sphere-packing-talk

David Wakeham

November 1, 2019

Contents

1 Introduction

I'll be talking about **Sphere packing and quantum gravity** (1905.01319) by Hartman, Mazáč and Rastelli. Perhaps a better title would have been "Sphere packing and the modular bootstrap"; the "quantum gravity" is mild clickbait. But the basic idea is that the number of balls you can fit into a large box is related, in a suprising and beautiful way, to the lightest non-vacuum operator in a 2D CFT. When your CFT is holographic, it's dual to some UV complete theory of quantum theory in AdS_3, and this lightest operator gets interpreted as the lightest black hole state. In particular, if you want a UV-complete theory of quantum gravity, you need a black hole below that mass. I'll start by describing the modular bootstrap, which constrains the spectrum of a 2D CFT; then review sphere packing and the recent progress, culminating in solution for $d = 8$ and $d = 24$ dimensions; and finally, we can gloriously unite the two, and I'll comment briefly on the implications for quantum gravity.

2 Modular bootstrap

2.1 Partition functions

We're going to take a 2D CFT, compactify space and put it on a unit radius circle, and finally heat it up so it has period β in imaginary time. Since we have a circle in space multiplied by a circle in imaginary time, the theory lives on a torus, described by some complex modulus $\tau \in \mathbb{H}_+$.

Let's consider a left-moving and a right-moving algebra on the circle, with respective central charges $(c, \bar{c})$. The partition function is a sum over states

$$Z(\tau, \bar{\tau}) = \mathrm{Tr}[q^{L_0 - c/24}\bar{q}^{\bar{L}_0 - \bar{c}/24}] = \sum_{(h,\bar{h})} q^{h-c/24}\bar{q}^{\bar{h}-\bar{c}/24}, \quad q = e^{2\pi i\tau}, \quad \bar{q} = e^{-2\pi i\bar{\tau}}.$$

Less abstractly, $h - c/24$ and $\bar{h} - \bar{c}/24$ are just the energies of the left and right-moving excitation.

We've also introduced a second modular parameter $\bar{\tau}$, which is just the complex conjugate $\bar{\tau} = \tau^*$. But as with z and $\bar{z}$ in complex analysis, we can treat $\bar{\tau} \in \mathbb{H}_-$ as an independent parameter, and the partition function will still converge, and in fact, be holomorphic in both these variables.

Now, geometrically, the same torus is described by any modulus τ' related to τ by an element of the *modular group* $\mathrm{SL}(2, \mathbb{Z})$, the 2×2 integer matrices of unit determinant, generated by the S and T transformations:

$$T : \tau \mapsto \tau + 1, \quad S : \tau \mapsto -\frac{1}{\tau}.$$

Modular invariance of the CFT just means the partition function is invariant under $\mathrm{SL}(2, \mathbb{Z})$:

$$Z\left(\frac{a\tau + b}{c\tau + d}, (\frac{a\bar{\tau} + b}{c\bar{\tau} + d}\right) = Z(\tau, \bar{\tau}), \quad \begin{bmatrix} a & c \\ b & d \end{bmatrix} \in \mathrm{SL}(2, \mathbb{Z}).$$

We have the freedom to choose whatever $\bar{\tau}$ we want, so let's make it the reflection across the real axis, $\bar{\tau} = \tau$. We then define the *spinless partition function*

$$\mathcal{Z}(\tau) = Z(\tau, -\tau) = \sum_{\Delta} q^{\Delta - c/12},$$

so-called because this sets $q = \bar{q}$, and only the scaling dimension $\Delta = h + \bar{h}$ appears in our sum over states. The usual finite temperature partition function, for instance, is given by $\tau = i\beta/2\pi$. Then S invariance is the usual high-low temperature duality, since

$$\tau = \frac{i\beta}{2\pi} \mapsto -\frac{1}{\tau} = \frac{i2\pi}{\beta} = \frac{i}{2\pi}\left(\frac{4\pi^2}{\beta}\right).$$

While our choice gives us a simpler sum over scaling dimension, it breaks the T-symmetry, since $\tau \mapsto \tau + 1, \bar{\tau} \mapsto \bar{\tau} + 1$ is inconsistent with $\bar{\tau} = -\tau$. Only the S-symmetry $\tau \mapsto -1\tau$ remains.

