

High-dimensional sphere packing and the modular bootstrap

Nima Afkhami-Jeddi,^a Henry Cohn,^b Thomas Hartman,^c David de Laat^d
and Amirhossein Tajdini^c

^a*Enrico Fermi Institute & Kadanoff Center for Theoretical Physics, University of Chicago, Chicago, Illinois, U.S.A.*

^b*Microsoft Research New England, Cambridge, Massachusetts, U.S.A.*

^c*Department of Physics, Cornell University, Ithaca, New York, U.S.A.*

^d*Delft Institute of Applied Mathematics, Delft University of Technology, Delft, The Netherlands*

E-mail: nimaaj@uchicago.edu, cohn@microsoft.com, hartman@cornell.edu, d.delaat@tudelft.nl, at734@cornell.edu

ABSTRACT: We carry out a numerical study of the spinless modular bootstrap for conformal field theories with current algebra $U(1)^c \times U(1)^c$, or equivalently the linear programming bound for sphere packing in $2c$ dimensions. We give a more detailed picture of the behavior for finite c than was previously available, and we extrapolate as $c \rightarrow \infty$. Our extrapolation indicates an exponential improvement for sphere packing density bounds in high dimensions. Furthermore, we study when these bounds can be tight. Besides the known cases $c = 1/2, 4$, and 12 and the conjectured case $c = 1$, our calculations numerically rule out sharp bounds for all other $c < 90$, by combining the modular bootstrap with linear programming bounds for spherical codes.

KEYWORDS: Conformal Field Theory, Field Theories in Lower Dimensions

ARXIV EPRINT: [2006.02560](https://arxiv.org/abs/2006.02560)

Contents

1	Introduction	1
1.1	Results from the spinless modular bootstrap for large c	2
1.2	New constraints on tight sphere packing bounds	4
2	The spinless modular bootstrap	5
2.1	Setting up the bootstrap	5
2.2	Uncertainty principle	7
2.3	Sphere packing	9
2.4	Numerics	10
3	Numerical results	12
3.1	Data and plots	12
3.2	Extrapolation	14
4	Properties of the spectrum and degeneracies	16
4.1	Convergence to the free fermion spectrum	17
4.2	Deviations from the free fermion spectrum	20
4.3	Growth rate of the degeneracies	22
4.4	Degeneracies for $c = 4$ and 12	23
4.5	Degeneracies for arbitrary c	24
4.6	Possible explanation of the high-energy spectrum	25
5	Spherical codes and implied kissing numbers	26
5.1	The linear programming bound	29
5.2	Levenshtein’s universal bound	31
5.3	The Kabatyanskii-Levenshtein bound	32
5.4	Implied kissing numbers	33
5.5	Bounds for the average kissing number	35
A	The limiting case of the spinless modular bootstrap	36
B	Convergence of the spinless bootstrap	37

1 Introduction

To what extent do self-consistency principles constrain, or even determine, the behavior of a system? This question underlies many topics in mathematics and physics. One notable example is the conformal bootstrap program [1–4] (see [5–7] for reviews), which seeks to map the space of possible conformal field theories (CFTs) and to identify those on the boundary of theory space (i.e., those with extremal properties, almost but not quite inconsistent). A more down to earth example is the sphere packing problem, in which the goal is to maximize the fraction of \mathbb{R}^d covered by congruent spheres whose interiors are not allowed to overlap. In low dimensions it is not hard to guess the optimal packings, but even

that remains mysterious in high dimensions. Proving upper bounds for the packing density is particularly difficult, and in most cases the best bounds currently known are obtained via the linear programming bound of Cohn and Elkies [8], which relies on harmonic analysis.

While these problems sound completely unrelated, Hartman, Mazáč, and Rastelli [9] discovered a surprising connection between them: the spinless modular bootstrap for two-dimensional CFTs is very nearly the same as the linear programming bound for sphere packing. The underlying optimization problems are exactly equivalent when the current algebra $U(1)_{\text{left}}^c \times U(1)_{\text{right}}^{\bar{c}}$ with total central charge $c_{\text{total}} = c + \bar{c}$ acts on the CFT and the sphere packing dimension is given by $d = c_{\text{total}}$, and they are closely related (but not equivalent) under the Virasoro algebra. The relationship between the modular bootstrap and linear programming bounds seems to be specific to these particular techniques, rather than being based on a direct connection between CFTs and sphere packings.

In this paper we will focus on the spinless modular bootstrap for $U(1)_{\text{left}}^c \times U(1)_{\text{right}}^{\bar{c}}$ with c_{total} large. The analysis depends only on c_{total} , not on the left and right central charges individually. To simplify the notation we set $\bar{c} = c$ and refer simply to the $U(1)^c$ modular bootstrap, parameterizing our results by $c = c_{\text{total}}/2$.

Neither the modular bootstrap nor the linear programming bound has been completely analyzed, either theoretically or numerically. Each depends on producing some additional information (namely, an auxiliary function or linear functional satisfying certain inequalities), which must be chosen carefully to optimize the resulting bound, and this optimization has proved difficult. The equivalence between these problems adds to the motivation for studying them, because any consequences will shed light on two seemingly disparate topics. A third application is to generalizations of the Bourgain-Clozel-Kahane uncertainty principle for signs of functions [10, 11]. Thus, these problems live at a particularly fruitful intersection of several fields.

In this paper, we carry out the first large-scale numerical study of the $U(1)^c$ spinless modular bootstrap with c large, or equivalently the linear programming bound on sphere packing in high dimensions, by adapting the numerical techniques introduced by Afkhami-Jeddi, Hartman, and Tajdini for the Virasoro case [12]. These techniques closely parallel the approach independently taken by Cohn, Elkies, Kumar, and Gonçalves [8, 11, 13, 14] in the sphere packing literature, but the paper [12] introduced better extrapolation techniques and achieved superior performance.

In CFT terms, the spinless modular bootstrap corresponds to constraints on the partition function at zero angular potential. A natural question is whether the spinning modular bootstrap, i.e., including an angular potential, also bounds the density of general sphere packings. The answer is that it does not, as this would contradict known packings. The spinning bootstrap analysis for CFTs with $U(1)_{\text{left}}^c \times U(1)_{\text{right}}^c$ symmetry has interesting implications for holographic duality and will appear in a separate paper [15].

1.1 Results from the spinless modular bootstrap for large c

For sphere packing in high dimensions, the central question is the asymptotic behavior of the packing density. It is at least 2^{-d} in \mathbb{R}^d , with only much lower-order improvements known [16–18], and it is at most $2^{-(\kappa+o(1))d}$ with $\kappa = 0.59905576\dots$. The latter bound was

found by Kabatyanskii and Levenshtein [19] in 1978, and the exponential decay rate has not been improved since then. Cohn and Zhao [20] showed how to obtain it via the linear programming bound, and a fundamental open question is whether the linear programming bound is capable of improving on this decay rate.

In terms of the $U(1)^c$ spinless modular bootstrap, bounding the packing density amounts to bounding the spectral gap of the CFT. Specifically, the Kabatyanskii-Levenshtein bound says that the scaling dimension of the lowest non-vacuum primary is at most $c/(K + o(1))$ as $c \rightarrow \infty$, where $K = e\pi 2^{2\kappa-1} = 9.79674646\dots$. No better bound is known for the spectral gap.

One of our primary results in this paper is a numerical estimate of the fully optimized $U(1)^c$ spinless modular bootstrap bound for the spectral gap (Conjecture 3.1). In sphere packing terms, it amounts to an upper bound of $2^{-(\lambda+o(1))d}$ for the sphere packing density in \mathbb{R}^d as $d \rightarrow \infty$ with $\lambda \approx 0.6044$; in modular bootstrap terms, it amounts to an upper bound of $c/(\Lambda + o(1))$ for the spectral gap as $c \rightarrow \infty$ with $\Lambda \approx 9.869$. This bound is based on numerical extrapolation, with no proof or even heuristic derivation, but we give a careful accounting of the potential error from the extrapolation. We furthermore guess that the exact value of Λ is π^2 (Conjecture 3.2), although that conjecture is much more speculative.

Conceptually, what our computations indicate is that the Kabatyanskii-Levenshtein upper bound can be decreased by an exponential factor through optimizing the linear programming bound. If proved, this bound would settle a longstanding open problem in discrete geometry. However, the improvement in the decay rate will be small.

The analytical [9] and numerical [12] results for Virasoro symmetry are quite a bit further away from each other (the analytical bound for the spectral gap is $c/8.503$, while the numerical bound is $c/9.08$). We have no conceptual explanation for why the Kabatyanskii-Levenshtein bound should come rather close to optimizing the $U(1)^c$ case, yet fall slightly short. Perhaps generalizing this bound will offer new techniques for optimizing the modular bootstrap more broadly, but we do not expect that it will lead to an exact solution without some new idea.

Sphere packings are error-correcting codes for a continuous communication channel, and they therefore play an important role in information theory. Their discrete counterpart is error-correcting codes for a binary channel, and these two theories are in many ways closely analogous [21], with substantial interplay between them, both in results and in techniques. The linear program bound originated in the discrete setting, in a fundamental paper by Delsarte [22], before being generalized to sphere packing by Cohn and Elkies [8], and the Kabatyanskii-Levenshtein bound was inspired by the MRRW bound, due to McEliece, Rodemich, Rumsey, and Welch [23].

Much like the case of sphere packing, the asymptotic rate in the MRRW bound has not been beaten by any method, and it is an open problem whether it optimizes the linear programming bound. Barg and Jaffe [24] examined this issue numerically, and they conjectured that it is the optimal rate in the linear programming bound. Their conjecture is widely believed, but the evidence is not conclusive. While our results have no direct implications for binary error-correcting codes, they suggest that the MRRW bound may not be optimal, because it is the discrete analogue of the Kabatyanskii-Levenshtein bound.

It would be valuable to perform a more extensive study than Barg and Jaffe were able to do in 2001, as well as to compare the data with section 3.2 of [25].

At the optimum, the linear programming approach provides not just a bound on the spectral gap, but a candidate spectrum for a CFT that saturates it. In sphere packing terms, this spectrum amounts to the pair correlation function of the packing. We study the spectrum numerically in section 4 and find some intriguing structure. For computational purposes, the infinite set of bootstrap constraints is truncated to a finite system of $2N$ equations, with N taken as large as possible. The corresponding spectrum has N states other than the vacuum, with scaling dimensions $\Delta_1 < \Delta_2 < \dots < \Delta_N$. We conjecture a formula for the ratio Δ_n/N in the limit $N \rightarrow \infty$ with n/N held fixed. The formula is piecewise smooth, with an abrupt transition from linear to nonlinear behavior at $n \sim (2/\pi)N$. We have no analytic explanation for this transition. The linear portion of the spectrum matches the spectrum of the generalized free fermion in one dimension, which was used to construct analytic functionals for CFT in [26] and adapted to sphere packing in [9].

1.2 New constraints on tight sphere packing bounds

In addition to studying the asymptotic behavior of the modular bootstrap, we also search for exceptional behavior at finite central charge. Four particular values are known to play a special role, namely $c = 1/2, 1, 4,$ and 12 . In sphere packing terms, these cases correspond to exact solutions of the sphere packing problem in dimensions 1, 2, 8, and 24. While $d = 1$ is trivial, $d = 8$ and $d = 24$ are far deeper, and they are the only cases in which the sphere packing problem has been solved above $d = 3$ (which was solved by Hales [27, 28], with no connection to the modular bootstrap). Dimension 8 was a breakthrough due to Viazovska [29], and dimension 24 built on her techniques [30]. The linear programming bound seems to be exact when $d = 2$ as well, but no proof is known, although the two-dimensional sphere packing problem can be solved directly [31, 32].

These cases are more subtle for CFTs than they are for sphere packings. For $(c, \bar{c}) = (4, 4)$, there is indeed a CFT invariant under $U(1)_{\text{left}}^c \times U(1)_{\text{right}}^{\bar{c}}$ that achieves the spinless modular bootstrap bound, namely eight free fermions [33]. No such CFT exists for $(c, \bar{c}) = (12, 12)$ (see [15]), but there is a chiral CFT with $(c, \bar{c}) = (24, 0)$, namely the 24 chiral bosons compactified using the quotient of \mathbb{R}^{24} by the Leech lattice.¹ The case $c = 1/2$ is not an integer, so $U(1)_{\text{left}}^c \times U(1)_{\text{right}}^{\bar{c}}$ symmetry does not even make sense, but again we can use a chiral boson with $(c, \bar{c}) = (1, 0)$. This time, however, it is not fully conformally invariant. Instead, it has a nontrivial phase under the action of the generator T of $SL_2(\mathbb{Z})$, but the spinless modular bootstrap nevertheless applies. Finally, in the case $c = 1$ no CFT invariant under $U(1)_{\text{left}}^c \times U(1)_{\text{right}}^{\bar{c}}$ achieves the spinless bound (see [15]), but it is achieved by two chiral bosons with $(c, \bar{c}) = (2, 0)$ and a nontrivial phase under the T transformation. Thus, the CFT picture encompasses all four exceptional cases, provided we allow chiral CFTs and are willing to relax conformal invariance.

¹Note that we do not take an orbifold quotient, as in the Monster CFT of Frenkel, Lepowsky, and Meurman [34], because we want current algebra $U(1)^{24}$. Similarly, 8 chiral bosons compactified using E_8 meet the bound with $(c, \bar{c}) = (8, 0)$, but the $(c, \bar{c}) = (4, 4)$ case is more noteworthy.

Why should the exceptional solutions of the sphere packing problem be limited to these specific dimensions? It comes as no surprise to see sporadic behavior tied to E_8 and the Leech lattice (for $c = 4$ and 12 , respectively), but it is difficult to explain why this behavior is not more widespread. For a provocative example, why shouldn't the linear programming bound solve the sphere packing problem in all sufficiently large dimensions? We do not know how to rule out this possibility, although it is utterly implausible.

To shed light on this problem, we examine the conditions that would have to hold to obtain a sphere packing meeting the linear programming bound. To do so, we incorporate additional constraints beyond the modular invariance of the partition function. Specifically, we study the *implied kissing number*, the average number of tangencies between spheres in a hypothetical packing with this property. In all dimensions up through $d = 250$ other than $1, 2, 8, 24, 180, 181,$ and 192 , we show that the implied kissing number from our numerically optimized bound is impossibly large. Thus, no sphere packing can attain the exact linear programming bound in these dimensions.² We do not expect optimal solutions in dimensions $180, 181,$ or 192 , and we see no sign of them, but our bounds do not rule them out.

As the unexpected occurrence of dimensions such as 181 indicates, this problem has a surprisingly intricate structure. While certain aspects behave in straightforward ways that are not hard to extrapolate, other aspects are far more subtle. One feature of our numerical solutions for which we have no conceptual explanation is a kind of periodicity: the degeneracies are especially well described by a Cardy-like entropy formula when c is a multiple of 8 , and the scaling dimensions are especially close to those for generalized one-dimensional free fermions when c is 4 more than a multiple of 8 . In other words, multiples of 4 behave particularly well, but no value of c looks equally simple from all perspectives. The reason for this behavior remains mysterious.

