for $\mathbf{B}_2$. This makes the cost of each sample $n^{o(n)}$.

For the radius, we cannot use the length of the first basis vector: by (5.19), this length is not within a constant factor of $\lambda_1(\mathcal{L})$. We therefore replace the cutoff $\|\boldsymbol{b}_1\|$ in [19, Sections 3–4] by $R = \sqrt{n}$ and retain the shortest nonzero difference found. Minkowski gives $\lambda_1(\mathcal{L}) \leq R$,

while the Haar shortest-vector bound gives $R/\lambda_1(\mathcal{L}) = O(1)$ with high probability. The short-translation argument therefore still applies, and grid cells contain at most $\exp(O(n \ln \ln n))$ lattice points by the same packing bound. Together with the sampling bound above, the collision analysis gives $n^{o(n)}\sqrt{\mathcal{P}(\mathbf{B}_2)} = n^{\tau_2 n + o(n)}$ time and polynomial space, with constant success probability and polynomial encoding factors suppressed. This bounds the search time once the basis is available; constructing the basis is part of preprocessing.

Further binary splits. The recursive splits in Section 5.3 were chosen to minimize the largest prefix determinant. For collision search, however, we need to minimize $\mathrm{TV}(p)/4$. Evaluating the binary profiles of Section 5.3 gives collision coefficients approximately 0.09504173 for four blocks and 0.11758925 for 16 384 blocks, compared with 0.08913266 for two blocks. Continuing the balanced refinement suggests that the collision coefficient approaches 1/8, even though the enumeration coefficient decreases towards 0.147397 .... The limits refer to the analytic profiles, taking $n \to \infty$ at each fixed depth before increasing the depth. The reason is that further splits lower the largest prefix peak while introducing additional rises and falls. These particular splits therefore do not appear to improve collision search.

The heuristic block construction. The analytic profiles of Section 6 also give a larger collision coefficient than the two-block construction. The limiting coefficient is $(2-\sqrt{2})/8 + (\sqrt{2}-1)/(4e) = 0.11131847\ldots$. In fact, in the ideal block-profile model it already has this value for every fixed total number $K \geq 3$ of blocks, with $K-2$ middle blocks. Smaller blocks reduce the height of each prefix ripple, but their total variation stays unchanged. This value exceeds both the single-split coefficient and the HKZ coefficient $1/(4e)$, although it is below the predicted limit 1/8 of the deeply refined binary profiles.

These are calculations for the analytic profiles: the first-moment assumption in Section 6 supplies prefix ceilings and does not by itself establish the corresponding Gram–Schmidt variation or search cost. These comparisons concern the specific enumeration constructions above; they do not establish optimality of the two-block basis among all configurations allowed by Assumption 1. Constructing bases specifically optimized for collision search, and determining whether its lower bound $(2-\sqrt{2})/8$ is attainable, remain questions for future work.

Proof (Theorem 1.8). We use the two-block basis from Section 5.2, together with the product and coordinate estimates established in Section 8.2. With probability $1 - o(1)$, this basis exists and the shortest-vector bounds in (2.7) hold for both the lattice and its dual. We work on this event, on which $\lambda_1(\mathcal{L}) \asymp \sqrt{n}$. By Section 8.1, the fixed radius $R = 10\sqrt{n}$ contains a closest vector to every target. To account for this radius, define:

$$\mathcal{P}_R(\mathbf{B}) = \prod_i \max\left\{1, \frac{R}{\sqrt{n}\, a_i}\right\}. \tag{8.8}$$

Thus $\mathcal{P}(\mathbf{B}) \leq \mathcal{P}_R(\mathbf{B}) \leq 10^n \mathcal{P}(\mathbf{B})$, and the maximum coordinate scale remains $n^{o(n)}$.

Fix any target $\boldsymbol{t}$. If $\boldsymbol{t} \in \mathcal{L}$, return $\boldsymbol{t}$. Otherwise, let $\boldsymbol{z}^*$ be a closest lattice point and put $\boldsymbol{v}^* = \boldsymbol{t} - \boldsymbol{z}^*$, so $\|\boldsymbol{v}^*\| \leq R$. We sample one implicit array in $\mathcal{L}$ and another in $\boldsymbol{t} + \mathcal{L}$, both using rounded Gram–Schmidt coordinates centered at the origin:

$$\kappa_i(\boldsymbol{x}) = \left\lfloor \frac{\langle \boldsymbol{x}, \boldsymbol{b}_i^* \rangle}{a_i^2} \right\rceil. \tag{8.9}$$

Here $\lfloor u \rceil = \lfloor u + 1/2 \rfloor$. Restricted to either coset, κ is a bijection with $\mathbb{Z}^n$. Indeed, for $\boldsymbol{x} = \boldsymbol{t} + \mathbf{B}\boldsymbol{z}$, each integer z_i is recovered uniquely from $\kappa_i(\boldsymbol{x})$ once $z_{i+1}, \ldots, z_n$ are known. The same applies

with $\boldsymbol{t} = \boldsymbol{0}$. We apply the FGL product distribution to these integer coordinates, using its rounded weights $\alpha_i \le \min\{1, na_i^2/R^2\}$. This gives distributions ν_0 on $\mathcal{L}$ and $\nu_{\boldsymbol{t}}$ on $\boldsymbol{t} + \mathcal{L}$.