What we're going to do now is expand our partition function into irreps of our symmetry algebra, namely primaries with scaling dimension Δ:

$$\mathcal{Z}(\tau) = \chi_{\mathrm{vac}}(\tau) + \sum_{\Delta > 0} \rho_\Delta \chi_\Delta(\tau).$$

Here, there is some integer degeneracy ρ_Δ of primaries with scaling dimension Δ. This can actually be negative if there are negative norm states, so we implement unitarity by requiring $\rho_\Delta > 0$.

To conclude this whirlwind review, I should actually tell you what symmetries we're considering. For quantum gravity, the relevant algebra is two copies of the Virasoro algebra $\mathrm{Diff}(\mathbb{S}^1)$, with $c = \bar{c}$, which I'll write $V = \mathrm{Vir}_c \times \mathrm{Vir}_c$. In this case, we can explicitly write the characters for our primaries:

$$\chi_\Delta^V(\tau) = \frac{q^{\Delta - c/12}}{\prod_{k\geq 1}(1-q^k)^2} = \frac{q^{\Delta-(c-1)/12}}{\eta(\tau)^2}, \quad \chi_{\mathrm{vac}}^V(\tau) = (1-q)^2\chi_0^V(\tau).$$

Here, $\eta(\tau) = q^{1/24}\prod_k(1-q^k)$ is the Dedekind eta function, arising as usual as the generating function for partitions. Note that the vacuum character is not simply the analytic continuation to $\Delta = 0$.

We can also extend the algebra to include conserved currents. The simplest example is to add c free bosons to the theory, which carry a $U = \mathrm{U}(1)^c \times \mathrm{U}(1)^c$ symmetry. The modular invariant characters, in this case, are

$$\chi_\Delta^U(\tau) = \frac{q^\Delta}{\eta(\tau)^{2c}}, \quad \chi_{\mathrm{vac}}^U(\tau) = \chi_0^U(\tau).$$

These can be derived from the character for the free boson, and the free boson partition on the torus. The Virasoro algebra is a subalgebra of this extended algebra U, since it is just the algebra obeyed by the generators of the stress-energy tensor. I won't go into the details now, but both this result and the various characters can be found in Di Francesco et al.

2.2 Linear programming and eigenfunctionals

The basic idea of the modular bootstrap is to write S-invariance as a constraint on the characters, then project using judiciously some nice basis of linear functionals. We're going to focus on a concrete question: namely, finding upper bounds on the dimension of the lightest, nontrivial primary. This is relevant to holography, for instance, if we want to understand to what extent pure gravity (just gravitons and black holes) is consistent.

So, we first define the antisymmetrised characters

$$\Phi(\tau) = \chi(\tau) - \chi(-1/\tau),$$

so S-invariance becomes

$$\Phi_{\mathrm{vac}}^A(\tau) + \sum_{\Delta>0}\rho_\Delta\Phi_\Delta^A(\tau) = 0$$

for $A \in \{V, U\}$.

Suppose we can choose a linear functional ω which is strictly positive on the vacuum, and nonneative above some cutoff dimension Δ_*, i.e.

$$\omega[\Phi_{\mathrm{vac}}^A] > 0, \quad \omega[\Phi_\Delta^A] \geq 0 \text{ for } \Delta \geq \Delta_*.$$

By linearity, acting on the RHS of the S-invariance constraint gives 0. To balance the positive contribution from the vacuum, we would then clearly need a negative contribution for some $\Delta < \Delta_*$. This tells us that a primary exists below Δ_*.