2 The spinless modular bootstrap

2.1 Setting up the bootstrap

In this section, we will briefly review the spinless modular bootstrap, which is a technique for proving bounds on the possible scaling dimensions of primary fields in a compact, unitary 2d CFT [33, 35, 36]. Given such a CFT, let $Z(\tau, \bar{\tau})$ be its partition function, i.e., the sum over all states of $q^{h-c/24} \bar{q}^{\bar{h}-\bar{c}/24}$, where h and \bar{h} are the conformal weights of the state, c and \bar{c} are the left and right central charges, and $q = e^{2\pi i\tau}$ and $\bar{q} = e^{-2\pi i\bar{\tau}}$ (with τ and $-\bar{\tau}$ in the upper half-plane).³ Because of conformal invariance, the partition function satisfies modular invariance:

$$Z\left(\frac{a\tau + b}{c\tau + d}, \frac{a\bar{\tau} + b}{c\bar{\tau} + d}\right) = Z(\tau, \bar{\tau}) \tag{2.1}$$

²Strictly speaking, our work does not amount to a proof, because it leaves open the possibility that we have not fully optimized the linear programming bound. However, we present strong numerical evidence that we have optimized it.

³Mathematicians should note that the bars do not denote complex conjugates.

whenever

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{SL}_2(\mathbb{Z}) \tag{2.2}$$

(where in these formulas c is of course not necessarily the left central charge). In terms of the usual generators

$$S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \quad \text{and} \quad T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \tag{2.3}$$

of $\text{SL}_2(\mathbb{Z})$, modular invariance amounts to

$$Z(-1/\tau, -1/\bar{\tau}) = Z(\tau + 1, \bar{\tau} + 1) = Z(\tau, \bar{\tau}). \tag{2.4}$$

For the spinless modular bootstrap, we specialize the partition function to have zero angular potential. In other words, we set $\bar{\tau} = -\tau$ (i.e., $\bar{q} = q$) and use the restricted partition function

$$\mathcal{Z}(\tau) = Z(\tau, -\tau). \tag{2.5}$$

The action of S on τ and $\bar{\tau}$ preserves the condition $\bar{\tau} = -\tau$, and thus

$$\mathcal{Z}(-1/\tau) = \mathcal{Z}(\tau), \tag{2.6}$$

but the action of T does not. Thus, we expect that usually $\mathcal{Z}(\tau + 1) \neq \mathcal{Z}(\tau)$.

The spinless modular bootstrap is based on the identity $\mathcal{Z}(-1/\tau) = \mathcal{Z}(\tau)$. Because we make no use of the action of T , the bound applies even to theories that are invariant only under S . A simple example is a single chiral boson at the self-dual radius, for which $Z(\tau, \bar{\tau}) = \theta_3(\tau)/\eta(\tau)$ in terms of the Jacobi theta function and Dedekind eta function. Such theories are not fully conformally invariant, but the spinless modular bootstrap still applies.

The combined contribution of the descendants of a primary field of scaling dimension $\Delta = h + \bar{h}$ to the partition function $\mathcal{Z}(\tau)$ is a character $\chi_\Delta(\tau)$ of a Verma module of the current algebra, and thus the partition function is given by a sum

$$\mathcal{Z}(\tau) = \sum_{\Delta} d_{\Delta} \chi_{\Delta}(\tau) \tag{2.7}$$

over the scaling dimensions of the primary fields, each with multiplicity given by the degeneracy d_{Δ} . The vacuum corresponds to $\Delta = 0$, with degeneracy $d_0 = 1$, and the other scaling dimensions are positive numbers $\Delta_1 < \Delta_2 < \dots$ that tend to infinity.

The precise form of the characters depends on the current algebra. Our main interest in this paper will be the algebra $U(1)_{\text{left}}^c \times U(1)_{\text{right}}^{\bar{c}}$ (more precisely, the corresponding affine Lie algebra), in which case

$$\chi_{\Delta}(\tau) = \frac{e^{2\pi i \tau \Delta}}{\eta(\tau)^{c+\bar{c}}}, \tag{2.8}$$

where η is again the Dedekind eta function. In particular, the only dependence on the central charges is through their sum $c + \bar{c}$.

The spectral gap of the CFT is the lowest scaling dimension Δ_1 of a primary other than the vacuum. We can obtain an upper bound for the spectral gap by producing a

linear functional that acts in a certain way on functions of τ . The key observation is that if we set

$$\Phi_{\Delta}(\tau) = \chi_{\Delta}(\tau) - \chi_{\Delta}(-1/\tau), \tag{2.9}$$

then we obtain the crossing equation

$$\sum_{\Delta} d_{\Delta} \Phi_{\Delta}(\tau) = 0 \tag{2.10}$$

by modular invariance. Now suppose ω is a linear functional such that

$$\omega(\Phi_0) > 0 \tag{2.11}$$

and

$$\omega(\Phi_{\Delta}) \geq 0 \tag{2.12}$$

whenever $\Delta \geq \Delta_{\text{gap}}$ for some constant Δ_{gap} . If we apply ω to the crossing equation, we find that

$$\omega(\Phi_0) + \sum_{\Delta > 0} d_{\Delta} \omega(\Phi_{\Delta}) = 0, \tag{2.13}$$

which would be impossible if all the non-zero scaling dimensions were at least Δ_{gap} , because the total would be positive. Thus, we conclude that some primary must have a scaling dimension strictly between 0 and Δ_{gap} . In other words, Δ_{gap} is a strict upper bound for the spectral gap. One can show that it is a weak upper bound even if $\omega(\Phi_0) = 0$, as long as $\omega(\Phi_{\Delta})$ is not identically zero (see appendix A), and that $\omega(\Phi_0) = 0$ must hold for the optimal choice of ω .

The optimal functional ω is not known, except in a handful of special cases discussed below. In sections 3 and 4, we will give the most detailed numerical study so far of how ω and Δ_{gap} behave.

As noted earlier, because the spinless modular bootstrap for $U(1)_{\text{left}}^c \times U(1)_{\text{right}}^{\bar{c}}$ depends only on $c + \bar{c}$, we will set $\bar{c} = c$ and refer just to c . Strictly speaking this notation is misleading when $c + \bar{c}$ is odd, because the current algebra $U(1)_{\text{left}}^c \times U(1)_{\text{right}}^{\bar{c}}$ makes sense only when c and \bar{c} are nonnegative integers. For example, the only physically meaningful cases with $c + \bar{c} = 1$ are $(c, \bar{c}) = (1, 0)$ or $(0, 1)$, and they are therefore what we mean when we refer to the $c = 1/2$ case. More generally, the abstract problem of optimizing the bound makes sense for any $c > 0$, but there are consequences for CFTs only when c is an integer or half-integer.

2.2 Uncertainty principle

Hartman, Mazáč, and Rastelli [9] reformulated the $U(1)^c$ spinless modular bootstrap in terms of an uncertainty principle for eigenfunctions of the Fourier transform as follows. Suppose $d = 2c$ is an integer, which is the meaningful case for CFTs. Given a functional ω as above, we define a radial function $f_{\omega}: \mathbb{R}^d \rightarrow \mathbb{R}$ by $f_{\omega}(x) = \omega(\Phi_{|x|^2/2})$. If we normalize the Fourier transform in \mathbb{R}^d by

$$\hat{f}(y) = \int_{\mathbb{R}^d} dx f(x) e^{-2\pi i \langle x, y \rangle}, \tag{2.14}$$

then f_ω is an eigenfunction of the Fourier transform with eigenvalue -1 ; in other words, $\widehat{f_\omega} = -f_\omega$. To see why, we start with

$$\begin{aligned} \Phi_{|x|^2/2}(\tau) &= \chi_{|x|^2/2}(\tau) - \chi_{|x|^2/2}(-1/\tau) \\ &= \frac{e^{\pi i \tau |x|^2}}{\eta(\tau)^d} - \frac{e^{\pi i (-1/\tau) |x|^2}}{\eta(-1/\tau)^d} \\ &= \frac{e^{\pi i \tau |x|^2} - (i/\tau)^{d/2} e^{\pi i (-1/\tau) |x|^2}}{\eta(\tau)^d}, \end{aligned} \tag{2.15}$$

because the Dedekind eta function satisfies the identity $\eta(-1/\tau) = (\tau/i)^{1/2} \eta(\tau)$. The complex Gaussian $x \mapsto e^{\pi i \tau |x|^2}$ on \mathbb{R}^d has Fourier transform $y \mapsto (i/\tau)^{d/2} e^{\pi i (-1/\tau) |y|^2}$. Thus, the function $x \mapsto e^{\pi i \tau |x|^2} - (i/\tau)^{d/2} e^{\pi i (-1/\tau) |x|^2}$ in the numerator of $\Phi_{|x|^2/2}(\tau)$ is a -1 eigenfunction of the Fourier transform for each τ , because it is the difference of a Gaussian and its Fourier transform. We conclude that f_ω also satisfies $\widehat{f_\omega} = -f_\omega$, by the linearity of ω . The same holds even when $2c$ is not an integer, if we interpret the radial Fourier transform in non-integral dimension as a Hankel transform.

Conversely, every radial -1 eigenfunction of the Fourier transform in \mathbb{R}^d arises as f_ω for some ω , which we can obtain as follows by constructing a basis. If we let

$$\omega_k = \left. \frac{\partial^k}{\partial \tau^k} \right|_{\tau=i}, \tag{2.16}$$

then $\omega_k(\Phi_\Delta)$ is the product of $e^{-2\pi\Delta}$ with a polynomial in Δ of degree at most k , in which the coefficient of Δ^k is

$$\left. \frac{(2\pi i)^k - (i/\tau)^{d/2} (2\pi i)^k (1/\tau^2)^k}{\eta(\tau)^d} \right|_{\tau=i} = \begin{cases} 0 & \text{if } k \text{ is even, and} \\ 2(2\pi i)^k / \eta(i)^d & \text{if } k \text{ is odd.} \end{cases} \tag{2.17}$$

For comparison, the Laguerre polynomials $L_k^{(d/2-1)}$ give a basis for radial functions on \mathbb{R}^d as $x \mapsto L_k^{(d/2-1)}(2\pi|x|^2)e^{-\pi|x|^2}$, with eigenvalues $(-1)^k$ under the Fourier transform. We conclude that the functions $x \mapsto \omega_k(\Phi_{|x|^2/2})$ with $k = 1, 3, 5, \dots, 2m-1$ must span the same space as the Laguerre eigenfunctions with these values of k , and thus they span the entire -1 eigenspace as $m \rightarrow \infty$.

We have seen that choosing the linear functional ω in the $U(1)^c$ spinless modular bootstrap amounts to choosing an integrable, radial function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ with $\widehat{f} = -f$ such that f is not identically zero. The constraints on ω say that $f(0) \geq 0$ and $f(x) \geq 0$ whenever $|x| \geq r$ for some radius r . Then $\Delta_{\text{gap}} = r^2/2$, and optimizing the bound means minimizing r . This optimization problem for signs of eigenfunctions was first studied by Cohn and Elkies [8], and it was placed in the context of more general uncertainty principles for signs of functions by Cohn and Gonçalves [11].

The corresponding problem for $+1$ eigenfunctions asks for an integrable, radial function $g: \mathbb{R}^d \rightarrow \mathbb{R}$ with $\widehat{g} = g$ such that $g(0) \leq 0$ and $g(x) \geq 0$ for $|x| \geq r$. Again the goal is to minimize r , without letting g vanish identically. This problem does not arise in the spinless modular bootstrap as set up above, but it would apply if the partition function satisfied

$\mathcal{Z}(-1/\tau) = -\mathcal{Z}(\tau)$ and $d_\Delta < 0$ for $\Delta > 0$ (see section 2.1 of [11]). It behaves much like the -1 case. For example, Cohn and Gonçalves [11] obtained an exact solution of the $+1$ problem for $c = 6$, which is analogous to the solutions of the -1 problem with $c = 4$ or 12 . The partition function in this case is given by $\mathcal{Z}(\tau) = \sqrt{j(\tau) - 1728}$, which also arises as the Norton series for a certain pair of elements in the Monster group (see equation (7.3.4c) in [37, p. 425]). Although we do not have a direct physical interpretation for this problem, generalized modular transformations do arise in theories with discrete anomalies [38] or fermions [39, 40], including sectors that obey $\mathcal{Z}(-1/\tau) = -\mathcal{Z}(\tau)$.

2.3 Sphere packing

The -1 eigenfunction uncertainty principle plays a fundamental role in discrete geometry, where it underlies the linear programming bound for the sphere packing density. Linear programming bounds are a powerful technique for proving upper bounds for packing density or error-correcting code rates. They were introduced for discrete error-correcting codes by Delsarte [22] in 1972, and extended to sphere packing in Euclidean space by Cohn and Elkies [8] in 2003. The connection with the spinless modular bootstrap for $U(1)^c$ was derived by Hartman, Mazáč, and Rastelli [9] in 2019.

The linear programming bound for sphere packing in \mathbb{R}^d converts an auxiliary function satisfying certain inequalities into a sphere packing density bound, as follows.⁴

Theorem 2.1 (Cohn and Elkies [8]). *Let $h: \mathbb{R}^d \rightarrow \mathbb{R}$ be an integrable, continuous, radial function such that \hat{h} is integrable, and let r be a positive real number. If $h(0) = \hat{h}(0) = 1$, $h(x) \leq 0$ whenever $|x| \geq r$, and $\hat{h}(y) \geq 0$ for all y , then every sphere packing in \mathbb{R}^d has density at most the volume of a sphere of radius $r/2$ in \mathbb{R}^d , i.e.,*

$$\frac{\pi^{d/2}}{(d/2)!} \left(\frac{r}{2}\right)^d. \tag{2.18}$$

The problem of choosing h so as to minimize r is clearly reminiscent of the -1 eigenfunction uncertainty principle, but not obviously equivalent to it. One direction is simple: if h satisfies the hypotheses of Theorem 2.1, then the function $f = \hat{h} - h$ satisfies $f(0) = 0$, $\hat{f} = -f$, and $f(x) \geq 0$ for $|x| \geq r$. Conversely, Cohn and Elkies conjectured that an optimal solution f of the -1 eigenfunction problem can always be lifted to a function h for use in Theorem 2.1 with the same value of r , such that $\hat{h} - h = f$. In other words, the linear programming bound for sphere packing should be identical to the spinless modular bootstrap. No proof is known, but no counterexample has been found, either numerically or analytically.

At first glance, it is not obvious that any auxiliary function satisfies the hypotheses of the linear programming bound. For a first example, let $\chi: \mathbb{R}^d \rightarrow \mathbb{R}$ be the characteristic function of a ball $B_{r/2}$ centered at the origin, with its radius $r/2$ chosen so that $\text{vol}(B_{r/2}) = 1$. Then the convolution $h := \chi * \chi$ has Fourier transform $\hat{h} = \hat{\chi}^2$. By construction, $h(x) = 0$ for $|x| \geq r$ and $\hat{h}(y) \geq 0$ for all y ; furthermore, $h(0) = \text{vol}(B_{r/2}) = 1$ and $\hat{h}(0) = \text{vol}(B_{r/2})^2 = 1$. Thus, the sphere packing density in \mathbb{R}^d is at most $\text{vol}(B_{r/2}) = 1$.