Translation by $\boldsymbol{v}^*$ maps $\mathcal{L}$ bijectively onto $\boldsymbol{t} + \mathcal{L}$. The rounding error in the difference of each coordinate is at most one, and hence, for every $\boldsymbol{x} \in \mathcal{L}$,

$$\sum_i \alpha_i |\kappa_i(\boldsymbol{x} + \boldsymbol{v}^*) - \kappa_i(\boldsymbol{x})|^2 \le \frac{2n}{R^2}\|\boldsymbol{v}^*\|^2 + 2n \le 4n. \tag{8.10}$$

The central-box argument of [19, Theorem 3.1] therefore applies to these two coset bijections without change. It gives maximum point masses at most $2^{O(n)}/\mathcal{P}_R(\mathbf{B})$ and

$$\sum_{\boldsymbol{x}\in\mathcal{L}} \nu_0(\boldsymbol{x})\nu_{\boldsymbol{t}}(\boldsymbol{x} + \boldsymbol{v}^*) \ge \frac{2^{-O(n)}}{\mathcal{P}_R(\mathbf{B})}. \tag{8.11}$$

Neither sampler requires knowledge of $\boldsymbol{z}^*$ or $\boldsymbol{v}^*$. Independent arrays of length $2^{O(n)}\mathcal{P}_R(\mathbf{B})$, generated from short seeds as in FGL, consequently have the same representation and second-moment bounds, up to exponential factors. Those arguments use only the point-mass and overlap bounds and remain valid for the two different marginal distributions.

A common scaling clears the input denominators and places both cosets in $\mathbb{Z}^n$. We scale the search radius by the same factor, leaving the coordinate scales and all ratios used below unchanged. This permits the exact discrete-grid implementation of FGL for arbitrary targets.

We then apply the FGL random grid at scale $W = O(R\sqrt{\ln n})$. A pair with difference $\boldsymbol{v}^*$ shares a cell with probability $2^{-O(n)}$. Packing each coset separately bounds the number of distinct points of their union in a cell by

$$2\left(1 + \frac{2W}{\lambda_1(\mathcal{L})}\right)^n = \exp(O(n \ln\ln n)). \tag{8.12}$$

This bound is uniform in $\boldsymbol{t}$, even when the two cosets are very close. The collision lemma [19, Lemma 4] can therefore be applied with the good pairs restricted to cross-array pairs whose difference is $\boldsymbol{v}^*$. We initialize the retained candidate with $\boldsymbol{0}$ and ignore pairs from the same array. For every reported cross-array pair $X \in \mathcal{L}$, $Y \in \boldsymbol{t} + \mathcal{L}$, the vector $\boldsymbol{t} - (Y - X)$ belongs to $\mathcal{L}$. Throughout the prescribed search budget, we retain the candidate closest to $\boldsymbol{t}$. With constant probability a pair with difference $\boldsymbol{v}^*$ is reported, making the retained candidate exact. A constant number of independent runs raises this probability above $2/3$.

As in FGL, the resulting query time is at most $n^{o(n)}\sqrt{\mathcal{P}_R(\mathbf{B})} = n^{(\tau_2 + o(1))n}$, with polynomial space. Every geometric bound is independent of the target, so one basis supports all targets, with fresh randomness for each query.

It remains to describe the preprocessing and the size of the retained advice. Preprocessing may enumerate primitive half-rank sublattices, perform the two block reductions, and test the resulting product and coordinate bounds. The estimates above have $e^{O(n)}$ slack: for example, thresholds $\mathcal{P}(\mathbf{B}) \le n^{2\tau_2 n}e^{n\sqrt{\ln n}}$ and $\max_i \max\{1, a_i^{-1}\} \le e^{n\sqrt{\ln n}}$ therefore suffice for all sufficiently large n. Rounding these thresholds upward to powers of two makes the tests exact and changes the query exponent only by $o(1)$.

The product bound also gives polynomial encoding length after size reduction: since $\sum_i y_i = 0$, one has $|y_i| \le \ln \mathcal{P}(\mathbf{B}) = O(n \ln n)$ and $\|\boldsymbol{b}_i\| \le \sqrt{n}\,\mathcal{P}(\mathbf{B})$. As in Section 8.1, the corresponding integer basis changes have polynomial bit length relative to the input. One may therefore restrict preprocessing to a finite search over such changes; the bounds above guarantee success on the stated event. Only the chosen basis is retained as advice. Preprocessing has unrestricted running time and is independent of the target. □

8.3 Quantum enumeration

The quantum enumeration algorithm of Aono–Nguyen–Shen [6] allows us to use the smaller trees directly. The following consequence of the uniform shifted bound isolates the dependence on the basis.