To say more about these functionals, Freidan and Keller (2013) observed that the S transform of a character can be expressed as a Fourier transform. We'll focus on Virasoro symmetry for simplicity. We can index primaries by vectors $x \in \mathbb{R}^2$, with

$$\Delta(x) = \frac{x^2}{2} + \frac{c-1}{12},$$

and a Gaussian integral (along with the S-transformation property $\eta(-1/\tau) = \sqrt{-i\tau}\eta(\tau)$) gives

$$S \cdot \chi^V_{\Delta(x)} = \int d^2y\, e^{-2\pi i x\cdot y} \chi^V_{\Delta(y)}.$$

The antisymmetrised characters Φ^A_Δ obey a similar property, since the Fourier transform is linear. But by construction, they're eigenfunctions, with eigenvalue -1, of the S transform, simply because they're the character minus the S-image of the character, and $S^2 = I$. It follows that antisymmetrised characters are eigenfunctions of the Fourier transform,

$$\Phi^V_{\Delta(x)} = -\int d^2y\, e^{-2\pi i x\cdot y} \Phi^V_{\Delta(y)},$$

and hence, by linearity, $g(x) = \omega[\Phi^V_{\Delta(x)}]$ is an eigenfunction:

$$\hat{g}(x) = \int d^2y\, e^{-2\pi i x\cdot y} g(y) = \omega \left[\int d^2y\, e^{-2\pi i x\cdot y} \Phi^A_{\Delta(y)} \right] = -\omega[\Phi^A_{\Delta(x)}] = -g(x). \tag{1}$$

Similarly, for the algebra U, you can show that $g(x)$ is a -1 eigenfunction for a $2c$-dimensional Fourier transform, and we index our primaries by $\Delta(x) = x^2/2$ for a vector $x \in \mathbb{R}^{2c}$.

The point of all this is that there is a nice basis of -1 eigenfunctions for the Fourier transform in d dimensions, built out of exponentials and associated Laguerre polynomials:

$$\omega_i(\Delta(x)) = L^{d/2-1}_{2i-1}(4\pi x^2) e^{-2\pi x^2}$$

I can now take linear combinations of these functionals, and try and make the upper bound Δ_* as small as possible. That's the modular bootstrap. I'm not going to go into the details, but let's talk briefly about the results you can obtain this way. Let $\Delta_A(c)$ denotes the smallest value of Δ_* you can obtain for the algebra A and central charge c. For the Virasoro algebra, Hellerman (2009) showed that

$$\Delta_V(c) \leq \frac{c}{6} + O(1).$$

We'll talk about results for the U(1) modular boostrap later.

It turns out that the values $c = 4$ and $c = 12$ are special. The bootstrap produces "optimal" functionals saturing the conditions on ω we specified above. The corresponding bounds on Δ_* are $\Delta_V(4) = 1$ and $\Delta_V(12) = 2$; moreover, these can be numerically reverse-engineered to give a candidate partition function. These partition functions are very special: the numerics suggest they are just theta functions for very special lattices in $2c = 8, 24$ dimensions. The theta function for a lattice $\Lambda \in \mathbb{R}^d$ (the integer span of d vectors) is

$$\Theta_\Lambda(\tau) = \sum_{x\in\Lambda} e^{i\pi\tau x^2}.$$

Then the $c = 4$ partition function involves the E_8 root lattice Λ_8. This has a simple description as vectors in $\mathbb{R}^8$ with integer or half-integer coordinates whose sum is even:

$$\Lambda_8 = \{x \in \mathbb{Z}^8 \cup (\mathbb{Z} + \tfrac{1}{2})^8 : \sum_i x_i \equiv 0 \bmod 2\}.$$

Specifically, we have

$$\mathcal{Z}_4(\tau) = \frac{1}{\eta(\tau)^8}\Theta_{\Lambda_8}(\tau).$$

Similarly, the $c = 12$ partition function involves a special 24-dimensional lattice called the Leech lattice Λ_{24}, which I won't describe explicitly. The partition function is

$$\mathcal{Z}_{12}(\tau) = \frac{1}{\eta(\tau)^{24}}\Theta_{\Lambda_{24}}(\tau) - 24.$$