⁴The original technical hypotheses in [8] were slightly stronger; see [20] and [41].

This bound is sharp when $d = 1$, but it is of course not an exciting packing bound. For $d > 1$ the linear programming bound is much better than this first attempt.

In fact, it is the best upper bound known for the sphere packing density in high dimensions [20], but it is generally far from a tight bound [41]. Only four cases seem to be sharp: $d = 1$ (as shown above), 2, 8, and 24. The case $d = 8$ was a breakthrough due to Viazovska [29], and the case $d = 24$ extended her techniques [30]. These are the only cases in which the sphere packing problem has been solved above three dimensions. The optimal auxiliary functions for $d = 8$ and 24 can also be derived from analytic functionals constructed that same year by Mazáč [26] in the four-point function bootstrap for 1d CFTs, as shown by Hartman, Mazáč, and Rastelli [9]. Remarkably, the case $d = 2$ remains unsolved analytically. There is no doubt that it matches the two-dimensional packing density,⁵ but no proof is known.

Linear programming bounds can be applied not just to sphere packing, but to understand ground states under pair potential functions more broadly [42–44]. We will not address that topic in this paper, except to note that our numerical results seem consistent with Conjecture 7.2 in [44], which says that the linear programming bound for sphere packing extends to the Gaussian core model and thereby proves a form of universal optimality, despite the failure of the analogous statement for binary error-correcting codes [45].

2.4 Numerics

To obtain numerical bounds for the $U(1)^c$ spinless modular bootstrap, we must choose a finite-dimensional space of functionals ω . We truncate at derivative order $4N - 1$; in other words, ω will be a linear combination of $\omega_1, \omega_3, \dots, \omega_{4N-1}$, where as above $\omega_k = \partial^k / \partial \tau^k \Big|_{\tau=i}$.

For convenience let $f(\Delta) = \omega(\Phi_\Delta)$, which differs from the -1 Fourier eigenfunction in being a function of Δ rather than x with $\Delta = |x|^2/2$. Then $f(\Delta)$ can be written in terms of the Laguerre eigenfunctions as

$$f(\Delta) = \sum_{j=1}^{2N} \alpha_j f_j(\Delta), \tag{2.19}$$

where

$$f_j(\Delta) = L_{2j-1}^{(c-1)}(4\pi\Delta) e^{-2\pi\Delta}. \tag{2.20}$$

For fixed Δ_{gap} and N , the question is whether f can be chosen to satisfy the positivity conditions $f(0) \geq 0$ and $f(\Delta) \geq 0$ for $\Delta \geq \Delta_{\text{gap}}$ without vanishing identically. This problem can be solved using semidefinite programming, or approximated using linear programming.

Let $\Delta_1^{\text{LP},N}(c)$ be the best bound that can be obtained for a fixed truncation order N , and let $\Delta_1^{\text{LP}}(c)$ be the best bound without restriction on ω . Increasing N improves the bound, and we expect that

$$\Delta_1^{\text{LP}}(c) = \lim_{N \rightarrow \infty} \Delta_1^{\text{LP},N}(c). \tag{2.21}$$

⁵With the techniques from appendix A of [13], we can prove rigorously that they agree to more than a thousand decimal places, by using 300 forced double roots at the vector lengths from the hexagonal lattice.

Numerical linear or semidefinite programming succeeds at small N , but has been limited to $N \lesssim 100$ by the computational cost. It is much faster to trade the linear program for a nonlinear optimization over the roots of $f(\Delta)$.

At the optimum, $f(\Delta)$ is found empirically to have single roots at $\Delta_0 = 0$ and $\Delta_1 = \Delta_{\text{gap}}$, and $N - 1$ double roots $\Delta_2, \Delta_3, \dots, \Delta_N$. Assuming this to hold in general, we can restate the optimization problem as follows: fix Δ_1 , and maximize $f(0)$ over the parameters α_j for $1 \leq j \leq 2N$ and Δ_n for $2 \leq n \leq N$, subject to the pattern of roots

$$\begin{aligned} f(\Delta_1) &= 0 \text{ and} \\ f(\Delta_n) &= f'(\Delta_n) = 0 \text{ for } 2 \leq n \leq N. \end{aligned} \tag{2.22}$$

If the optimized function has $f(0) > 0$, then this value of Δ_1 is excluded. The marginal bound has $f(0) = 0$, and the corresponding Δ_1 gives $\Delta_1^{\text{LP},N}(c)$.

If this problem can be solved with some mild non-degeneracy conditions, then it probably gives $\Delta_1^{\text{LP},N}(c)$ (see section 5 of [11]). However, there is no guarantee that the optimum must be of this form, and it fails for $c = 3/2$. Specifically, when $c = 3/2$ it works for $1 \leq N \leq 21$ and $27 \leq N \leq 32$, but for $22 \leq N \leq 26$ and $N = 33$ there is a different pattern of roots. We do not know what happens as $N \rightarrow \infty$.

Aside from $c = 3/2$, this method has always worked in practice.⁶ The resulting bound can be made rigorous simply by proving that the optimal functional satisfies the positivity conditions.

This is essentially the method used by Cohn and Elkies [8]. In the conformal bootstrap, a similar approach was first discussed by El-Showk and Paulos in the context of 1d correlation functions [46]. Recently, their methods were improved by Afkhami-Jeddi, Hartman, and Tajdini and applied to the modular bootstrap [12]. The first step is to dualize the optimization problem. The dual problem, together with the equation $f(0) = 0$, leads to the equations

$$f_k(0) + \sum_{n=1}^N d_n f_k(\Delta_n) = 0 \text{ for } 1 \leq k \leq 2N. \tag{2.23}$$

Here the unknowns are d_1, \dots, d_N and $\Delta_1, \dots, \Delta_N$. These equations are nothing but the original crossing equation (2.10) truncated to N states. If we absorb the factors of $e^{-2\pi\Delta_n}$ from $f_k(\Delta_n)$ into the coefficients d_n , we are left with $2N$ polynomial equations in $2N$ unknowns.

Solutions to (2.23) with $d_n > 0$ place a lower bound $\Delta_1^{\text{LP},N}(c) \geq \Delta_1$, almost by definition. What is more surprising is that the solution saturating $\Delta_1 = \Delta_1^{\text{LP},N}(c)$ has exactly N states and can be found efficiently by Newton's method. The last ingredient we need in the algorithm is a procedure to generate the initial guess for Newton's method. We start at small N , where it is easy to find a guess by hand, and then gradually increase N , using the results from lower N to generate the next guess as in appendix B of [12]. This method allows for large jumps in N , while still converging to the bound within a few Newton steps.

⁶Our primary interest is in large c , and the main role of $c = 3/2$ is to dash any hope of proving that the method works for all c . We have no conceptual explanation for why this case seems to differ from all the others.




Figure 1. The linear programming bound for the sphere packing density.

3 Numerical results

3.1 Data and plots

The linear programming bound for the sphere packing density in \mathbb{R}^d is shown in figure 1 for $d \leq 48$, along with the record packing densities from [21, table I.1, pp. xix–xx], and table 1, while figure 2 shows $\Delta_1^{\text{LP}}(c)/c$. We believe that the truncation order in our calculations is high enough that these plots and table are indistinguishable from the fully optimized bounds; see appendix B for how we calibrated the convergence rate. All of our data is available from <https://hdl.handle.net/1721.1/125646>, including truncation orders, scaling dimensions, and degeneracies.

The linear programming bound is sharp when $d = 1, 2$ (conjecturally), 8, or 24, and seemingly nowhere else. From the perspective of discrete geometry, one of the most mysterious aspects is the role of these special dimensions. Unlike some other cases of the conformal bootstrap (see, for example, [47, section V.B.4]), the bound itself shows no sign of kinks or other non-analytic behavior at these points in figure 1. Instead, the upper bound looks much the same there as elsewhere, and it is the packing densities that show anomalous behavior.

Figures 1 and 2 faithfully reflect the qualitative properties of these bounds, but they are of limited use in extrapolating to high dimensions. For that purpose, a log-log plot is more effective, as in figure 3. This figure shows the linear programming bound and the record sphere packing densities from [21] in black. The green and red curves show




Figure 2. The ratio $\Delta_I^{LP}(c)/c$, together with lower bounds from sphere packings.




Figure 3. A comparison of the linear programming bound with other bounds. The annotations on the right show the limits of the colored curves in high dimensions, including our conjectured limit for the linear programming bound from Conjecture 3.1.

the best bounds that have been analytically derived: the green curve is Levenshtein’s bound [48], and the red curve is the Kabatyanskii-Levenshtein bound [19], computed using Levenshtein’s universal bound for spherical codes [49, 50] and the approach of Cohn and Zhao [20] (we review these bounds in sections 5.2 and 5.3). It is known that the linear programming bound is at least as strong as these bounds [8, 20], but no further analytic results are known. As shown in figure 3, our numerical calculations indicate that the linear programming bound is not much stronger than the better of these two bounds.

The blue curve in figure 3 is a lower bound for the linear programming bound due to Scardicchio, Stillinger, and Torquato [51], which is a variant of a bound obtained by Torquato and Stillinger [52]. No better lower bound is known for the linear programming bound in the limit as $d \rightarrow \infty$, and the known lower bounds for the sphere packing density are much worse. For example, the Minkowski-Hlawka bound says that there exist sphere packings of density at least 2^{-d} in \mathbb{R}^d , which amounts to the bottom edge of figure 3; this exponential decay rate is the same as in the ideal glass phase of the hard sphere model (see [53, p. 247]), and no known lower bound improves on this rate. These lower bounds are obtained from probabilistic or averaging arguments, and as far as we are aware, no explicit construction in dimension 2048 or greater has been shown to achieve the Minkowski-Hlawka bound.

3.2 Extrapolation

The annotations on the right side of figure 3 show the limits of the various curves as $d \rightarrow \infty$. The limits of the green, red, and blue curves are known explicitly, and one can see that even $d = 2048$ is not especially close to the asymptotic limit as $d \rightarrow \infty$. We know that the black curve always lies below the red curve, and it appears to be getting steadily closer. Thus, we expect the linear programming bound to be at most slightly better than the Kabatyanskii-Levenshtein bound in the limit as $d \rightarrow \infty$. What figure 3 does not reveal is whether the gap in fact tends to zero.

To estimate the asymptotic gap between the red and black curves, it is helpful to examine numerical data. Table 1 shows the difference $KL - LP$ between these two bounds, as well as the differences between consecutive values of $KL - LP$ and their ratios. The $KL - LP$ column does indeed seem to decrease for $d \geq 2$, and we expect that it is converging towards its limit at a rate proportional to $1/d$. In that case, the difference column to its right should be decreasing towards zero at the same rate $1/d$, and so the ratios in the last column should tend to 2. The behavior of the ratio column is not absolutely clear, but it does look like it may be increasing towards 2 beyond $d = 32$.

We therefore hypothesize that the difference column will continue to decrease by a factor between 1.92 and 2 in each additional row. In that case, the sum of all the entries below 0.00070 in the difference column must lie between

$$\sum_{n=1}^{\infty} \frac{0.00070}{2^n} = 0.00070 \quad \text{and} \quad \sum_{n=1}^{\infty} \frac{0.00070}{1.92^n} < 0.00077. \quad (3.1)$$

In other words, the $KL - LP$ column will fall below 0.00611 by an amount between 0.00070 and 0.00077, and so its limiting value will be between 0.00534 and 0.00541. We conclude

d	LP	KL	KL – LP	Difference	Ratio
1	0.00000	0.00000	0.00000		
2	-0.07049	0.29248	0.36297		
4	-0.15665	0.17511	0.33176	0.03122	0.23
8	-0.24737	-0.04879	0.19858	0.13318	1.59
16	-0.33192	-0.21692	0.11501	0.08357	1.87
32	-0.40382	-0.33342	0.07040	0.04461	1.55
64	-0.46101	-0.41947	0.04154	0.02885	1.73
128	-0.50432	-0.47947	0.02485	0.01669	1.82
256	-0.53589	-0.52023	0.01566	0.00919	1.87
512	-0.55824	-0.54749	0.01075	0.00491	1.90
1024	-0.57370	-0.56553	0.00816	0.00259	1.92
2048	-0.58418	-0.57737	0.00682	0.00135	1.92
4096	-0.59120	-0.58508	0.00611	0.00070	
∞	?	-0.59906	?		

Table 1. Numerical comparison of the linear programming bound in \mathbb{R}^d (denoted LP) with the Kabatyanskii-Levenshtein bound (denoted KL). The bounds are in the form $(\log_2 \text{density})/d$, and are rounded rather than truncated.

that the limit of the LP column must be -0.6044 ± 0.0001 (i.e., between -0.6045 and -0.6043). In other words, the linear programming bound will be approximately $2^{-0.6044d}$ when d is large.

Conjecture 3.1. *There exists a constant λ with $0.604 < \lambda < 0.605$ such that the linear programming bound for the sphere packing density in \mathbb{R}^d is $2^{-(\lambda+o(1))d}$ as $d \rightarrow \infty$ when the auxiliary function is fully optimized.*

Of course we have no proof of this conjecture, or even a heuristic derivation. It is possible that the numbers could behave entirely differently when d is much larger, but that does not seem plausible. We are very confident in the first three decimal places of the estimate 0.6044 for λ , and fairly confident in the fourth. In fact, we have a proposal for the exact constant:

Conjecture 3.2. *The constant λ in Conjecture 3.1 is given by $2^{-\lambda} = \sqrt{e/(2\pi)}$. Equivalently,*

$$\lim_{c \rightarrow \infty} \frac{\Delta_1^{\text{LP}}(c)}{c} = \frac{1}{\pi^2} \tag{3.2}$$

and

$$\lim_{d \rightarrow \infty} \frac{A_-(d)}{\sqrt{d}} = \frac{1}{\pi}, \tag{3.3}$$

where $A_-(d)$ denotes the optimal radius for the -1 eigenfunction uncertainty principle in \mathbb{R}^d .