Lemma 8.2 (Quantum CVPP from prefix bounds). *With probability* $1-e^{-\Omega(n)}$ *over* $\mathcal{L}\sim\mu_n$*, the following holds simultaneously for every basis* $\mathbf{B}$ *of polynomial encoding length and every target* $\boldsymbol{t}$*. There is a quantum algorithm finding an exact closest lattice vector with probability at least* $2/3$*, using polynomial classical and quantum space and query time*

$$Q_{\mathrm{quant}}(\mathbf{B},\boldsymbol{t}) \le \exp\left(\frac{1}{2}\max_{0\le k\le n} P_k(\mathbf{B}) + O(n)\right), \tag{8.13}$$

with polynomial factors in the encoded input length suppressed. The constant implicit in $O(n)$ *is independent of the basis and target.*

Proof. Work on the common high-probability event used in Section 8.1, on which the shifted counting bound holds and $R = 10\sqrt{n}$ is at least the covering radius. Let $T_R(\mathbf{B},\boldsymbol{t})$ be the number of nodes in the full enumeration tree for the ball centered at $\boldsymbol{t}$ with radius R. Summing the bound (3.6) over its levels gives

$$T_R(\mathbf{B},\boldsymbol{t}) \le \exp\left(\max_k P_k(\mathbf{B}) + O(n)\right). \tag{8.14}$$

At least one full-depth leaf represents a closest vector.

The tree has depth n, and the admissible values of the next integer coefficient form an interval. Its children can therefore be counted and indexed in polynomial time. Writing $a_i = \|\boldsymbol{b}_i^*\|$, its maximum degree is at most $1 + 2R/\min_i a_i$, whose logarithm is polynomial in the input length. Binary tree conversion consequently has polynomial depth and increases the number of nodes by at most a polynomial factor.

Apply the minimum-finding version of quantum backtracking from [6, Theorem 6] to the full-depth leaves, with objective $\|\boldsymbol{v}-\boldsymbol{t}\|^2$. In the encoded model, clearing denominators makes this an integer objective with polynomial bit length. The algorithm returns a minimizer with probability at least $2/3$ in $\mathrm{poly}(\text{input length})\sqrt{T_R(\mathbf{B},\boldsymbol{t})}$ time and polynomial space. This proves the bound. □

Proof (Proof of Theorem 1.9). Choose the basis supplied by Theorem 1.3. Its maximum logarithmic prefix determinant is at most $(c_{\mathrm{rec}} + o(1))n\ln n$. Section 8.1 shows that preprocessing can retain such a basis with polynomial encoding length. Lemma 8.2 gives the first query bound.

For the conditional bound, fix $\varepsilon > 0$. Under Assumption 1, Theorem 1.4 supplies a basis with maximum logarithmic prefix determinant at most $(c_{\mathrm{flat}} + 2\varepsilon + o(1))n\ln n$. The number of blocks is chosen as a function of ε and is fixed as n grows. The same lemma gives the second bound. In both cases, the advice is the chosen classical basis, preprocessing is independent of the target, and the geometric bounds hold uniformly over all targets. The success probability refers to each query separately. □

SVP on a supplied basis. Taking the target to be zero and excluding the zero leaf gives the same square-root search bounds for SVP on an available basis. A radius $R = \sqrt{n}$ suffices by Minkowski. Basis construction remains part of preprocessing.

Comparison with collision search. Let p_n be the piecewise-linear normalized prefix profile. Suppressing polynomial encoding factors, the two geometric expressions in our query analyses satisfy

$$\frac{\ln Q_{\mathrm{quant}}(\mathbf{B}, \boldsymbol{t})}{n \ln n} \leq \frac{1}{2} \max_x p_n(x) + o(1), \qquad \frac{\ln \sqrt{\mathcal{P}(\mathbf{B})}}{n \ln n} = \frac{1}{4} \mathrm{TV}(p_n). \tag{8.15}$$

Since $p_n(0) = p_n(1) = 0$, we have $\mathrm{TV}(p_n) \geq 2 \max_x p_n(x)$, with equality for a profile that rises to its maximum and then falls, with possible constant plateaus in between. Additional rises and falls increase the FGL factor without increasing the maximum height.

For quantum enumeration, the prefix ceilings of Section 6 are sufficient: their maximal excess above c_{flat} tends to zero as the fixed number of blocks grows. No control of the limiting total variation is required. The conditional quantum exponent therefore approaches $c_{\mathrm{flat}}/2$, the same coefficient as the lower bound on the FGL geometric factor in Theorem 1.6. This equality of coefficients is not a lower bound on quantum CVPP, nor does it show that FGL attains the same coefficient on these bases.

8.4 Discrete Gaussian sampling with polynomial space

To prove Theorem 1.10, we use Klein's sequential sampler [36] with the rejection correction of Brakerski et al. [11, full version, Section 5.2]. The correction gives the desired distribution at every width; its cost depends on the Gram–Schmidt lengths of the supplied basis.

Extending the notation $\rho = \rho_1$ from Section 2.3, write

$$\rho_s(S) = \sum_{\boldsymbol{x} \in S} e^{-\pi \|\boldsymbol{x}\|^2 / s^2}, \qquad D_{\mathcal{L},s}(\boldsymbol{x}) = \frac{e^{-\pi \|\boldsymbol{x}\|^2 / s^2}}{\rho_s(\mathcal{L})} \qquad (\boldsymbol{x} \in \mathcal{L}). \tag{8.16}$$

Throughout this subsection, $\mathcal{L}$ is unimodular and $s > 0$. Recall that $a_i = \|\boldsymbol{b}_i^*\|$ and $\mathcal{P}(\mathbf{B}) = \prod_i \max\{1, a_i^{-1}\}$.