Both have primaries at nonnegative integers, so for $c = 12$, we have a primary below the general bound. These partition functions are remarkable for a couple of reasons. First, these lattices turn out to provide the densest sphere packings in 8 and 24 dimensions. As I will discuss shortly, this is not a coincidence! Both partition functions are also fully modular-invariant, so we have magically recovered T symmetry. A curious property of the $\mathcal{Z}_{12}$ partition function is that it is precisely the partition function of the *chiral monster CFT*, with $c = 24$ and $\bar{c} = 0$.

3 Sphere packing

3.1 Basics

The appearance of these special lattices suggests that there might be a connection between the modular bootstrap for 2D CFTs and sphere-packing, with central charge playing the role of spatial dimension. To make this connection concrete, we need to learn a bit more about sphere packing.

The basic problem is simple to state: if I have an infinite supply of identical, d-dimensional spheres, what's the most efficient way I can pack them together without overlap? Given a packing, the *density* is just the fraction of space occupied by spheres.

You're not going to win any prizes for solving the 1D problem, since you can just tile the real line with intervals of the same length. So the maximum density is 1. In 2D, the answer is to put spheres on the vertices of a hexagonal lattice, which you can easily show has density $\pi/\sqrt{12}$. In 3D, the optimal packing is to make layers of spheres in a hexagonal lattice, then place them on top of each other with a slight shift, so that the next layer sits in the "dips" of the previous layer. This has density $\pi/\sqrt{18}$. If I do this layering in a regular way, I get either a face-centred cubic lattice, or a hexagonal close-packed lattice, but there are an uncountable number of ways to layer irregularly, and all have the same density.

Proving that these densities are optimal is hard. It's not hard to show they're optimal among lattice packings, which Gauss proved in the 19th century, but it wasn't until 1940 that Toth proved that the hexagonal lattice is optimal among all arrangements, including non-lattice packings, and it wasn't until 1998 that Hales proved (after a few hundred pages) that the hexagonal layers are optimal in 3D. As already mentioned, the E_8 root lattice and Leech lattice Λ_{24} give optimal packings in $d = 8$ and $d = 24$ dimensions using methods I'll outline shortly, but in general, non-lattice packings are optimal.

To simplify the problem, note that we can approximate any non-lattice packing arbitarily well using a *periodic packing*, where we have some arrangement of spheres in the unit cell of a lattice, and we simply periodically repeat that. You can approximate a non-periodic packing using a periodic packing with a large unit cell.

Given some periodic packing on a lattice $\Lambda \in \mathbb{R}^d$, with N spheres of radius r per unit cell, the density is

$$\rho = \frac{NV_d r^d}{|\Lambda|},$$

where $V_d = \pi^{d/2}/\Gamma(d/2+1)$ is the volume of the unit ball, and $|\Lambda|$ is the volume of the unit cell, or equivalently, the determinant of the lattice basis. The "crystal basis" $\{v_1, \ldots, v_N\}$ sets the position of spheres within each unit cell.

The fraction is obviously invariant under rescaling, so for simplicity, we set the diameter of the spheres to be 1. This means the density becomes

$$\rho = \frac{NV_d}{2^d|\Lambda|},$$

and there is a lower bound $|v_i - v_j| \geq 1$ on vectors in the unit cell. After rescaling, we see that a bound on ρ is a bound on $N/|\Lambda|$ for all periodic sets of vectors with unit minimal distance.