Of course Conjecture 3.2 is speculative, and four digits of accuracy is far from enough to make a definitive argument for this value. The equivalent limits are much more appealing

d	-1 radius	+1 radius	Ratio
4	1.2038	0.9660	1.2462
8	1.4142	1.2173	1.1618
16	1.7393	1.5812	1.1000
32	2.2241	2.1000	1.0591
64	2.9317	2.8359	1.0338
128	3.9515	3.8787	1.0188
256	5.4109	5.3565	1.0102
512	7.4905	7.4504	1.0054
1024	10.446	10.417	1.0028
2048	14.640	14.619	1.0014

Table 2. The ratio of the -1 eigenfunction uncertainty principle radius to the $+1$ radius.

than the formula for λ , and their simplicity justifies going out on a limb. We have no great faith in this conjecture, but it is worth noting that a simple formula fits the data beautifully. For comparison, $\lambda = 0.6044 \pm 0.0001$ would amount to

$$0.318287 < \lim_{d \rightarrow \infty} \frac{A_-(d)}{\sqrt{d}} < 0.318333, \tag{3.4}$$

and $1/\pi = 0.318309\dots$

Our calculations also support Conjecture 1.5 from [11], which says that the sign change radii for the $+1$ and -1 eigenfunction uncertainty principles in \mathbb{R}^d are the same asymptotically as $d \rightarrow \infty$. See table 2 and <https://hdl.handle.net/1721.1/125646> for the numerical data.⁷ In the notation of [11],

$$\lim_{d \rightarrow \infty} \frac{A_+(d)}{\sqrt{d}} = \lim_{d \rightarrow \infty} \frac{A_-(d)}{\sqrt{d}}. \tag{3.5}$$

Specifically, the ratio $A_+(d)/A_-(d)$ seems to be $1 + O(1/d)$.

4 Properties of the spectrum and degeneracies

The spectra of our numerically optimized solutions of the spinless modular bootstrap behave remarkably regularly when the central charge c is large. Figure 4 shows the scaling dimensions for $c \leq 64$, with the sharp cases $c = 1/2, 1, 4$, and 12 highlighted. When $c = 4$ or 12 , the scaling dimensions are positive integers, excluding 1 when $c = 12$. The green lines in the figure extrapolate these arithmetic progressions to other values of c . In other words, the n -th green line from the bottom amounts to $\Delta_n = n + (c - 4)/8$. While this equation never holds exactly except when $c = 4$ or 12 , it is an excellent approximation when n is large and $c \geq 4$. Later in this section we will examine how close this approximation is.

The equation $\Delta_n = n + (c - 4)/8$ amounts to the 1d generalized free fermion spectrum. This spectrum arose in analytic functionals for the 1d conformal bootstrap constructed by

⁷The reason why this table omits $d = 1$ and 2 is that we do not have good numerical data in these dimensions. See the end of section 4 in [11].




Figure 4. The spectra $\Delta_n^{\text{LP}}(c)$ for $2 \leq c \leq 64$ and $c \in \{1/2, 1\}$, drawn in black. The dots highlight the spectra for $c \in \{1/2, 1, 4, 12\}$, and the green lines show the 1d generalized free fermion spectra.

Mazáč [26], which were generalized to a basis by Mazáč and Paulos [54]. Hartman, Mazáč, and Rastelli [9] discovered that these functionals could be adapted to the 2d modular bootstrap with $U(1)^c$ or Virasoro symmetry, and special cases were independently constructed by Rolin and Wagner [55] and by Feigenbaum, Grabner, and Hardin [56].

Figure 5 shows the sphere packing bounds obtained from these functionals. No proof is known that the functionals satisfy the required inequalities, and indeed they do not for $8 < d < 24$; in particular, they would prove an impossibly good linear programming bound for $d = 16$ (see the dual bounds in [41]). However, the inequalities seem to hold for other values of d . Unfortunately, the resulting bounds are disappointing. The fact that the $\Delta_1^{\text{LP}}(c)$ curve in figure 4 bends below the free fermion line $\Delta_1 = 1 + (c - 4)/8$ is crucial for obtaining a strong bound, and the quality of the bound depends on the degree of deflection.

In contrast to the behavior for large c , the spectra for small c are much less regularly spaced. As we decrease c below 4, the scaling dimensions in figure 4 start to diverge unpredictably from the green lines, and the behavior for $c \in \{1/2, 1\}$ is entirely different. Our numerical techniques break down at $c = 3/2$, and it is presumably not a coincidence that this failure occurs at the transition between different regimes. It would be interesting to explore this transition for $1 < c < 2$.

4.1 Convergence to the free fermion spectrum

Figure 6 shows how the spectrum converges as we increase the truncation order N . The limiting values for $\Delta_n^{\text{LP},N}(c)$ as $N \rightarrow \infty$ are quite close to the 1d generalized free fermion




Figure 5. The linear programming bound and the hypothetical bound based on the 1d generalized free fermion spectrum $\Delta_n(c) = n + (c - 4)/8$.




Figure 6. The spectrum for $c = 128$ computed using truncation order $N = 1, 2, 4, \dots, 128$ (red through black) and the free fermion spectrum $\Delta_n(c) = n + (c - 4)/8$ (green).




Figure 7. The probability density function for the normalized spectrum is 1 on $[0, 2/\pi]$ and $x \mapsto x^{-1/2}(4/\pi - x)^{1/2}$ on $[2/\pi, 4/\pi]$.




Figure 8. The cumulative distribution functions for the normalized spectrum $\Delta_n^{\text{LP},N}(4)/N$ with $1 \leq n \leq N$ when $N = 128$ (red) and in the limit as $N \rightarrow \infty$ (black).

values, but there is a substantial divergence when n is large enough relative to N , as noted in [57, section 3] and [12, Figure 2]. Although we have no proof, we can predict the form of this divergence. It occurs starting at $n \sim (2/\pi)N$, with shape determined by the following conjecture:

Conjecture 4.1. *For each $c \geq 2$, the distribution of the normalized spectrum $\Delta_n^{\text{LP},N}(c)/N$ with $1 \leq n \leq N$ converges as $N \rightarrow \infty$ to the probability distribution on the interval $[0, 4/\pi]$ with density function 1 on $[0, 2/\pi]$ and $x \mapsto x^{-1/2}(4/\pi - x)^{1/2}$ on $[2/\pi, 4/\pi]$, shown in figure 7.*

See figure 8 for a comparison with numerical data for $c = 4$ and $N = 128$. The behavior is similar for all $c \geq 2$, with slower convergence as c grows.




Figure 9. The gaps $\log_{10}(1/|\Delta_n(c) - (n + (c - 4)/8)|)$ for $c = 16$ (black), $c = 20$ (red), $c = 24$ (green), and $c = 28$ (blue).

Conjecture 4.1 says that the distribution shifts from uniform to a beta distribution around $n \sim (2/\pi)N$. The uniform distribution corresponds to the 1d generalized free fermion spectrum, and the beta distribution describes the root distribution of high-degree Laguerre polynomials (see [58, Theorem 1]). Specifically, if we normalize the roots of the highest-degree polynomial $L_{4N-1}^{(c-1)}(4\pi\Delta)$ from the truncated crossing equation (2.23) by dividing by a factor of N , then their distribution converges as $N \rightarrow \infty$ to the beta function on $[0, 4/\pi]$ with density $x \mapsto x^{-1/2}(4/\pi - x)^{1/2}/2$. From this perspective, the transition in Conjecture 4.1 is between the uniform limiting behavior as $N \rightarrow \infty$ and a generic root distribution corresponding to high-degree Fourier eigenfunctions.

Conjecture 4.1 gives the following description of the curves in figure 6 via Theorem 4 in [59]. We wish to approximate $\Delta_n^{\text{LP},N}(c)/N$ as $N \rightarrow \infty$ with $n/N \rightarrow \alpha$ for some constant $\alpha \in [0, 1]$. The conjecture says that

$$\frac{\Delta_n^{\text{LP},N}(c)}{N} \rightarrow \begin{cases} \alpha & \text{if } \alpha \leq 2/\pi, \text{ and} \\ (2/\pi)(1 + \cos \beta) & \text{if } \alpha \geq 2/\pi, \end{cases} \quad (4.1)$$

where β is the solution of $\beta - \sin \beta = (1 - \alpha)\pi/2$ with $0 \leq \beta \leq \pi/2$. We will give some motivation for the high-energy portion of this formula in section 4.6, after discussing degeneracies.

4.2 Deviations from the free fermion spectrum

Aside from $c = 4$ or 12 , the spectrum $\Delta_n^{\text{LP}}(c)$ is not exactly equal to the 1d generalized free fermion spectrum $\Delta_n(c) = n + (c - 4)/8$. Instead, we always find some error in this approximation, and one natural question is whether the error tends to 0 as $n \rightarrow \infty$. Figure 9 shows four test cases, namely $c \in \{16, 20, 24, 28\}$. In each of these cases, the error $|\Delta_n^{\text{LP}}(c) - (n + (c - 4)/8)|$ becomes small, no more than 10^{-5} when n is large. However,




Figure 10. The maximum $\varepsilon(c)$ of $|\Delta_n^{\text{LP}}(c) - (n + (c - 4)/8)|$ over $2c \leq n \leq 128$.

there is a striking difference between two different scenarios: when $c = 16$ or 24 , the error stabilizes above zero, while it seems to converge to 0 when $c = 20$ or 28 . We have no conceptual explanation of this behavior, which seems to be periodic modulo 8 as c varies, with multiples of 8 being the worst case and 4 modulo 8 being the best case (and the only case with convergence to zero). Note that $c = 4$ and 12 fall into the latter category.

Figure 10 shows another aspect of this periodicity. In this figure, we plot

$$\varepsilon(c) = \max_{2c \leq n \leq 128} |\Delta_n^{\text{LP}}(c) - (n + (c - 4)/8)|. \quad (4.2)$$

Here the upper bound of 128 for n should be viewed as a stand-in for infinity, and the lower bound $2c$ is intended to exclude sporadic behavior for small n before convergence; in other words, we view $\varepsilon(c)$ as an approximation to

$$\limsup_{n \rightarrow \infty} |\Delta_n^{\text{LP}}(c) - (n + (c - 4)/8)|. \quad (4.3)$$

In the plot, there are singularities at $c = 4$ and 12 that correspond to $\varepsilon(c) = 0$, with similar cusps at $c = 20, 28$, and 36 , and again we see the largest error terms when c is a multiple of 8. This unexpected periodicity shows that the spinless modular bootstrap has a much richer number-theoretic structure than is apparent from the smooth plot of $\Delta_1^{\text{LP}}(c)$ in figure 2.

We can summarize figures 9 and 10 with the following conjecture:

Conjecture 4.2. *For $c \geq 4$, the quantity*

$$\limsup_{n \rightarrow \infty} |\Delta_n^{\text{LP}}(c) - (n + (c - 4)/8)| \tag{4.4}$$

vanishes if and only if c is an integer congruent to 4 modulo 8, and there is a constant α such that

$$\limsup_{n \rightarrow \infty} |\Delta_n^{\text{LP}}(c) - (n + (c - 4)/8)| \leq \alpha 2^{-c} \tag{4.5}$$

for all $c \geq 4$. Furthermore, we can take $\alpha = 1/\pi + o(1)$ as $c \rightarrow \infty$.

4.3 Growth rate of the degeneracies

Solving the truncated crossing equation (2.23) yields not just scaling dimensions $\Delta_1^{\text{LP},N}(c)$ but also corresponding degeneracies $d_n^{\text{LP},N}(c)$, which converge as $N \rightarrow \infty$ to the degeneracies $d_n^{\text{LP}}(c)$ of a hypothetical CFT that attains the spinless modular bootstrap bound. For $c \leq 50$ and $c \notin \{1/2, 1, 4, 12\}$, numerical calculations show that these degeneracies are not integers, and thus they cannot come from an actual CFT.⁸ For larger c , it is difficult to assess integrality, because the degeneracies grow exponentially as $c \rightarrow \infty$ and must therefore be computed to high precision.

The cumulative growth rate of the degeneracies is determined by modularity as follows. Because $\eta(-1/\tau) = (\tau/i)^{1/2}\eta(\tau)$, the modular invariance of the partition function $\mathcal{Z}(\tau) = \eta(\tau)^{-2c} \sum_{\Delta} d_{\Delta} e^{2\pi i \tau \Delta}$ implies the identity

$$\sum_{\Delta} d_{\Delta} e^{2\pi i \tau \Delta} = (i/\tau)^c \sum_{\Delta} e^{-2\pi i \Delta/\tau}. \tag{4.6}$$

If we set $\tau = i\beta/(2\pi)$ and let $\beta \rightarrow 0$, we find that

$$\sum_{\Delta} d_{\Delta} e^{-\beta \Delta} \sim (2\pi/\beta)^c. \tag{4.7}$$

Now the Karamata Tauberian theorem [60, Theorem 4.3 of Chapter V] implies that

$$\sum_{\Delta \leq A} d_{\Delta} \sim \frac{(2\pi A)^c}{\Gamma(c+1)} \tag{4.8}$$

as $A \rightarrow \infty$.

In other words, the function

$$\rho_c(\Delta) = \frac{(2\pi)^c \Delta^{c-1}}{\Gamma(c)} \tag{4.9}$$

is the $U(1)^c$ analogue of the Cardy formula [61] for degeneracies, because

$$\int_0^A d\Delta \rho_c(\Delta) = \frac{(2\pi A)^c}{\Gamma(c+1)}. \tag{4.10}$$

⁸They could still come from a sphere packing. In that case d_n would be the average number of sphere centers at distance $\sqrt{2\Delta_n}$ from a given sphere center, which need not be an integer.




Figure 11. The normalized degeneracies $d_n^{\text{LP}}(4)/\rho_4(\Delta_n^{\text{LP}}(4))$ for $1 \leq n \leq 256$. As $n \rightarrow \infty$, they fluctuate between $1/\zeta(4) = 0.9239\dots$ and $\zeta(3)/\zeta(4) = 1.1106\dots$, shown with gray lines, and they are bounded away from 1.

A similar formula applies to operator-product coefficients that appear in the bootstrap equations for conformal correlators [62].

Because the scaling dimensions $\Delta_n^{\text{LP}}(c)$ are uniformly spaced with distance 1 asymptotically, we expect that $d_n^{\text{LP}}(c)$ will be roughly $\rho_c(\Delta_n^{\text{LP}}(c))$ as $n \rightarrow \infty$. The asymptotic formula (4.8) gives a sense in which this approximation is true on average, but we will see that the precise behavior is far more delicate. In the Virasoro case, a more fine-grained understanding of the asymptotic spectrum has been obtained recently using complex Tauberian theorems [63–67]. It would be interesting to do the same for sphere packing. This could perhaps explain the linear portion of the large- c spectrum, where the level spacing is very close to 1.

4.4 Degeneracies for $c = 4$ and 12

The degeneracies $d_n^{\text{LP}}(4)$ and $d_n^{\text{LP}}(12)$ are the coefficients of the theta series of the E_8 and Leech lattices, respectively. Much is known about these modular forms, including precise descriptions of their coefficients (see, for example, [21, p. 122 and p. 134]). The degeneracies are well understood, but far more subtle than the scaling dimensions $\Delta_n^{\text{LP}}(4) = n$ and $\Delta_n^{\text{LP}}(12) = n + 1$.

Figure 11 shows the normalized degeneracies $d_n^{\text{LP}}(4)/\rho_4(\Delta_n^{\text{LP}}(4))$ for $c = 4$. They are bounded above and below, but do not converge to 1; instead, they are strictly bounded away from 1. The most noteworthy aspect of figure 11 is that the normalized degeneracies are almost, but not quite, periodic.⁹ This near periodicity is explained by a classical formula for coefficients of Eisenstein series: $d_n^{\text{LP}}(4)/\rho_4(\Delta_n^{\text{LP}}(4)) = \sigma_3(n)/(\zeta(4)n^3)$, where $\sigma_k(n)$ denotes the sum of the k -th powers of the divisors of n , and ζ is the Riemann zeta

⁹See [68, Theorem 5.13A.1] for the theory of almost periodic functions.