Lemma 8.3 (Sampling cost from a basis). *For every basis* $\mathbf{B}$ *of* $\mathcal{L}$*, corrected Klein sampling outputs* $D_{\mathcal{L},s}$ *exactly when the one-dimensional Gaussian draws and acceptance coins are exact. The expected number of complete proposals is*

$$K(\mathbf{B}, s) = \frac{M_{\mathbf{B}}(s)}{\rho_s(\mathcal{L})} \leq 2^n \mathcal{P}(\mathbf{B}), \qquad M_{\mathbf{B}}(s) = \prod_{i=1}^{n} \rho_{s/a_i}(\mathbb{Z}). \tag{8.17}$$

Proof. Klein's sampler chooses the coefficients of $\boldsymbol{x} = \mathbf{B}\boldsymbol{z}$ in reverse order. Given $z_{i+1}, \ldots, z_n$, it draws $z_i \in \mathbb{Z}$ with weight proportional to

$$\exp\left(-\frac{\pi a_i^2 (z_i - c_i)^2}{s^2}\right), \qquad c_i = -\sum_{j>i} \mu_{j,i} z_j.$$

Orthogonality of the Gram–Schmidt vectors gives the proposal probability

$$q_{\mathbf{B},s}(\boldsymbol{x}) = \frac{e^{-\pi \|\boldsymbol{x}\|^2 / s^2}}{A_{\mathbf{B},s}(\boldsymbol{x})}, \qquad A_{\mathbf{B},s}(\boldsymbol{x}) = \prod_{i=1}^{n} \rho_{s/a_i}(\mathbb{Z} - c_i). \tag{8.18}$$

Poisson summation gives $\rho_u(\mathbb{Z} - c) \leq \rho_u(\mathbb{Z})$, so the ratio $A_{\mathbf{B},s}(\boldsymbol{x})/M_{\mathbf{B}}(s)$ is a valid acceptance probability. The probability of proposing and accepting $\boldsymbol{x}$ is therefore $e^{-\pi \|\boldsymbol{x}\|^2 / s^2} / M_{\mathbf{B}}(s)$. Summing over $\mathcal{L}$ proves the output distribution and the equality in (8.17).

The one-dimensional estimate $\rho_u(\mathbb{Z}) \le 1 + u \le 2\max\{1, u\}$ bounds the numerator. For the denominator, the zero vector and Poisson summation give $\rho_s(\mathcal{L}) \ge \max\{1, s^n\}$. Hence

$$K(\mathbf{B}, s) \le 2^n \frac{\prod_i \max\{1, s/a_i\}}{\max\{1, s^n\}} \le 2^n \mathcal{P}(\mathbf{B}). \tag{8.19}$$

For $s \le 1$, bound each numerator factor by $\max\{1, a_i^{-1}\}$. For $s \ge 1$, divide each factor by s and use $\max\{s^{-1}, a_i^{-1}\} \le \max\{1, a_i^{-1}\}$. □

Proof (Theorem 1.10). Choose the two-block basis $\mathbf{B}_2$ and size-reduce it, preserving its Gram–Schmidt lengths. The product estimate of Theorem 1.7 gives $\mathcal{P}(\mathbf{B}_2) \le n^{2\tau_2 n + o(n)}$. The encoding bound in Section 8.2 ensures that this basis has polynomial size. Choose a power of two U satisfying

$$2^n \mathcal{P}(\mathbf{B}_2) \le U < 2^{n+1} \mathcal{P}(\mathbf{B}_2), \qquad N = U(n^3 + 2).$$

Run at most N corrected proposals, returning zero if none is accepted. Lemma 8.3 bounds the ideal timeout probability by

$$(1 - U^{-1})^N \le e^{-N/U} = e^{-n^3 - 2} < \frac{1}{4} e^{-n^3}.$$

Set $\delta = e^{-n^3}/(4(n+1)N)$. Implement each one-dimensional draw with statistical error at most δ and each acceptance coin with probability error at most δ. The one-dimensional sampler and theta evaluations in [11, full version, Sections 5.1–5.2] allow this with polynomial cost in the input lengths and $\ln(1/\delta)$, after truncating tails and internal rejection loops. Each draw's error budget includes these truncations and numerical errors. For the acceptance coin, compute each of the n theta ratios to error at most $\delta/(4n)$, using the primal sum when $s/a_i < 1$ and its Poisson form otherwise.

There are at most $(n + 1)N$ such calls. Coupling the implementation to the ideal procedure bounds their total contribution to the statistical distance by $(n + 1)N\delta = e^{-n^3}/4$. Together with the timeout error, this is less than e^{-n^3}. Moreover,

$$\ln N = O(n \ln n), \qquad \ln(1/\delta) = n^3 + O(n \ln n),$$

so the required precision contributes only a polynomial factor to the running time.

For the space bound, truncate each one-dimensional draw so that

$$|z_i - c_i| \le 1 + O\left(\frac{s}{a_i}\sqrt{\ln(1/\delta)}\right).$$

Since $|\mu_{j,i}| \le 1/2$, reverse induction gives

$$\max_i |z_i| \le 2^{O(n)} \max_i\{1, s/a_i\}\, O\left(\sqrt{\ln(1/\delta)}\right).$$

All stored coefficients therefore have bit length polynomial in ℓ. Storing the basis, one proposal, the numerical workspace and a counter suffices; the workspace is reused after each rejection. □

For comparison, on Haar-random lattices the HKZ product estimate [19, Lemma 1], combined with Lemma 8.3, gives the exponent $1/(2e) = 0.18393\ldots$ for the same sampler. The two-block basis lowers this bound without increasing the sampling space. Basis construction is part of preprocessing and is excluded from (1.14).