3.2 Cohn-Elkies theorem

- To proceed, we need to define a Theta function associated to a periodic packing. This is similar to the theta function for a lattice, but now we perform a (weighted) sum over all distances between spheres in the packing:

$$\Theta(\tau) = \sum_{ij=1}^{N} \sum_{x\in\Lambda} e^{\pi i\tau(x+v_i-v_j)^2} = \sum_{ij=1}^{N} \sum_{x\in\Lambda} q^{(x+v_i-v_j)^2/2}.$$

  We then divide by the η function to get the "partition function" of the packing:

$$\mathcal{Z}(\tau) = \frac{\Theta(\tau)}{\eta(\tau)^d} = \sum_{ij=1}^{N} \sum_{x\in\Lambda} \chi^U_{(x+v_i-v_j)^2/2}(\tau),$$

  where we've combined these powers of q and the η to get $\mathrm{U}(1)^c$ characters for central charge $c = d/2$.

- Let's see what happens after an S transformation. The Poisson summation formula on a lattice states that

$$\sum_{x\in\Lambda} F(x) = \frac{1}{|\Lambda|} \sum_{y\in\Lambda^*} \hat{F}(y),$$

  where Λ^* is the dual lattice,

$$\Lambda^* = \{y \in \mathbb{R}^d : \forall x \in \Lambda, e^{2\pi i x\cdot y} = 1\}.$$

  Applying this to our packing partition function, we learn after some algebra that

$$\mathcal{Z}(\tau) = \frac{1}{|\Lambda|} \sum_{y\in\Lambda^*} \left| \sum_{i=1}^{N} e^{2\pi i v_i\cdot y} \right|^2 \chi^U_{y^2/2}(-1/\tau).$$

  An important point is that the coefficients of the "crossed" (S-transformed) characters are positive.

- Now consider a linear functional ω acting on functions of τ, and define the map $f : \mathbb{R}^d \to \mathbb{R}$ by

$$f(x) = \omega\left[\chi^U_{x^2/2}(\tau)\right].$$

Since the S transform of the character can be written as a d-dimensional Fourier transform, the action of ω on the crossed-channel characters can also be written as a Fourier transform:

$$\omega\left[\chi^U_{x^2/2}(-1/\tau)\right] = \int_{\mathbb{R}^d} d^d x\, e^{-2\pi i x\cdot y} f(x) = \hat{f}(y).$$

Applying ω to the Poisson summation formula above then gives

$$\sum_{ij=1}^N \sum_{x\in\Lambda} f(x + v_i - v_j) = \frac{1}{|\Lambda|}\sum_{y\in\Lambda^*}\left|\sum_{i=1}^N e^{2\pi i v_i\cdot y}\right|^2 \hat{f}(y).$$

In fact, this formula follows directly from the Poisson summation formula, but introducing ω will aid is in drawing parallels to the modular bootstrap shortly.

- To get bounds on density, we subtract off the identity character from both sides. On the LHS, this is given by $x = 0, i = j$, and on the right by $y = 0$. Then, shifting these zero terms onto the LHS, and the rest onto the RHS, we have

$$N^2\left[f(0) - \frac{N}{|\Lambda|}\hat{f}(0)\right] = -\sum_{x\neq 0, i\neq j} f(x + v_i - v_j) + \frac{1}{|\Lambda|}\sum_{y\in\Lambda^*\{0\}}\left|\sum_{i=1}^N e^{2\pi i v_i\cdot y}\right|^2 \hat{f}(y).$$

This is useful because the LHS involves the ratio, $N/|\Lambda|$, we are trying to bound. Suppose we can choose a linear functional ω so that the RHS is positive. More precisely, we would like f to have the following properties:

  - $\hat{f}(y) \geq 0$ for all $y \in \mathbb{R}^d$.
  - $f(x) \leq 0$ for all $|x| \leq 1$.

- Two subtleties. First, the values at $x, y = 0$ are positive by definition. Second, since the argument of f is the distance between sphere centres, and this is at minimum 1, this guarantees the RHS is positive.