Figure 12. The maximum $\delta(c)$ of $|d_n^{\text{LP}}(c)/\rho_c(\Delta_n^{\text{LP}}(c)) - 1|$ over $2c \leq n \leq 128$.

function. This function is not periodic, but for each $\varepsilon > 0$, there exists a natural number m such that if $n_1 \equiv n_2 \pmod{m}$, then $|\sigma_3(n_1)/n_1^3 - \sigma_3(n_2)/n_2^3| < \varepsilon$. In other words, for each $\varepsilon > 0$ it is approximately periodic to within ε , with the period length growing as $\varepsilon \rightarrow 0$. Similarly, the function $n \mapsto d_n^{\text{LP}}(12)/\rho_{12}(\Delta_n^{\text{LP}}(12))$ is the sum of the almost periodic function $n \mapsto \sigma_{11}(n)/(\zeta(12)n^3)$ and a term converging to zero as $n \rightarrow \infty$.

4.5 Degeneracies for arbitrary c

The unexpected periodicity modulo 8 for scaling dimensions has a counterpart for degeneracies, as shown in figure 12. Here the multiples of 8 are the best case for the accuracy of the $U(1)^c$ Cardy formula, while integers that are 4 modulo 8 are the worst case. Unlike the case of scaling dimensions, the cusps in figure 12 do not seem to correspond to zero error in the limit as $n \rightarrow \infty$. Instead, see figure 13. Perhaps this discrepancy indicates that $\rho_c(\Delta)$ should be replaced with some better approximation.

Based on this data and the cases $c = 4$ and 12, we make the following conjecture, which is a more precise analogue of Conjecture 4.2:

Conjecture 4.3. For $c > 2$,

$$\limsup_{n \rightarrow \infty} \frac{d_n^{\text{LP}}(c)}{\rho_c(\Delta_n^{\text{LP}}(c))} \leq \frac{\zeta(c-1)}{\zeta(c)} = 1 + 2^{-c} + O(3^{-c}) \tag{4.11}$$

and

$$\liminf_{n \rightarrow \infty} \frac{d_n^{\text{LP}}(c)}{\rho_c(\Delta_n^{\text{LP}}(c))} \geq \frac{1}{\zeta(c)} = 1 - 2^{-c} + O(3^{-c}), \tag{4.12}$$




Figure 13. The gaps $\log_{10}(1/|d_n^{LP}(c)/\rho_c(\Delta_n^{LP}(c)) - 1|)$ for $c = 8$ (black), $c = 12$ (red), $c = 16$ (green), and $c = 20$ (blue).

with equality whenever c is an integer that is congruent to 4 modulo 8.

Equality provably holds for $c = 4$ or 12.

4.6 Possible explanation of the high-energy spectrum

The formula for the upper portion of the high-energy spectrum, given by the second line in (4.1), was motivated by the following calculation. It does not fully explain the formula, but it indicates why it is a reasonable guess.

The counterpart of the Cardy formula in terms of functionals acting on the crossing equation is the integral identity

$$\int_0^\infty d\Delta \frac{(2\pi)^c \Delta^{(c-1)}}{\Gamma(c)} e^{-2\pi\Delta} L_k^{(c-1)}(4\pi\Delta) = (-1)^k L_k^{(c-1)}(0), \quad (4.13)$$

which is the evaluation at 0 of the condition of being a radial Fourier eigenfunction. We will use this identity to generate an exact solution of the truncated crossing equations, which has too many states but is otherwise suggestive of the optimal solution.

For small enough k , we can evaluate the integral exactly using the Gauss-Laguerre quadrature formula

$$\int_0^\infty dx x^{c-1} e^{-x} p(x) = \sum_{m=1}^n w_m p(x_m), \quad (4.14)$$

which holds for all polynomials p of degree at most $2n - 1$ if we use the roots of $L_n^{(c-1)}(x)$ as the quadrature nodes x_m with weights

$$w_m = \frac{\Gamma(n+c)x_m}{n!(n+1)^2(L_{n+1}^{(c-1)}(x_m))^2}. \quad (4.15)$$

If we take $n = 2N$ and $p(x) = L_k^{(c-1)}(2x)$, we find that for $k \leq 4N - 1$,

$$\int_0^\infty d\Delta \frac{(2\pi)^c \Delta^{(c-1)}}{\Gamma(c)} e^{-2\pi\Delta} L_k^{(c-1)}(4\pi\Delta) = \sum_{m=1}^{2N} d_m L_k^{(c-1)}(4\pi\Delta_m), \quad (4.16)$$

where $d_m = w_m/\Gamma(c)$ and the dimensions Δ_m satisfy $L_{2N}^{(c-1)}(2\pi\Delta_m) = 0$. In other words, we have generated a solution to the equations

$$L_k^{(c-1)}(0) + \sum_{m=1}^{2N} d_m L_k^{(c-1)}(4\pi\Delta_m) = 0 \quad (4.17)$$

for odd $k \leq 4N - 1$. This solution has $2N + 1$ states, while our numerical method involves finding a solution to these truncated crossing equations with only $N + 1$ states, so this solution has no direct bearing on the bound. However we observe numerically that the high energy spectrum approximately agrees. The roots of the Laguerre polynomial $L_{2N}^{(c-1)}(2\pi\Delta)$ for large N and Δ are given by the beta distribution described in section 4.1, which motivates its appearance.

5 Spherical codes and implied kissing numbers

In sphere packing terms, the scaling dimensions Δ_n measure the distances $\sqrt{2\Delta_n}$ between distinct sphere centers in the packing, and the corresponding degeneracy is the average number of centers at that distance from a given center. In particular, the first degeneracy d_1 is the average number of tangencies for spheres in the packing, i.e., the *average kissing number* of the packing. De Laat, Oliveira, and Vallentin [69] showed how to strengthen the linear programming bound for sphere packing by incorporating geometric bounds for the degeneracies, which go beyond the modular invariance of the partition function. In this section, we will use this idea systematically to explore when the linear programming bound can be sharp and how to improve on it. Along the way, we will review the Kabatyanskii-Levenshtein bound.

As motivation for this line of work, consider the *implied kissing number* $d_1^{\text{LP}}(d/2)$, which is the average kissing number in a hypothetical d -dimensional packing achieving the linear programming bound. When $d = 1, 2, 8,$ or 24 , we of course obtain the kissing number of the optimal packing, but in general we obtain unrealistic numbers. For example, when $d = 4$ the implied kissing number is $26.43\dots$, which exceeds Musin's optimal bound of 24 for the four-dimensional kissing number [70]; it is therefore impossible for any packing to achieve the exact linear programming bound in \mathbb{R}^4 . As figure 14 shows, the implied kissing number is impossibly high for every $d \leq 24$ except the known sharp cases. Within this range of dimensions, it perfectly delineates which cases are sharp.

Figure 15 shows how the implied kissing number grows in high dimensions. Comparing it with upper bounds turns out to be surprisingly subtle, and we will do so in figure 16 once we have explained more about the needed bounds. Aside from low dimensions, figure 15 looks similar to figure 3, and that is not a coincidence: table 3 indicates that the implied kissing number is $2^{d+o(d)}$ times the linear programming bound as $d \rightarrow \infty$. This relationship




Figure 14. The implied kissing number from the linear programming bound, compared with the best upper bound known for the actual kissing number [71] and the current record [21].

is easily explained using the $U(1)^c$ Cardy formula, because the sphere packing density is $\rho_c(\Delta_1)\Delta_1/(c4^c)$ in terms of the spectral gap Δ_1 and $c = d/2$. We can approximate the implied kissing number $d_1^{LP}(c)$ by $\rho_c(\Delta_1^{LP}(c))$ using the Cardy formula; because of the size of $\Delta_1^{LP}(c)$ we expect some error, but the error factor should be subexponential in c , in fact roughly $\Delta_2^{LP}(c) - \Delta_1^{LP}(c)$. We conclude that the linear programming bound for the packing density is $d_1^{LP}(c)/(4 + o(1))^c$ as $c \rightarrow \infty$, which is the desired relationship.

To prove bounds for kissing numbers, the relevant optimization problem is the spherical code problem, which is a compact analogue of the sphere packing problem. In dimension d and with minimal angle θ this problem asks how large a subset \mathcal{C} of the unit sphere in \mathbb{R}^d can be if $\langle x, y \rangle \geq s$ for all distinct $x, y \in \mathcal{C}$, where $s = \cos \theta$. In other words, all points in \mathcal{C} must be separated by at least a distance of θ along the surface of the sphere, and so \mathcal{C} yields a packing with spherical caps of radius $\theta/2$. Such a set is called a *spherical code* with minimal angle θ . The kissing problem amounts to the case $\theta = \pi/3$; note that here we are considering the kissing problem for a single sphere, rather than averaged over a packing in Euclidean space.

Let $A(d, s)$ be the largest possible size of such a code. Delsarte, Goethals, and Seidel [74] introduced a linear programming bound for $A(d, s)$, which Kabatyanskii and Levenshtein [19] used to obtain the best sphere packing density bounds known in Euclidean space.

After briefly reviewing this linear programming bound, we will discuss two applications of spherical codes to the sphere packing problem, followed by a new average kissing bound. First we discuss the Kabatyanskii-Levenshtein bound, using the approach from [50]. Then




Figure 15. The implied kissing number in high dimensions, compared with the record kissing number [21, 72] and the excluded volume lower bound for the kissing number from [73].

d	LP	IK	$1 + \text{LP} - \text{IK}$
1	0.00000	1.00000	0.00000
2	-0.07049	1.29248	-0.36297
4	-0.15665	1.18103	-0.33768
8	-0.24737	0.98836	-0.23573
16	-0.33192	0.81510	-0.14702
32	-0.40382	0.68279	-0.08661
64	-0.46101	0.58837	-0.04937
128	-0.50432	0.52325	-0.02757
256	-0.53589	0.47929	-0.01518
512	-0.55824	0.45003	-0.00827
1024	-0.57370	0.43077	-0.00447
2048	-0.58418	0.41822	-0.00240
4096	-0.59120	0.41009	-0.00128
∞	-0.6044	0.3956	0.00000

Table 3. Numerical comparison of the linear programming bound for the sphere packing density in \mathbb{R}^d (denoted LP) with the implied kissing number (denoted IK), both in the form $\log_2(\cdot)/d$.

we discuss a strengthening of the linear programming bound for the sphere packing problem through bounds for spherical codes, and its implications for when the bound can be tight. We conclude with a linear programming bound for the average kissing number.

The results of this section rely on the geometry of the sphere packing problem and do not appear to have any direct application to CFTs. On the other hand, conceptually, upper bounds on the average kissing number are similar to upper bounds on operator-product coefficients often considered in the bootstrap literature (e.g., [75, 76]). It is also our hope that the methods in this section (in particular, the application of the Christoffel-Darboux formula to produce positive auxiliary functions) will inspire new analytic approaches to the conformal bootstrap.

5.1 The linear programming bound

The analogue of the radial Fourier transform on the surface of a sphere is the expansion in terms of zonal spherical harmonics, which uses the following orthogonal polynomials.¹⁰ Let d be the dimension of the spherical code, and let a and b be nonnegative integers. (We can take $a = b = 0$ for now, but we will make use of a and b in the Kabatyanski-Levenshtein bound.)

Let

$$w^{a,b}(t) = \frac{\Gamma(a+b+d-1)}{2^{a+b+d-2}\Gamma(a+\frac{d-1}{2})\Gamma(b+\frac{d-1}{2})} (1-t)^a(1+t)^b(1-t^2)^{(d-3)/2}, \quad (5.1)$$

where the normalization is chosen so that $\int_{-1}^1 dt w^{a,b}(t) = 1$. Define the orthogonal polynomials $Q_i^{a,b}(t)$ with $\deg(Q_i^{a,b}) = i$ and positive leading coefficients by

$$\int_{-1}^1 dt Q_i^{a,b}(t) Q_j^{a,b}(t) w^{a,b}(t) = \delta_{i,j} \quad (5.2)$$

for $i, j \geq 0$. Up to normalization, these polynomials are the Jacobi polynomials with parameters $(d-3)/2 + a$ and $(d-3)/2 + b$.

The parameters $a = b = 0$ are particularly important, and the polynomials $Q_i := Q_i^{0,0}$ are known as the ultraspherical polynomials in dimension d . For a continuous function $f: [-1, 1] \rightarrow \mathbb{R}$,

$$f_0 := \int_{-1}^1 dt f(t) w^{0,0}(t) = \int_{S^{d-1}} d\mu(x) f(\langle x, e \rangle), \quad (5.3)$$

where $e \in S^{d-1}$ is an arbitrary point and μ is the surface measure on the sphere S^{d-1} , normalized so that $\mu(S^{d-1}) = 1$. Therefore the polynomials Q_i are orthogonal if we think of them as zonal functions on S^{d-1} , i.e., functions $x \mapsto Q_i(\langle x, e \rangle)$ invariant under the stabilizer subgroup of $O(d)$ with respect to e . Moreover, these polynomials are of positive type: for all finite $\mathcal{C} \subseteq S^{d-1}$ and all coefficients $c_x \in \mathbb{R}$ for $x \in \mathcal{C}$,

$$\sum_{x,y \in \mathcal{C}} c_x c_y Q_i(\langle x, y \rangle) \geq 0. \quad (5.4)$$

¹⁰One key conceptual difference between spheres and Euclidean space is that \mathbb{R}^d is its own Pontryagin dual, which means Fourier eigenfunctions make sense, while that concept does not apply to spheres.

This inequality follows from the addition formula

$$\frac{Q_i(\langle x, y \rangle)}{Q_i(1)} = \frac{1}{r_i} \sum_{j=1}^{r_i} v_{i,j}(x)v_{i,j}(y), \tag{5.5}$$

where the functions $v_{i,j}: \mathbb{R}^d \rightarrow \mathbb{R}$ for $j = 1, \dots, r_i$ are an orthonormal basis of the spherical harmonics of degree i (see [77, Theorem 9.6.3]); specifically, the addition formula shows that the left side of (5.4) is a square and therefore nonnegative.

The linear programming bound for spherical codes converts an auxiliary function into an upper bound on the greatest size of a spherical code:

Theorem 5.1 ([74]). *Let $f: [-1, 1] \rightarrow \mathbb{R}$ be a continuous function and $s \in [-1, 1]$. If f is of positive type as a zonal function on S^{d-1} , $f(t) \leq 0$ for $t \in [-1, s]$, and $f_0 \neq 0$, then $A(d, s)$ is at most $f(1)/f_0$.*

By Schoenberg’s theorem [78], every continuous function $f: [-1, 1] \rightarrow \mathbb{R}$ of positive type is of the form $f(x) = \sum_{n=0}^{\infty} f_n Q_n(x)$ with $f_n \geq 0$, where convergence is uniform and absolute. Therefore the linear programming bound can be approximated arbitrarily well by the minimum of $f(1)/f_0$ over functions f of the form

$$f(t) = \sum_{n=0}^N f_n Q_n(t) \tag{5.6}$$

with $f_0 > 0$, $f_1, \dots, f_N \geq 0$, and $f(t) \leq 0$ for $t \in [-1, s]$.