The dependence on $\mathcal{P}(\mathbf{B})$ also explains why smaller enumeration trees need not improve this sampler. At $s = 1$, $\mathcal{P}(\mathbf{B}) \le M_{\mathbf{B}}(1) \le 2^n \mathcal{P}(\mathbf{B})$. Siegel's mean-value formula gives $\rho_1(\mathcal{L}) \le e^n$ with probability at least $1 - 2e^{-n}$. On this event, simultaneously for all bases,

$$\frac{\ln K(\mathbf{B}, 1)}{n \ln n} = \frac{1}{2} \operatorname{TV}(p_n) + O(1/\ln n), \tag{8.20}$$

where p_n is the normalized prefix profile used in (8.5). Thus the relevant quantity is its total variation, rather than its maximum height.

9 Discussion on worst-case lattices

Our results separate the performance of HKZ bases from the best performance attainable by any basis of the same lattice. This distinction also matters for worst-case examples. Hanrot–Stehlé [27] construct HKZ-reduced bases whose enumeration trees attain the exponent $1/(2e)$. This establishes sharpness of the HKZ analysis, but leaves open whether different bases of the underlying lattices permit smaller trees.

9.1 The Hanrot–Stehlé worst-case HKZ examples

The explicit Gram–Schmidt lengths of Hanrot–Stehlé [27, Section 4], after normalization to determinant one, give $\lambda_1(\mathcal{L}_n) = \Theta(\sqrt{n})$ and the same leading logarithmic prefix profile p_{H} as in (4.3). In particular, their designated HKZ basis has a middle prefix of determinant $n^{((\ln 2)/4 + o(1))n}$. This profile alone does not exclude a denser half-rank primitive sublattice spanned by other combinations of lattice vectors.

Following the construction of Hanrot–Stehlé [27, Sections 3–4], let ν_n be the distribution obtained by fixing these Gram–Schmidt lengths, drawing the size-reduction coefficients independently and uniformly in $[-1/2, 1/2]$, conditioning on the full basis being HKZ-reduced, and rescaling the lattice to determinant one.

We expect the dense-sublattice principle used for Haar-random lattices to extend to this ensemble: for every fixed $\eta > 0$, a lattice drawn from ν_n should, with probability $1 - o(1)$, contain a primitive rank-$\lfloor n/2 \rfloor$ sublattice of determinant at most $n^{(1/8+\eta)n}$. For sufficiently small fixed $\eta > 0$, the split identities (5.14) and Lemma 4.1 give $\max_k P_k(\mathbf{B}) \le (c_2 e^{4\eta} + o(1)) n \ln n$ for the resulting split basis $\mathbf{B}$. Since $\lambda_1 = \Theta(\sqrt{n})$, packing and Poisson summation give $\rho(\mathcal{L}^*) \le e^{O(n)}$, so the upper estimate in (3.4) applies as well. Choosing η sufficiently small therefore gives, for every fixed $\varepsilon > 0$, the predicted tree bound

$$n^{(c_2+\varepsilon)n + o(n)}, \qquad c_2 = \frac{1}{4\sqrt{e}} < \frac{1}{2e}. \tag{9.1}$$

This is a heuristic prediction for the conditioned ensemble, rather than a consequence of Assumption 1, which concerns Haar measure. In particular, the correlations between candidate sublattices and the effect of conditioning on HKZ reduction remain to be justified. The known Hanrot–Stehlé examples therefore appear compatible with improving splits under this extension of the heuristic; they do not presently provide an obstruction to smaller trees for other bases.

9.2 Rankin constants and the remaining worst-case question

The general question is related to the unresolved asymptotics of Rankin's constants [23]. For a unimodular rank-n lattice, put

$$d_m(\mathcal{L}) = \min_{\substack{\mathcal{M} \subset \mathcal{L} \text{ primitive} \\ \operatorname{rank} \mathcal{M} = m}} \det \mathcal{M}, \qquad \gamma_{n,m} = \sup_{\det \mathcal{L} = 1} d_m(\mathcal{L})^2. \tag{9.2}$$

For $m = \lfloor n/2 \rfloor$, the sublattice first-moment formula and the deterministic HKZ prefix bound give

$$n^{(1/8-o(1))n} \le \sqrt{\gamma_{n,m}} \le n^{((\ln 2)/4+o(1))n}. \tag{9.3}$$

The lower estimate follows from (2.12), since, for every fixed $\eta > 0$, the Haar expected count tends to zero below $n^{(1/8-\eta)n}$; the upper estimate follows by taking a middle HKZ prefix and applying Lemma 4.1. The gap leaves open whether every lattice admits half-rank sublattices on the Haar scale. Even a sharp half-rank estimate would not by itself supply the compatible nested sublattices needed for the best multiblock constructions.

To formulate the tree-size question independently of any particular construction, define

$$C_{\mathrm{wc}} = \limsup_{n\to\infty} \sup_{\det \mathcal{L}=1} \inf_{\mathbf{B} \text{ basis of } \mathcal{L}} \frac{\ln T(\mathbf{B}, \lambda_1(\mathcal{L}))}{n \ln n}. \tag{9.4}$$

Theorem 1.5 and the worst-case HKZ upper bound [28, Theorem 2.5] give $c_{\mathrm{flat}} \le C_{\mathrm{wc}} \le c_{\mathrm{HKZ}} = 1/(2e)$.