- If the RHS is positive, then the LHS is positive. This is equivalent to the celebrated bound of Cohn and Elkies (2016):

$$\frac{N}{|\Lambda|} \leq \frac{f(0)}{\hat{f}(0)} \quad \Longrightarrow \quad \rho \leq \frac{V_d}{2^d}\frac{f(0)}{\hat{f}(0)}.$$

Beautiful!

- Optimising the bound from the Cohn-Elkies method is "just" an infinite-dimensional linear programming problem. If we normalise our linear functional so that $f(0) = 1$, then the optimal bound on density comes from *maximising* $\hat{f}(0)$, subject to the constraints above.

3.3 Nonlinear optimisation and some conjectures

- Solving the LP problem directly is cumbersome. We can make it simpler (at the cost of making it nonlinear) by translating it into eigenfunctions of the Fourier transform where we have an explicit basis, just as for the modular bootstrap.

- Let's decompose f into Fourier-even and -odd parts:

$$f = h - g, \quad \hat{f} = h + g, \quad \hat{g} = -g, \quad \hat{h} = h.$$

  (This is possible for any f, since we can define $2h = \hat{f} + f$, $2g = \hat{f} - f$.)

- We can expand these functions in a basis of Laguerre polynomials, of even and odd degree respectively:

$$g(x) = \sum_{i=1}^{M} \beta_i L^{\nu}_{2i-1}(2\pi|x|^2)e^{-\pi|x|^2}, \quad h(x) = \sum_{i=1}^{M} \alpha_i L^{\nu}_{2i-2}(2\pi|x|^2)e^{-\pi|x|^2},$$

  where $\nu = d/2 - 1$. (Note these form an orthogonal basis for *radial* even and odd eigenfunctions.) The basis is of course infinite-dimensional, but we can solve the LP for any finite M and get a rigorous bound. The bound gets bigger as we make M larger.

- To optimise, we are going to impose a different set of constraints on g and h. We are basically going to massage it into a form which First, let's discuss the odd eigenfunction g. We require it to have a root at 0 and double roots at some points $r_1, \ldots, r_P$, where $P = M/2 - 1$. These give M constraints which fix all β_i coefficients. We try to choose the r_i to *minimise the position of the last sign change* r_0. This fixes g. We then require h to have double roots at the positions r_i, and for the *sum* $g + h$ to have a double root at r_0. These M constraints fix the coefficients α_i of h.

- You might hope that, somehow, there is a sort of "repulsion" between sign changes on g, so that minimising r_0 squeezes out all the others until only one is left. You might hope that imposing a double root in $\hat{f} = g + h$ at r_0 forces it to have constant sign, which we can choose to be positive. In this case, $f = h - g$ has a sign change at r_0 with twice the gradient at r_0, and is therefore nonpositive for $|x| \geq r_0$. If these somewhat miraculous expectations are borne out, the f constructed this way (modulo some rescaling) satisfies the conditions of the Cohn-Elkies theorem.

- And indeed, Cohn and Elkies found that, in all the examples they computed, the nonlinear optimum solved the linear problem! It's not clear that this works in general, it's still a conjecture, but you can produce a candidate f this way, and it always works.

- There was another miracle. In general, the LP bounds need not be sharp; they can be greater than the true maximum. But Cohn and Elkies numerically found that, in 8 and 24 dimensions, the LP bound appeared to be converging to the packing densities of the E_8 root lattice and Leech lattice respectively. If you're interested, these are given by $\pi^4/384$ and $\pi^{12}/12!$ respectively. They conjectured that in the limit $M \to \infty$, where we optimise over the full set of Laguerre polynomials, the LP procedure would converge on "magic functions" f_8 and f_{24} which would precisely agree with the lattice packing densities. That would solve the packing problem in 8 and 24 dimensions.