One can optimize this bound numerically for given d and N using semidefinite programming. Specifically, we can create a semidefinite program where f_0, \dots, f_N are one-by-one positive semidefinite matrices, and the inequality constraint is modeled as

$$f(t) = \begin{cases} -v_0(t)^\top X v_0(t) - (t+1)(s-t)v_1(t)^\top Y v_1(t) & \text{if } N \text{ is even, and} \\ -(t+1)v_0(t)^\top X v_0(t) - (s-t)v_0(t)^\top Y v_0(t) & \text{if } N \text{ is odd,} \end{cases} \tag{5.7}$$

where X and Y are positive semidefinite matrices and

$$v_k(t) = (Q_0(t), \dots, Q_{\lfloor N/2 \rfloor - k}(t)). \tag{5.8}$$

If we additionally require $f_0 = 1$, then the objective is $f(1)$, which means we are minimizing a linear functional over positive semidefinite matrices with linear constraints. This semidefinite program can be solved numerically on a computer.

Shtrom [79] computed the exact linear programming bound for the kissing number $A(d, 1/2)$ when $d \leq 146$, by determining the optimal value of N , beyond which there is no improvement. We have extended these computations to $d \leq 424$ using $N = 95$, which appears to be high enough in this range of dimensions and in any case should closely approximate the optimum. Figure 16 shows the ratio of the implied kissing number to this upper bound. They are very close to each other in size, but their precise ratio seems difficult to predict, and we do not know what happens as $d \rightarrow \infty$. No sphere packing can match the linear programming bound for density when this ratio is strictly greater than




Figure 16. The ratio of the implied kissing number to the linear programming bound for the kissing number. It dips below 1 at dimensions 180–181, 192, 282–283, 294–296, and many beyond that.

1. Our initial hope was that this condition would rule out every dimension $d > 24$, but it does not. Instead, further progress may depend on more powerful bounds for the kissing number, such as semidefinite programming bounds [71, 80].

This approach can rule out exact equality in the linear programming bound for packing density, but it does not give a quantitative improvement. We will return to the problem of improving the density bound, once we explain Levenshtein’s universal bound and the Kabatyanskii-Levenshtein bound.

5.2 Levenshtein’s universal bound

Levenshtein’s universal bound is the best bound that has been analytically derived from Theorem 5.1; our explanation follows section 5.4 in [50]. Let $t_k^{a,b} \in [-1, 1)$ be the largest root of $Q_k^{a,b}(t)$, where we set $t_0^{1,1} = -1$. One can show that

$$t_{k-1}^{1,1} < t_k^{1,0} < t_k^{1,1} \tag{5.9}$$

for all $k \geq 1$ (see (5.89) in [50]). For $s \in [-1, 1)$, define the function

$$f^{(s)}(t) = \begin{cases} (t-s)(K_{k-1}^{1,0}(t,s))^2 & \text{if } t_{k-1}^{1,1} \leq s < t_k^{1,0}, \text{ and} \\ (t+1)(t-s)(K_{k-1}^{1,1}(t,s))^2 & \text{if } t_k^{1,0} \leq s < t_k^{1,1}, \end{cases} \tag{5.10}$$

where

$$K_{k-1}^{a,b}(t,s) = \sum_{i=0}^{k-1} Q_i^{a,b}(t)Q_i^{a,b}(s). \tag{5.11}$$

This function $f^{(s)}$ will prove a bound in Theorem 5.1, once we verify the hypotheses of the theorem (see Theorem 5.42 in [50] for more details than we give here). By construction, $f^{(s)}(t) \leq 0$ for $-1 \leq t \leq s$. To show $f^{(s)}(t)$ is of positive type for $t_{k-1}^{1,1} \leq s < t_k^{1,0}$, one can check that for such s ,

$$Q_0^{1,0}(s) > 0, \quad \dots, \quad Q_{k-1}^{1,0}(s) > 0, \quad \text{and} \quad Q_k^{1,0}(s) < 0. \tag{5.12}$$

Since $Q_k^{1,0}(t)$ is of positive type, these inequalities show $K_{k-1}^{1,0}(t, s)$ is of positive type as a function of t . Moreover, they also show $(t - s)K_{k-1}^{1,0}(t, s)$ is of positive type as a function of t by using the Christoffel-Darboux formula, which says

$$(t - s)K_{k-1}^{1,0}(t, s) = (c_{k-1}/c_k)(Q_k^{1,0}(t)Q_{k-1}^{1,0}(s) - Q_k^{1,0}(s)Q_{k-1}^{1,0}(t)), \quad (5.13)$$

where $c_k > 0$ is the leading coefficient of $Q_k^{1,0}$. Since the product of functions of positive type is also of positive type, it follows that $f^{(s)}(t)$ is of positive type for $t_{k-1}^{1,1} \leq s < t_k^{1,0}$. By using the property that

$$(t + 1)Q_i^{1,1}(t)Q_j^{1,1}(t) \quad (5.14)$$

is of positive type, this argument can be extended to show $f^{(s)}(t)$ is of positive type for all $s \in [-1, 1]$.

Thus, these function can be used as auxiliary functions for Theorem 5.1, which gives Levenshtein's universal bound for the sphere. In terms of the normalized polynomials $\mathcal{Q}_i(s) = Q_i(s)/Q_i(1)$, we arrive at the bound

$$A(d, s) \leq \begin{cases} \binom{k+d-3}{k-1} \left(\frac{2k+d-3}{d-1} - \frac{\mathcal{Q}_{k-1}(s) - \mathcal{Q}_k(s)}{(1-s)\mathcal{Q}_k(s)} \right) & \text{if } t_{k-1}^{1,1} \leq s < t_k^{1,0}, \text{ and} \\ \binom{k+d-2}{k} \left(\frac{2k+d-1}{d-1} - \frac{(1+s)(\mathcal{Q}_k(s) - \mathcal{Q}_{k+1}(s))}{(1-s)(\mathcal{Q}_k(s) + \mathcal{Q}_{k+1}(s))} \right) & \text{if } t_k^{1,0} \leq s < t_k^{1,1}. \end{cases} \quad (5.15)$$

In certain cases this bound is the best that can be obtained from Theorem 5.1, but in general it does not fully optimize the choice of auxiliary function.

5.3 The Kabatyanskii-Levenshtein bound

The following geometric inequality shows how the sphere packing density $\Delta_{\mathbb{R}^d}$ in \mathbb{R}^d can be bounded using $A(d, s)$:

$$\Delta_{\mathbb{R}^d} \leq \min_{\pi/3 \leq \theta \leq \pi} \sin^d(\theta/2) A(d, \cos \theta) = \min_{-1 \leq s \leq 1/2} \left(\frac{1-s}{2} \right)^{d/2} A(d, s) \quad (5.16)$$

(see (6.9) in [50] or Proposition 2.1 in [20]).

To obtain a good bound for fixed d , this inequality can be combined with Levenshtein's universal bound for $A(d, s)$, where the best value of s can be found by optimizing a piecewise differentiable function. The resulting bound is the one shown as the Kabatyanskii-Levenshtein bound in figure 3, although Kabatyanskii and Levenshtein [19] used a slightly worse bound for $A(d, s)$ as well as for $\Delta_{\mathbb{R}^d}$ in terms of $A(d, s)$.

To obtain an asymptotic bound as $d \rightarrow \infty$, we can use the inequality

$$A(d, s) \leq A(d, t_k^{1,1}) \leq \frac{f^{(t_k^{1,1})}(1)}{f_0^{(t_k^{1,1})}} = 2 \binom{d+k-1}{k} \quad (5.17)$$

when $s \leq t_k^{1,1}$ (see (6.13) in [50]). If $k, d \rightarrow \infty$ with $k/d \rightarrow \alpha$, then

$$t_k^{1,1} \rightarrow \frac{2\sqrt{\alpha(1+\alpha)}}{2\alpha+1} \quad (5.18)$$

by Corollary 5.17 in [50]. For $\theta < \pi/2$, taking

$$\alpha = \frac{1 - \sin \theta}{2 \sin \theta} \tag{5.19}$$

ensures that $t_k^{1,1} \rightarrow \cos \theta$, and applying Stirling's formula shows that

$$\frac{1}{d} \log_2 A(d, \cos \theta) \leq \frac{1 + \sin \theta}{2 \sin \theta} \log_2 \frac{1 + \sin \theta}{2 \sin \theta} - \frac{1 - \sin \theta}{2 \sin \theta} \log_2 \frac{1 - \sin \theta}{2 \sin \theta} + o(1) \tag{5.20}$$

as $d \rightarrow \infty$. It now follows from (5.16) that the sphere packing density is at most $2^{-(\kappa+o(1))d}$, where

$$\kappa = -\log_2 \sin \frac{\theta}{2} - \frac{1 + \sin \theta}{2 \sin \theta} \log_2 \frac{1 + \sin \theta}{2 \sin \theta} + \frac{1 - \sin \theta}{2 \sin \theta} \log_2 \frac{1 - \sin \theta}{2 \sin \theta}. \tag{5.21}$$

Optimizing for the best choice of θ between $\pi/3$ and $\pi/2$ yields the root $\theta = 1.09951240\dots$ of

$$\sec \theta + \tan \theta = e^{(\tan \theta + \sin \theta)/2}, \tag{5.22}$$

at which point we obtain the Kabatyanskii-Levenshtein bound $\kappa = 0.59905576\dots$

Cohn and Zhao [20] gave a general transformation showing that any bound obtained from Theorem 5.1 and (5.16) can also be obtained directly from the Euclidean linear programming bound. Thus, there is no need to use spherical codes to obtain the Kabatyanskii-Levenshtein bound. However, the transformation sheds little additional light on this bound, and it is difficult to see how someone might think of it without using spherical codes.

5.4 Implied kissing numbers

One can strengthen the linear programming bound by taking into account constraints on spherical codes. The following relaxation of the kissing number will prove useful in doing so. Let \mathcal{C} be the set of sphere centers in a packing. For $x \in \mathcal{C}$ and $r \geq 0$, let

$$N_x(r) = \#\{y \in \mathcal{C} : 0 < |x - y| \leq r\}, \tag{5.23}$$

and let $N(r)$ be the average of $N_x(r)$ over $x \in \mathcal{C}$ (we can restrict our attention to periodic packings, so that this average is well defined). If r_0 is the minimal distance in \mathcal{C} , then $N(r_0)$ is the average kissing number of \mathcal{C} , and for $t > 0$ we define the *average t -neighbor number* to be $N(tr_0)$. In other words, the average kissing number is the average 1-neighbor number.

The following strengthening of the linear programming bound is a special case of Theorem 1.4 in [69], where it was used to give improved bounds in dimensions 4 through 7 and 9. The proof in [69] amounts to retaining more terms in the Poisson summation argument from [8].

Theorem 5.2 ([69]). *Let $g: \mathbb{R}^d \rightarrow \mathbb{R}$ be a radial Schwartz function, and suppose that g satisfies the following inequalities for some $\eta \geq 0$ and $s > r > 0$:*

- (1) $g(0) > 0$ and $\widehat{g}(0) > 0$,
- (2) $g(x) \leq \eta g(0)$ for $r \leq |x| \leq s$,

(3) $g(x) \leq 0$ for $|x| \geq s$, and

(4) $\widehat{g}(y) \geq 0$ for all y .

Suppose furthermore that every sphere packing in \mathbb{R}^d has average s/r -neighbor number at most M . Then the sphere packing density in \mathbb{R}^d is at most

$$\frac{\pi^{d/2}}{(d/2)!} \left(\frac{r}{2}\right)^d \frac{g(0)}{\widehat{g}(0)} (1 + \eta M). \quad (5.24)$$

The linear programming bound is equivalent to taking $\eta = 0$. The extra flexibility of being able to choose $g(0)$ and $\widehat{g}(0)$ is irrelevant, because we can rescale g and its input variable, but it will be convenient below.

When the implied kissing number is impossibly large, we can apply Theorem 5.2 to improve on the linear programming bound from Theorem 2.1 as follows. Suppose h is a Schwartz function satisfying the hypotheses of the linear programming bound with radius r . One can check by a rescaling argument that the average kissing number of any sphere packing that achieves this bound must be

$$K = -\frac{d}{rh'(r)}, \quad (5.25)$$

where $h'(r)$ denotes the radial derivative at radius r (see Lemma 5.1 in [57]). Thus, if h is the optimal auxiliary function in the bound, then K must be the implied kissing number.

Suppose furthermore that for some $t > 1$ we can prove an upper bound B for the average t -neighbor number in every packing in \mathbb{R}^d with $B < K$. For example, B could be an upper bound for $A(d, \cos \theta)$ for some $\theta > \pi/3$, which can be arbitrarily close to $\pi/3$.

Given such a bound $B < K$, Theorem 5.2 proves a strictly stronger density bound than Theorem 2.1 using h , as follows. Let $s = (1 + \varepsilon)r$ for some small $\varepsilon > 0$, and define g by $g(x) = h(x/(1 + \varepsilon))$. Then g satisfies the hypotheses of Theorem 5.2 with $\eta = -rh'(r)\varepsilon + O(\varepsilon^2)$, $g(0) = 1$, and

$$\widehat{g}(0) = (1 + \varepsilon)^d = 1 + d\varepsilon + O(\varepsilon^2). \quad (5.26)$$

If ε is small enough, then we can take $M = B$ in Theorem 5.2, and

$$\frac{g(0)}{\widehat{g}(0)} (1 + \eta M) = \frac{1 - rh'(r)B\varepsilon + O(\varepsilon^2)}{1 + d\varepsilon + O(\varepsilon^2)} = 1 + d(B/K - 1)\varepsilon + O(\varepsilon^2), \quad (5.27)$$

which is less than 1 when ε is sufficiently small, because $B < K$. The improvement here is not large, but the resulting density bound is strictly better than that from Theorem 2.1. Thus, figure 16 shows that we can extend the improved density bound from [69] to $10 \leq d \leq 23$, $25 \leq d \leq 179$, and a number of larger values of d . However, we do not know what happens as $d \rightarrow \infty$.

5.5 Bounds for the average kissing number

The implied kissing number has a concrete geometric meaning, beyond being the average kissing number of a hypothetical packing. It turns out to be an upper bound for the average kissing number of any sphere packing, subject to some conjectures about interpolation. The key tool is the following theorem, which is the Euclidean analogue of Proposition 4.1 in [81] by Bourque and Petri.

Theorem 5.3. *Let $f: \mathbb{R}^d \rightarrow \mathbb{R}$ be a radial Schwartz function and $r > 0$, and suppose that f satisfies the following inequalities:*

- (1) $f(r) < 0$,
- (2) $f(x) \leq 0$ for $|x| \geq r$, and
- (3) $\widehat{f}(y) \geq 0$ for all y .