Conjecture 1.11 is equivalent to $C_{\mathrm{wc}} < c_{\mathrm{HKZ}} = 1/(2e)$. This is an existence conjecture, without a claim about the cost of finding the basis. The stronger possibility $C_{\mathrm{wc}} = c_{\mathrm{flat}}$ remains an open question. The heuristic discussion of the Hanrot–Stehlé family motivates investigating this direction, but does not settle the general worst-case question.

AI usage

OpenAI's ChatGPT 6 Astra and Anthropic's Claude Opus 5.5 were used in preparing this paper, including literature searches, the development and checking of proof arguments, drafting and editing the text and LaTeX source, generating figures, and repeated review cycles of the paper. The author verified the correctness and is responsible for the mathematical claims, references, and final presentation.

References

1. Aggarwal, D., Espitau, T., Peters, S., Stephens-Davidowitz, N.: Recursive lattice reduction—a framework for finding short lattice vectors. In: Proceedings of the 2025 Symposium on Simplicity in Algorithms (SOSA). pp. 533–573. SIAM (2025). https://doi.org/10.1137/1.9781611978315.40, https://arxiv.org/abs/2311.15064
2. Ajtai, M., Kumar, R., Sivakumar, D.: A sieve algorithm for the shortest lattice vector problem. In: Proceedings of the thirty-third annual ACM symposium on Theory of computing. pp. 601–610. STOC01, ACM (2001). https://doi.org/10.1145/380752.380857
3. Albrecht, M.R., Bai, S., Fouque, P.A., Kirchner, P., Stehlé, D., Wen, W.: Faster Enumeration-Based Lattice Reduction: Root Hermite Factor $k^{1/(2k)}$ in Time $k^{k/8+o(k)}$. In: Advances in Cryptology – CRYPTO 2020, Part II. Lecture Notes in Computer Science, vol. 12171, pp. 186–212. Springer (2020). https://doi.org/10.1007/978-3-030-56880-1_7
4. Albrecht, M.R., Ducas, L., Herold, G., Kirshanova, E., Postlethwaite, E.W., Stevens, M.: The General Sieve Kernel and New Records in Lattice Reduction. In: Advances in Cryptology – EUROCRYPT 2019. pp. 717–746. Springer International Publishing (2019). https://doi.org/10.1007/978-3-030-17656-3_25
5. Aono, Y., Nguyen, P.Q., Seito, T., Shikata, J.: Lower Bounds on Lattice Enumeration with Extreme Pruning. In: Advances in Cryptology – CRYPTO 2018, Part II. Lecture Notes in Computer Science, vol. 10992, pp. 608–637. Springer (2018). https://doi.org/10.1007/978-3-319-96881-0_21
6. Aono, Y., Nguyen, P.Q., Shen, Y.: Quantum lattice enumeration and tweaking discrete pruning. In: Advances in Cryptology – ASIACRYPT 2018, Part I. Lecture Notes in Computer Science, vol. 11272, pp. 405–434. Springer (2018). https://doi.org/10.1007/978-3-030-03326-2_14, https://eprint.iacr.org/2018/546
7. Bai, S., Laarhoven, T., Stehlé, D.: Tuple lattice sieving. LMS Journal of Computation and Mathematics **19**(A), 146–162 (2016). https://doi.org/10.1112/s1461157016000292