Then the average kissing number of any d -dimensional sphere packing is at most

$$-\frac{f(0)}{f(r)}. \tag{5.28}$$

Here $f(r)$ denotes the value of $f(x)$ when $|x| = r$.

Proof. It suffices to prove the inequality for finite packings and take a limit. Let \mathcal{C} be any finite subset of \mathbb{R}^d with minimal distance r , and let

$$N = \#\{(x, y) \in \mathcal{C}^2 : |x - y| = r\} / |\mathcal{C}| \tag{5.29}$$

be its average kissing number. Then Fourier inversion implies that

$$\sum_{x_1, x_2 \in \mathcal{C}} f(|x_1 - x_2|) = \int_{\mathbb{R}^d} dy \widehat{f}(y) \left| \sum_{x \in \mathcal{C}} e^{2\pi i \langle x, y \rangle} \right|^2 \geq 0, \tag{5.30}$$

while

$$\sum_{x_1, x_2 \in \mathcal{C}} f(|x_1 - x_2|) \leq |\mathcal{C}|f(0) + N|\mathcal{C}|f(r) \tag{5.31}$$

thanks to the inequalities for f . By combining these two bounds, we conclude that $N \leq -f(0)/f(r)$. \square

One can also prove this theorem using Poisson summation, along the lines of [8] or [81]. The conditions for equality are similar to those for the linear programming bound for packing density, if we assume self-duality.¹¹ Specifically, equality holds iff f vanishes at radius $r_n := \sqrt{2\Delta_n}$ for $n \geq 2$, \widehat{f} vanishes at r_n for $n \geq 0$ (with $r_0 = 0$), $r = r_1$, and $f'(r_1) = 0$ (because otherwise shifting r would improve the bound). For comparison, the equality conditions for h in Theorem 2.1 are identical, except that the conditions $f'(r_1) = 0$ and $\widehat{f}(0) = 0$ are replaced with $h(r_1) = 0$.

¹¹This is the same issue as the conjectured agreement between the linear programming bound and the -1 eigenfunction uncertainty principle.

Suppose $r_n = \sqrt{2\Delta_n^{\text{LP}}(d/2)}$ and $d_n = d_n^{\text{LP}}(d/2)$ come from the optimal solution to the linear programming bound in \mathbb{R}^d . The crossing equation says that

$$\sum_{n \geq 0} d_n f(r_n) = \sum_{n \geq 0} d_n \hat{f}(r_n) \tag{5.32}$$

for every radial Schwartz function $f: \mathbb{R}^d \rightarrow \mathbb{R}$. If f satisfies the equality conditions for Theorem 5.3, then this equation reduces to

$$f(0) + d_1 f(r_1) = 0. \tag{5.33}$$

In other words, the bound $-f(0)/f(r_1)$ for the average kissing number is d_1 , as desired.

When should such a function f exist and satisfy the hypotheses of Theorem 5.3? First, note that the condition $\hat{f}(0) = 0$ is redundant, for the following reason. If we let $F(x) = |x|f'(x)$, then $\hat{F}(y) = -d\hat{f}(y) - |y|\hat{f}'(y)$. The other conditions on f guarantee that $F(0) = F(r_n) = \hat{F}(r_n) = 0$ for $n \geq 1$, and then the crossing equation implies that $\hat{F}(0) = 0$ and hence $\hat{f}(0) = 0$. What we need is for f to satisfy the same equality conditions as h , except for changing $h(r_1) = 0$ to $f'(r_1) = 0$.

These conditions arise naturally in interpolation problems [44, 82]. Specifically, Open Problem 7.3 from [44] raised the question of whether radial Schwartz functions $g: \mathbb{R}^d \rightarrow \mathbb{R}$ are uniquely determined by the values and radial derivatives of g and \hat{g} at the radii r_n for $n \geq 1$. While this assertion fails for $d \leq 2$ and is difficult to test for $d = 3$, it seems to hold numerically for $d \geq 4$. Proving or disproving it would be an important step forward in our understanding of the modular bootstrap.

The conditions on f and h mean they are part of an interpolation basis for reconstructing g from these values, since all but one of the values must vanish for f and h . Thus, Theorem 5.3 gives a natural geometric interpretation for one of the basis functions, just as Theorem 2.1 does.

Aside from $d = 8$ or 24 (in which case [44] proves an interpolation theorem), we do not know how to prove that an interpolation basis exists, or that the basis functions satisfy the right sign conditions for these theorems. However, the numerical evidence indicates that both are true. If so, the implied kissing number is an upper bound for the average kissing number of every sphere packing.

This relationship has a pleasing consequence: in each dimension, either the implied kissing number is the best bound known for the average kissing number, or we can use a better bound in Theorem 5.2 to improve on the packing density bound.¹² In other words, if we fail to improve on the linear programming bound for density, it can only be because we have obtained an excellent bound for the average kissing number.

A The limiting case of the spinless modular bootstrap

In this appendix, we explain why Δ_{gap} is an upper bound for the spectral gap in the spinless modular bootstrap even if $\omega(\Phi_0) = 0$ (in the notation of section 2.1). We expect

¹²Technically we need a bound for the average t -neighbor number for some $t > 1$, but that is practically the same as a bound for $t = 1$.

that any such functional ω is the limit of functionals with $\omega(\Phi_0) > 0$, but it is not clear how to justify that expectation. Instead, we can use essentially the same argument as the proof of Proposition 2.4 in [11]. We will translate it into modular bootstrap terms for the convenience of the reader.

First, we note that the spectrum $0 = \Delta_0 < \Delta_1 < \Delta_2 < \dots$ must satisfy

$$\lim_{j \rightarrow \infty} \frac{\Delta_{j+1}}{\Delta_j} = 1, \tag{A.1}$$

since otherwise (4.8) could not hold for $\Delta_j < A < \Delta_{j+1}$ with j large.

Now suppose ω is a linear functional such that $\omega(\Phi_0) \geq 0$, $\omega(\Phi_\Delta) \geq 0$ whenever $\Delta \geq \Delta_{\text{gap}}$, and $\omega(\Phi_\Delta)$ is not identically zero. We wish to obtain a contradiction if $\Delta_1 > \Delta_{\text{gap}}$. The crossing equation

$$\omega(\Phi_0) + \sum_{\Delta > 0} d_\Delta \omega(\Phi_\Delta) = 0 \tag{A.2}$$

shows that $\omega(\Phi_{\Delta_j})$ must vanish for each j , but that does not directly yield a contradiction.

Instead, for each constant $\lambda \geq 1$ we will replace ω with a modified functional ω_λ such that $\omega_\lambda(\Phi_\Delta) = \omega(\Phi_{\lambda\Delta}) + \lambda^{-c}\omega(\Phi_{\Delta/\lambda})$. To see why this is possible, let $f(r) = \omega(\Phi_{r^2/2})$. As explained in section 2, f is a -1 eigenfunction for the radial Fourier transform in $2c$ dimensions, and conversely any such function arises for a suitable choice of ω . Let $f_\lambda(r) = f(\lambda^{1/2}r)$, so that $\widehat{f}_\lambda(r) = \lambda^{-c}\widehat{f}(r/\lambda^{1/2}) = -\lambda^{-c}f(r/\lambda^{1/2})$, and define ω_λ by $\omega_\lambda(\Phi_{r^2/2}) = f_\lambda(r) - \widehat{f}_\lambda(r) = f(\lambda^{1/2}r) + \lambda^{-c}f(r/\lambda^{1/2})$, which works because $f_\lambda - \widehat{f}_\lambda$ is again a -1 eigenfunction.

This new functional satisfies $\omega_\lambda(\Phi_0) \geq 0$ and $\omega_\lambda(\Phi_\Delta) \geq 0$ whenever $\Delta \geq \lambda\Delta_{\text{gap}}$. Furthermore, the only way $\omega_\lambda(\Phi_\Delta)$ can vanish at a point $\Delta \geq \lambda\Delta_{\text{gap}}$ is if $\omega(\Phi_{\lambda\Delta}) = \omega(\Phi_{\Delta/\lambda}) = 0$.

Now suppose $\Delta_1 > \Delta_{\text{gap}}$, and let $\lambda_{\text{gap}} = \Delta_1/\Delta_{\text{gap}}$. Then we conclude that $\omega(\Phi_\Delta) = 0$ whenever Δ is in one of the intervals $(\lambda_{\text{gap}}^{-1}\Delta_j, \lambda_{\text{gap}}\Delta_j)$. Because $\lim_{j \rightarrow \infty} \Delta_{j+1}/\Delta_j = 1$, these intervals cover an entire half-line $[R, \infty)$ for some $R > 0$. However, an eigenfunction of the Fourier transform cannot have compact support unless it vanishes identically, because it must be an entire function if it (equivalently, its Fourier transform) has compact support. Thus, because $\omega(\Phi_\Delta)$ does not vanish identically, we conclude that $\Delta_1 \leq \Delta_{\text{gap}}$, as desired.

B Convergence of the spinless bootstrap

In this appendix, we examine the convergence rate of the spinless modular bootstrap as a function of the truncation order, and in particular explain why we are confident that the numerical calculations in sections 3 and 4 have been fully optimized.

Figure 17 shows the density bounds obtained using truncation orders $N = 1, 2, 4, \dots, 512$ for dimensions $d \leq 2048$, in the same format as figure 3.¹³ Each fixed N seems to lead to the same limit as $d \rightarrow \infty$, analogously to [36, section 3.2], but they closely approximate the optimal linear programming bound over increasingly large ranges of d . In

¹³Technically, when d is small we plot only the smaller values of N , to avoid the failure of our numerical method for $d = 3$ and $N > 32$, which we noted in section 2.4.




Figure 17. Density bounds based on $N = 1, 2, 4, 8, \dots, 512$ (colored red through black).




Figure 18. Minimal N such that $\Delta_1^{\text{LP},N}(c) \leq (1 + \epsilon)\Delta_1^{\text{LP}}(c)$ for $\epsilon = 0.1, 0.01,$ and 0.001 , together with the least squares regression lines.

particular, doubling N more or less doubles the range of dimensions over which we obtain a close approximation.

What figure 17 indicates is that N should be chosen proportionally to d if we wish to obtain comparably accurate results. Figure 18 makes this assertion more precise, and shows that the required truncation order is remarkably close to linear in d . We obtained the limiting values $\Delta_1^{\text{LP}}(c)$ in figure 18 by taking N quite large, in particular more than twice as large as needed to make the plotted values stop changing.

We do not know a formula for the slopes in figure 18. When we need to estimate the convergence rate in high dimensions, we extrapolate from lower dimensions and then make




Figure 19. Convergence of $\Delta_1^{\text{LP},N}(c)$ as a function of N for $c = 1/2, 1, 2, 4, \dots, 256$ (colored red through black).

a conservative underestimate. What makes this procedure reliable is how close to linear figure 18 is. By contrast, figure 19 fixes c and examines how many digits of $\Delta_1^{\text{LP},N}(c)$ have converged as a function of N . Here, the behavior is not nearly as linear, and it is more difficult to extrapolate.

The most delicate numerical estimation in this paper occurs in obtaining the number -0.6044 as the infinite-dimensional limit of the LP column in table 1. Table 4 gives evidence that the values in table 1 are correctly extrapolated to infinite truncation order. Specifically, for each dimension table 4 lists the largest truncation order N we have computed, together with the smallest order N_k that agrees with order N to k decimal places. The numbers in black are exact, meaning that truncation order $N_k - 1$ is not enough. In each such case N is at least $2N_k$, and often much larger than that; this margin of safety gives us confidence that these values do reflect the limit as $N \rightarrow \infty$. The red numbers are obtained by doubling the numbers above them, which seems to produce an overestimate and would work in every other case with $d > 2$. Even for the red numbers, $N_5 < N$, and therefore we believe that our truncation orders are high enough for all the numbers in table 1 to have stabilized.

Acknowledgments

We are grateful to Ganesh Ajjanagadde, Matthew de Courcy-Ireland, Abhinav Kumar, Dalimil Mazáč, Stephen D. Miller, Danylo Radchenko, Leonardo Rastelli, Peter Sarnak,

d	N	N_1	N_2	N_3	N_4	N_5
1	∞	1	1	2	2	2
2	33	1	2	3	4	5
4	308	1	2	4	5	7
8	∞	2	3	5	7	9
16	933	2	4	7	10	13
32	997	3	7	11	15	19
64	1025	5	11	18	25	33
128	780	8	20	33	46	59
256	844	14	37	61	86	111
512	972	25	71	118	167	215
1024	852	48	138	232	327	422
2048	1700	93	271	458	647	836
4096	1750	183	536	916	1294	1672

Table 4. The truncation order N_k required to approximate $(\log_2 \text{density})/d$ to within $0.5 \cdot 10^{-k}$, so that rounding to k decimal places leads to error at most 10^{-k} .

and Maryna Viazovska for helpful conversations. The work of TH and AT is supported by the Simons Foundation (Simons Collaboration on the Nonperturbative Bootstrap). NA is supported by the Leo Kadanoff Fellowship.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] S. Ferrara, A.F. Grillo and R. Gatto, *Tensor representations of conformal algebra and conformally covariant operator product expansion*, *Annals Phys.* **76** (1973) 161 [[INSPIRE](#)].
- [2] A. Polyakov, *Non-Hamiltonian approach to conformal quantum field theory*, *Zh. Eksp. Teor. Fiz.* **66** (1974) 23 [*Sov. Phys. JETP* **39** (1974) 10].
- [3] G. Mack, *Duality in quantum field theory*, *Nucl. Phys. B* **118** (1977) 445 [[INSPIRE](#)].
- [4] R. Rattazzi, V.S. Rychkov, E. Tonni and A. Vichi, *Bounding scalar operator dimensions in 4D CFT*, *JHEP* **12** (2008) 031 [[arXiv:0807.0004](#)] [[INSPIRE](#)].
- [5] S. Rychkov, *EPFL lectures on conformal field theory in $D \geq 3$ dimensions*, SpringerBriefs in Physics (2016), [[DOI](#)] [[arXiv:1601.05000](#)] [[INSPIRE](#)].
- [6] D. Simmons-Duffin, *The conformal bootstrap*, in *Theoretical advanced study institute in elementary particle physics: New frontiers in fields and strings*, pp. 1–74, (2017), [[DOI](#)] [[arXiv:1602.07982](#)] [[INSPIRE](#)].
- [7] D. Poland and D. Simmons-Duffin, *The conformal bootstrap*, *Nature Phys.* **12** (2016) 535.
- [8] H. Cohn and N. Elkies, *New upper bounds on sphere packings I*, *Annals Math.* **157** (2003) 689 [[math/0110009](#)].

- [9] T. Hartman, D. Mazáč and L. Rastelli, *Sphere packing and quantum gravity*, *JHEP* **12** (2019) 048 [[arXiv:1905.01319](#)] [[INSPIRE](#)].
- [10] J. Bourgain, L. Clozel and J.-P. Kahane, *Principe d’Heisenberg et fonctions positives*, *Annales Inst. Fourier* **60** (2010) 1215.
- [11] H. Cohn and F. Gonçalves, *An optimal uncertainty principle in twelve dimensions via modular forms*, *Invent. Math.* **217** (2019) 799 [[arXiv:1712.04438](#)].
- [12] N. Afkhami-Jeddi, T. Hartman and A. Tajdini, *Fast conformal bootstrap and constraints on 3d gravity*, *JHEP* **05** (2019) 087 [[arXiv:1903.06272](#)] [[INSPIRE](#)].
- [13] H. Cohn and A. Kumar, *Optimality and uniqueness of the Leech lattice among lattices*, *Annals Math.* **170** (2009) 1003 [[math/0403263](#)].
- [14] H. Cohn and F. Gonçalves, *Data for “An optimal uncertainty principle in twelve dimensions via modular forms”*, data set, DSpace@MIT, (2018), <https://hdl.handle.net/1721.1/118165>.
- [15] N. Afkhami-Jeddi, H. Cohn, T. Hartman and A. Tajdini, *Free partition functions and an averaged holographic duality*, [arXiv:2006.04839](#) [[INSPIRE](#)].
- [16] K. Ball, *A lower bound for the optimal density of lattice packings*, *Int. Math. Res. Not.* **1992** (1992) 217.
- [17] S. Vance, *Improved sphere packing lower bounds from Hurwitz lattices*, *Adv. Math.* **227** (2011) 2144 [[arXiv:1105.3779](#)].
- [18] A. Venkatesh, *A note on sphere packings in high dimension*, *Int. Math. Res. Not.* **2013** (2013) 1628.
- [19] G.A. Kabatyanskiĭ and V.I. Levenshtein, *Bounds for packings on a sphere and in space*, *Prob. Peredachi Info.* **14** (1978) 3 [*Probl. Info. Transm.* **14** (1978) 1].
- [20] H. Cohn and Y. Zhao, *Sphere packing bounds via spherical codes*, *Duke Math. J.* **163** (2014) 1965 [[arXiv:1212.5966](#)].
- [21] J.H. Conway and N.J.A. Sloane, *Sphere packings, lattices and groups*, vol. 290 of *Grundlehren der Mathematischen Wissenschaften*, third ed., Springer-Verlag, New York, U.S.A., (1999), [[DOI](#)].
- [22] P. Delsarte, *Bounds for unrestricted codes, by linear programming*, *Philips Res. Rept.* **27** (1972) 272.
- [23] R.J. McEliece, E.R. Rodemich, H. Rumsey Jr. and L.R. Welch, *New upper bounds on the rate of a code via the Delsarte-MacWilliams inequalities*, *IEEE Trans. Inform. Theory* **IT-23** (1977) 157.
- [24] A. Barg and D.B. Jaffe, *Numerical results on the asymptotic rate of binary codes*, in *Codes and association schemes (Piscataway, NJ, 1999)*, vol. 56 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pp. 25–32. Amer. Math. Soc., Providence, RI, (2001). [[DOI](#)].
- [25] P. Boyvalenkov, D. Danev and M. Stoyanova, *Refinements of the Levenshtein bounds in q-ary Hamming spaces*, *Prob. Peredachi Info.* **54** (2018) 35 [*Probl. Info. Transm.* **54** (2018) 329].
- [26] D. Mazáč, *Analytic bounds and emergence of AdS₂ physics from the conformal bootstrap*, *JHEP* **04** (2017) 146 [[arXiv:1611.10060](#)] [[INSPIRE](#)].
- [27] T.C. Hales, *A proof of the Kepler conjecture*, *Annals Math.* **162** (2005) 1065.

- [28] T. Hales, M. Adams, G. Bauer, T.D. Dang, J. Harrison, L.T. Hoang, C. Kaliszzyk, V. Magron, S. McLaughlin, T.T. Nguyen, Q.T. Nguyen, T. Nipkow, S. Obua, J. Pleso, J. Rute, A. Solovyev, T.H.A. Ta, N.T. Tran, T.D. Trieu, J. Urban, K. Vu and R. Zumkeller, *A formal proof of the Kepler conjecture*, *Forum Math. Pi* **5** (2017) 1 [[arXiv:1501.02155](#)].
- [29] M.S. Viazovska, *The sphere packing problem in dimension 8*, *Annals Math.* **185** (2017) 991 [[arXiv:1603.04246](#)].
- [30] H. Cohn, A. Kumar, S.D. Miller, D. Radchenko and M. Viazovska, *The sphere packing problem in dimension 24*, *Annals Math.* **185** (2017) 1017 [[arXiv:1603.06518](#)].
- [31] A. Thue, *Om nogle geometrisk-taltheoretiske Theoremer*, *Forhandlingerne ved de Skandinaviske Naturforskeres* **14** (1892) 352.
- [32] T.C. Hales, *Cannonballs and honeycombs*, *Not. Amer. Math. Soc.* **47** (2000) 440.
- [33] S. Collier, Y.-H. Lin and X. Yin, *Modular bootstrap revisited*, *JHEP* **09** (2018) 061 [[arXiv:1608.06241](#)] [[INSPIRE](#)].
- [34] I. Frenkel, J. Lepowsky and A. Meurman, *Vertex operator algebras and the Monster*, vol. 134 of *Pure and Applied Mathematics*, Academic Press, Inc., Boston, MA, U.S.A. (1988).
- [35] S. Hellerman, *A universal inequality for CFT and quantum gravity*, *JHEP* **08** (2011) 130 [[arXiv:0902.2790](#)] [[INSPIRE](#)].
- [36] D. Friedan and C.A. Keller, *Constraints on 2d CFT partition functions*, *JHEP* **10** (2013) 180 [[arXiv:1307.6562](#)] [[INSPIRE](#)].
- [37] T. Gannon, *Moonshine beyond the Monster: The bridge connecting algebra, modular forms and physics*, Cambridge Monographs on Mathematical Physics, Cambridge University Press (2010), [<https://doi.org/10.1017/CBO9780511535116>] [[INSPIRE](#)].
- [38] Y.-H. Lin and S.-H. Shao, *Anomalies and bounds on charged operators*, *Phys. Rev. D* **100** (2019) 025013 [[arXiv:1904.04833](#)] [[INSPIRE](#)].
- [39] C.A. Keller and H. Ooguri, *Modular constraints on Calabi-Yau compactifications*, *Commun. Math. Phys.* **324** (2013) 107 [[arXiv:1209.4649](#)] [[INSPIRE](#)].
- [40] N. Benjamin and Y.-H. Lin, *Lessons from the Ramond sector*, *SciPost Phys.* **9** (2020) 065 [[arXiv:2005.02394](#)] [[INSPIRE](#)].
- [41] H. Cohn and N. Triantafillou, *Dual linear programming bounds for sphere packing via modular forms*, [arXiv:1909.04772](#).
- [42] H. Cohn and A. Kumar, *Universally optimal distribution of points on spheres*, *J. Am. Math. Soc.* **20** (2007) 99 [[math/0607446](#)].
- [43] H. Cohn and M. de Courcy-Ireland, *The Gaussian core model in high dimensions*, *Duke Math. J.* **167** (2018) 2417 [[arXiv:1603.09684](#)].
- [44] H. Cohn, A. Kumar, S.D. Miller, D. Radchenko and M. Viazovska, *Universal optimality of the E_8 and Leech lattices and interpolation formulas*, [arXiv:1902.05438](#).
- [45] H. Cohn and Y. Zhao, *Energy-minimizing error-correcting codes*, *IEEE Trans. Inform. Theory* **60** (2014) 7442 [[arXiv:1212.1913](#)].
- [46] S. El-Showk and M.F. Paulos, *Extremal bootstrapping: go with the flow*, *JHEP* **03** (2018) 148 [[arXiv:1605.08087](#)] [[INSPIRE](#)].

- [47] D. Poland, S. Rychkov and A. Vichi, *The conformal bootstrap: Theory, numerical techniques, and applications*, *Rev. Mod. Phys.* **91** (2019) 015002 [[arXiv:1805.04405](#)] [[INSPIRE](#)].
- [48] V.I. Levenshtein, *On bounds for packings in n -dimensional Euclidean space*, *Dokl. Akad. Nauk SSSR* **245** (1979) 1299 [*Sov. Math. Dokl.* **20** (1979) 417].
- [49] V.I. Levenshtein, *Designs as maximum codes in polynomial metric spaces*, *Acta Appl. Math.* **29** (1992) 1.
- [50] V.I. Levenshtein, *Universal bounds for codes and designs*, in *Handbook of coding theory*, Vol. I, pp. 499–648. North-Holland, Amsterdam, The Netherlands (1998).
- [51] A. Scardicchio, F.H. Stillinger and S. Torquato, *Estimates of the optimal density of sphere packings in high dimensions*, *J. Math. Phys.* **49** (2008) 043301 [[arXiv:0705.1482](#)].
- [52] S. Torquato and F.H. Stillinger, *New conjectural lower bounds on the optimal density of sphere packings*, *Exper. Math.* **15** (2006) 307 [[math/0508381](#)] [<http://projecteuclid.org/euclid.em/1175789761>].
- [53] G. Parisi, P. Urbani and F. Zamponi, *Theory of simple glasses: exact solutions in infinite dimensions*, Cambridge University Press, Cambridge, (2020), [[DOI](#)].
- [54] D. Mazáč and M.F. Paulos, *The analytic functional bootstrap. Part II. Natural bases for the crossing equation*, *JHEP* **02** (2019) 163 [[arXiv:1811.10646](#)] [[INSPIRE](#)].
- [55] L. Rolin and I. Wagner, *A note on Schwartz functions and modular forms*, *Arch. Math. (Basel)* **115** (2020) 35 [[arXiv:1903.05737](#)].
- [56] A.S. Feigenbaum, P.J. Grabner and D.P. Hardin, *Eigenfunctions of the Fourier Transform with specified zeros*, [arXiv:1907.08558](#).
- [57] H. Cohn and S.D. Miller, *Some properties of optimal functions for sphere packing in dimensions 8 and 24*, [arXiv:1603.04759](#).
- [58] W. Gawronski, *On the asymptotic distribution of the zeros of Hermite, Laguerre, and Jonquière polynomials*, *J. Approx. Theory* **50** (1987) 214.
- [59] L. Gatteschi, *Asymptotics and bounds for the zeros of Laguerre polynomials: a survey*, *J. Comput. Appl. Math.* **144** (2002) 7.
- [60] D.V. Widder, *The Laplace transform*, vol. 6 of *Princeton Mathematical Series*, Princeton University Press, Princeton, NJ, U.S.A. (1941).
- [61] J.L. Cardy, *Operator content of two-dimensional conformally invariant theories*, *Nucl. Phys. B* **270** (1986) 186 [[INSPIRE](#)].
- [62] J. Qiao and S. Rychkov, *A tauberian theorem for the conformal bootstrap*, *JHEP* **12** (2017) 119 [[arXiv:1709.00008](#)] [[INSPIRE](#)].
- [63] B. Mukhametzhanov and A. Zhiboedov, *Modular invariance, tauberian theorems and microcanonical entropy*, *JHEP* **10** (2019) 261 [[arXiv:1904.06359](#)] [[INSPIRE](#)].
- [64] S. Ganguly and S. Pal, *Bounds on the density of states and the spectral gap in CFT_2* , *Phys. Rev. D* **101** (2020) 106022 [[arXiv:1905.12636](#)] [[INSPIRE](#)].
- [65] S. Pal and Z. Sun, *Tauberian-Cardy formula with spin*, *JHEP* **01** (2020) 135 [[arXiv:1910.07727](#)] [[INSPIRE](#)].
- [66] B. Mukhametzhanov and S. Pal, *Beurling-Selberg extremization and modular bootstrap at high energies*, *SciPost Phys.* **8** (2020) 088 [[arXiv:2003.14316](#)] [[INSPIRE](#)].

- [67] S. Pal and Z. Sun, *High energy modular bootstrap, global symmetries and defects*, *JHEP* **08** (2020) 064 [[arXiv:2004.12557](#)] [[INSPIRE](#)].
- [68] B. Simon, *Szegő's theorem and its descendants: spectral theory for L^2 perturbations of orthogonal polynomials*, M.B. Porter Lectures, Princeton University Press, Princeton, NJ, U.S.A. (2011).
- [69] D. de Laat, F.M. de Oliveira Filho and F. Vallentin, *Upper bounds for packings of spheres of several radii*, *Forum Math. Sigma* **2** (2014) 1 [[arXiv:1206.2608](#)].
- [70] O.R. Musin, *The kissing number in four dimensions*, *Annals Math.* **168** (2008) 1 [[math/0309430](#)].
- [71] F.C. Machado and F.M. de Oliveira Filho, *Improving the semidefinite programming bound for the kissing number by exploiting polynomial symmetry*, *Exp. Math.* **27** (2018) 362 [[arXiv:1609.05167](#)].
- [72] K. Kallal, T. Kan and E. Wang, *Improved lower bounds for kissing numbers in dimensions 25 through 31*, *SIAM J. Discrete Math.* **31** (2017) 1895 [[arXiv:1608.07270](#)].
- [73] A.D. Wyner, *Capabilities of bounded discrepancy decoding*, *Bell Syst. Tech. J.* **44** (1965) 1061.
- [74] P. Delsarte, J.M. Goethals and J.J. Seidel, *Spherical codes and designs*, *Geom. Dedicata* **6** (1977) 363.
- [75] S. Hellerman and C. Schmidt-Colinet, *Bounds for state degeneracies in 2D conformal field theory*, *JHEP* **08** (2011) 127 [[arXiv:1007.0756](#)] [[INSPIRE](#)].
- [76] S. El-Showk, M.F. Paulos, D. Poland, S. Rychkov, D. Simmons-Duffin and A. Vichi, *Solving the 3d Ising Model with the Conformal Bootstrap II. c-Minimization and Precise Critical Exponents*, *J. Stat. Phys.* **157** (2014) 869 [[arXiv:1403.4545](#)] [[INSPIRE](#)].
- [77] G.E. Andrews, R. Askey and R. Roy, *Special functions*, vol. 71 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, U.K. (1999), [[DOI](#)].
- [78] I.J. Schoenberg, *Positive definite functions on spheres*, *Duke Math. J.* **9** (1942) 96.
- [79] D.V. Shtrom, *The Delsarte method in the problem of the contact numbers of Euclidean spaces of high dimensions*, *Proc. Steklov Inst. Math.* **2** (2002) S162.
- [80] C. Bachoc and F. Vallentin, *New upper bounds for kissing numbers from semidefinite programming*, *J. Am. Math. Soc.* **21** (2008) 909 [[math/0608426](#)].
- [81] M.F. Bourque and B. Petri, *Kissing numbers of closed hyperbolic manifolds*, [[arXiv:1905.11083](#)].
- [82] D. Radchenko and M. Viazovska, *Fourier interpolation on the real line*, *Publ. Math. Inst. Hautes Études Sci.* **129** (2019) 51 [[arXiv:1701.00265](#)].