8. Banaszczyk, W.: New Bounds in Some Transference Theorems in the Geometry of Numbers. Mathematische Annalen **296**, 625–635 (1993). https://doi.org/10.1007/BF01445125
9. Becker, A., Ducas, L., Gama, N., Laarhoven, T.: New directions in nearest neighbor searching with applications to lattice sieving. In: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 10–24. Society for Industrial and Applied Mathematics (2016). https://doi.org/10.1137/1.9781611974331.ch2
10. Blömer, J.: Closest Vectors, Successive Minima, and Dual HKZ-Bases of Lattices. In: Automata, Languages and Programming. pp. 248–259. Springer Berlin Heidelberg (2000). https://doi.org/10.1007/3-540-45022-x_22
11. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Proceedings of the 45th Annual ACM Symposium on Theory of Computing. pp. 575–584. STOC '13, ACM (2013). https://doi.org/10.1145/2488608.2488680, https://arxiv.org/abs/1306.0281, full version: arXiv:1306.0281, Section 5
12. Brassard, G., Høyer, P., Tapp, A.: Quantum cryptanalysis of hash and claw-free functions. In: LATIN '98: Theoretical Informatics. Lecture Notes in Computer Science, vol. 1380, pp. 163–169. Springer (1998). https://doi.org/10.1007/BFb0054319, https://arxiv.org/abs/quant-ph/9705002
13. Chen, Y., Nguyen, P.Q.: BKZ 2.0: Better Lattice Security Estimates. In: Advances in Cryptology – ASIACRYPT 2011. pp. 1–20. Springer Berlin Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_1
14. Dadush, D., Micciancio, D.: Algorithms for the Densest Sub-Lattice Problem. In: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 1103–1122. Society for Industrial and Applied Mathematics (2013). https://doi.org/10.1137/1.9781611973105.79
15. Dadush, D., Regev, O., Stephens-Davidowitz, N.: On the Closest Vector Problem with a Distance Guarantee. In: Proceedings of the 29th Conference on Computational Complexity. pp. 98–109 (2014), https://arxiv.org/abs/1409.8063
16. Doulgerakis, E., Laarhoven, T., de Weger, B.: Finding Closest Lattice Vectors Using Approximate Voronoi Cells. In: Ding, J., Steinwandt, R. (eds.) Post-Quantum Cryptography – PQCrypto 2019. Lecture Notes in Computer Science, vol. 11505, pp. 3–22. Springer (2019). https://doi.org/10.1007/978-3-030-25510-7_1
17. Doulgerakis, E., Laarhoven, T., de Weger, B.: Sieve, enumerate, slice, and lift: Hybrid lattice algorithms for SVP via CVPP. In: Nitaj, A., Youssef, A. (eds.) Progress in Cryptology – AFRICACRYPT 2020. Lecture Notes in Computer Science, vol. 12174, pp. 301–320. Springer (2020). https://doi.org/10.1007/978-3-030-51938-4_15, https://eprint.iacr.org/2020/487
18. Ducas, L., Loyer, J.: Lattice Reduction via Dense Sublattices: A Cryptanalytic No-Go. IACR Communications in Cryptology **2**(3), 12:1–12:15 (2025). https://doi.org/10.62056/ae0fhsfg
19. Feng, Y., Gao, Y., Liu, J.: Faster SVP in Polynomial Space (2026), https://arxiv.org/abs/2609.21612, arXiv:2609.21612
20. Fincke, U., Pohst, M.: Improved methods for calculating vectors of short length in a lattice, including a complexity analysis. Mathematics of Computation **44**(170), 463–471 (1985). https://doi.org/10.1090/s0025-5718-1985-0777278-8
21. Folland, G.B.: Real Analysis: Modern Techniques and Their Applications. John Wiley & Sons, 2 edn. (1999)
22. Fu, L.: Improvements in the analysis of Kannan's CVP algorithm. Theoretical Computer Science **621**, 103–112 (2016). https://doi.org/10.1016/j.tcs.2016.01.038
23. Gama, N., Howgrave-Graham, N., Koy, H., Nguyen, P.Q.: Rankin's constant and blockwise lattice reduction. In: Advances in Cryptology – CRYPTO 2006. Lecture Notes in Computer Science, vol. 4117, pp. 112–130. Springer (2006). https://doi.org/10.1007/11818175_7
24. Gama, N., Nguyen, P.Q., Regev, O.: Lattice Enumeration Using Extreme Pruning. In: Advances in Cryptology – EUROCRYPT 2010. pp. 257–278. Springer Berlin Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_13
25. Hanrot, G., Pujol, X., Stehlé, D.: Analyzing blockwise lattice algorithms using dynamical systems. In: Rogaway, P. (ed.) Advances in Cryptology – CRYPTO 2011. Lecture Notes in Computer Science, vol. 6841, pp. 447–464. Springer (2011). https://doi.org/10.1007/978-3-642-22792-9_25
26. Hanrot, G., Stehlé, D.: Improved Analysis of Kannan's Shortest Lattice Vector Algorithm. In: Advances in Cryptology - CRYPTO 2007. pp. 170–186. Springer Berlin Heidelberg (2007). https://doi.org/10.1007/978-3-540-74143-5_10
27. Hanrot, G., Stehlé, D.: Worst-Case Hermite–Korkine–Zolotarev Reduced Lattice Bases. Tech. Rep. RR-6422, INRIA (2008), https://arxiv.org/abs/0801.3331
28. Hanrot, G., Stehlé, D.: A Complete Worst-Case Analysis of Kannan's Shortest Lattice Vector Algorithm (2010), https://perso.ens-lyon.fr/damien.stehle/downloads/KANNAN_EXTENDED.pdf, extended manuscript
29. Haque, M.M., Pieprzyk, J.: Analysing recursive preprocessing of BKZ lattice reduction algorithm. IET Information Security **11**(2), 114–120 (2017). https://doi.org/10.1049/iet-ifs.2016.0049
30. Helfrich, B.: Algorithms to Construct Minkowski Reduced and Hermite Reduced Lattice Bases. Theoretical Computer Science **41**(2–3), 125–139 (1985). https://doi.org/10.1016/0304-3975(85)90067-2

31. Hunkenschröder, C., Reuland, G., Schymura, M.: On Compact Representations of Voronoi Cells of Lattices. Mathematical Programming **183**, 337–358 (2020). https://doi.org/10.1007/s10107-019-01463-3, https://arxiv.org/abs/1811.08532
32. Kannan, R.: Improved algorithms for integer programming and related lattice problems. In: Proceedings of the fifteenth annual ACM symposium on Theory of computing - STOC '83. pp. 193–206. STOC '83, ACM Press (1983). https://doi.org/10.1145/800061.808749
33. Kannan, R.: Minkowski's Convex Body Theorem and Integer Programming. Mathematics of Operations Research **12**(3), 415–440 (1987). https://doi.org/10.1287/moor.12.3.415
34. Karenin, A., Kirshanova, E., Nowakowski, J., May, A.: Fast slicer for batch-CVP: Making lattice hybrid attacks practical. In: Hanaoka, G., Yang, B.Y. (eds.) Advances in Cryptology – ASIACRYPT 2025, Part III. Lecture Notes in Computer Science, vol. 16247, pp. 100–132. Springer (2026). https://doi.org/10.1007/978-981-95-5099-9_4, https://eprint.iacr.org/2025/1910
35. Kim, S.: Mean value formulas on sublattices and flags of the random lattice. Journal of Number Theory **241**, 330–351 (2022). https://doi.org/10.1016/j.jnt.2022.03.013
36. Klein, P.N.: Finding the closest lattice vector when it's unusually close. In: Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 937–941. Society for Industrial and Applied Mathematics (2000), https://dl.acm.org/doi/10.5555/338219.338661
37. Korkine, A., Zolotareff, G.: Sur les formes quadratiques. Mathematische Annalen **6**(3), 366–389 (1873). https://doi.org/10.1007/BF01442795
38. Laarhoven, T.: Approximate Voronoi cells for lattices, revisited. Journal of Mathematical Cryptology **15**(1), 60–71 (2021). https://doi.org/10.1515/jmc-2020-0074
39. Lenstra, A.K., Lenstra, H.W., Lovász, L.: Factoring polynomials with rational coefficients. Mathematische Annalen **261**(4), 515–534 (1982). https://doi.org/10.1007/bf01457454
40. Li, J., Nguyen, P.Q.: Approximating the densest sublattice from Rankin's inequality. LMS Journal of Computation and Mathematics **17**(A), 92–111 (2014). https://doi.org/10.1112/s1461157014000333
41. Li, J., Nguyen, P.Q.: A complete analysis of the BKZ lattice reduction algorithm. Journal of Cryptology **38** (2025). https://doi.org/10.1007/s00145-024-09527-0, article 12
42. Meyer, B.: Generalised Hermite constants, Voronoi theory and heights on flag varieties. Bulletin de la Société Mathématique de France **137**(1), 127–158 (2009). https://doi.org/10.24033/bsmf.2571
43. Montanaro, A.: Quantum-walk speedup of backtracking algorithms. Theory of Computing **14**(15), 1–24 (2018). https://doi.org/10.4086/toc.2018.v014a015, https://theoryofcomputing.org/articles/v014a015/
44. Neumaier, A.: Bounding basis reduction properties. Designs, Codes and Cryptography **84**, 237–259 (2017). https://doi.org/10.1007/s10623-016-0273-9
45. Nguyen, P.Q.: Hermite's Constant and Lattice Algorithms. In: The LLL Algorithm, pp. 19–69. Springer Berlin Heidelberg (2009). https://doi.org/10.1007/978-3-642-02295-1_2
46. Pellet-Mary, A., Hanrot, G., Stehlé, D.: Approx-SVP in ideal lattices with pre-processing. In: Ishai, Y., Rijmen, V. (eds.) Advances in Cryptology – EUROCRYPT 2019, Part II. Lecture Notes in Computer Science, vol. 11477, pp. 685–716. Springer (2019). https://doi.org/10.1007/978-3-030-17656-3_24, https://eprint.iacr.org/2019/215
47. Schneider, M., Gama, N.: SVP Challenge (2010), https://www.latticechallenge.org/svp-challenge/, benchmark website; accessed 30 September 2026
48. Schnorr, C.P., Euchner, M.: Lattice basis reduction: Improved practical algorithms and solving subset sum problems. Mathematical Programming **66**(1-3), 181–199 (1994). https://doi.org/10.1007/bf01581144
49. Schnorr, C.P.: Lattice Reduction by Random Sampling and Birthday Methods. In: STACS 2003. pp. 145–156. Springer Berlin Heidelberg (2003). https://doi.org/10.1007/3-540-36494-3_14
50. Shapira, U., Weiss, B.: A volume estimate for the set of stable lattices. Comptes Rendus Mathématique **352**(11), 875–879 (2014). https://doi.org/10.1016/j.crma.2014.08.019
51. Siegel, C.L.: A Mean Value Theorem in Geometry of Numbers. The Annals of Mathematics **46**(2), 340–347 (1945). https://doi.org/10.2307/1969027
52. Södergren, A.: On the Poisson distribution of lengths of lattice vectors in a random lattice. Mathematische Zeitschrift **269**(3-4), 945–954 (2011). https://doi.org/10.1007/s00209-010-0772-8
53. Strömbergsson, A.: On the probability of a random lattice avoiding a large convex set. Proceedings of the London Mathematical Society **103**(6), 950–1006 (2011). https://doi.org/10.1112/plms/pdr021
54. Thunder, J.L.: Higher-dimensional analogs of Hermite's constant. Michigan Mathematical Journal **45**(2), 301–314 (1998). https://doi.org/10.1307/mmj/1030132184
55. Wagner, U., Maze, G.: Improvements in Closest Point Search Based on Dual HKZ-Bases. Theoretical Computer Science **520**, 62–74 (2014). https://doi.org/10.1016/j.tcs.2013.10.005, https://arxiv.org/abs/1201.5273
56. Yu, Y., Ducas, L.: Second order statistical behavior of LLL and BKZ. In: Selected Areas in Cryptography – SAC 2017. Lecture Notes in Computer Science, vol. 10719, pp. 3–22. Springer (2018). https://doi.org/10.1007/978-3-319-72565-9_1