- Now, a sharp bound requires that f applied to nonzero characters vanishes, or equivalently, the magic functions have roots (in fact, double zeros by the nonlinear construction) at nonzero lattice points and all dual lattice points. For Λ_8, this means

$$f_8(\sqrt{2k}) = \hat{f}_8(\sqrt{2k}) = 0, \quad k = 1, 2, 3, \ldots$$

while for Λ_{24},

$$f_{24}(\sqrt{2k}) = \hat{f}_{24}(\sqrt{2k}) = 0, \quad k = 2, 3, \ldots$$

Note we are using a convention where the shortest vector in Λ_8 has length $r_0 = \sqrt{2}$, while in Λ_{24}, $r_0 = 2$.

3.4 Magic functions

- In fact, if we can build functions with these roots, and which satisfy the LP bounds, they are the magic functions! This is what Maryna Viazovska did for $d = 8$ in 2016, and later in the year, with Cohn, Kumar, Miller, and Radchenko, for $d = 24$. I'm going to sketch the construction, and talk about the ingredients that go into the answer, but we won't derive it.

- So, we start with an ansatz based on the Cohn-Elkies conjecture:

$$h(r) = i\sin^2(\pi r^2/2)\int_0^\infty d\tau\, H_d(\tau)e^{i\pi r^2} \tag{2}$$

$$g(r) = i\sin^2(\pi r^2/2)\int_0^\infty d\tau\, G_d(\tau)e^{i\pi r^2}. \tag{3}$$

Here, big H and big G are functions we have to determine, but we want h and g to be Fourier-even and -odd as before. The role of the $\sin^2$ term is to give double roots at the places we want them, though there will be extra zeros we need to cancel out.

- These properties let us derive functional equations for H and G:

$$H_d(-1/\tau) = -(-i\tau)^{2-d/2}\Delta^{(2)}[H_d(\tau)] \tag{4}$$

$$H_d(\tau+1) = -(-i\tau)^{d/2-2}H_d(-1-1/\tau) \tag{5}$$

$$G_d(-1/\tau) = (-i\tau)^{2-d/2}\Delta^{(2)}[H_d(\tau)] \tag{6}$$

$$G_d(\tau+1) = (-i\tau)^{d/2-2}H_d(-1-1/\tau). \tag{7}$$

Here, $\Delta^{(2)}$ is the *second finite difference operator*

$$\Delta^{(2)}[X(\tau)] = \frac{1}{2}\left[X(\tau+1) + X(\tau-1) - 2X(\tau)\right].$$

- I'm just going skip to the answer. The basic thrust is that you can use known properties of nice, modular-covariant objects called modular forms to solve these functional equations in $d = 8$ and $d = 24$. For $d = 8$, the functions H_8 and G_8 are given by

$$H_8(\tau) = \frac{4\pi\tau^2(E_2E_4 - E_6)^2}{5(E_6^2 - E_4^3)} \tag{8}$$

$$G_8(\tau) = -\frac{32\theta_4^4(4\theta_3^8 - 5\theta_3^4\theta_4^4 + 2\theta_4^8)}{15\pi\theta_3^8\theta_2^8}. \tag{9}$$

The E_k are (weight k) modular forms called *Eiseinstein series*, while the θ_i are *theta functions.* In a similar fashion, you can write down H_{24} and G_{24} for $d = 24$ using Eisenstein series and theta functions, but I'm not going to do it now.

- But the point is that by constructing these magic functions, and showing they produce a density bound which exactly meets the density of the E_8 and Leech lattice packings, we have a direct that these are optimal.

4 Connections

4.1 Isodual lattices

- I've dropped enough hints that you can probably already guess how the modular bootstrap and sphere are related. We have partition functions built out of $\mathrm{U}(1)^c$ characters with arguments in the upper half-plane, and constraints imposed by the S transformation.

- First, let's consider the case of a lattice which is self-dual or *isodual*, which means it's congruent to its dual. The partition function for a lattice is simply

$$\mathcal{Z}(\tau) = \sum_{x\in\Lambda} \chi^U_{x^2/2}(\tau),$$

and in the special case that it's isodual, $\Lambda = \Lambda^*$ and $|\Lambda| = 1$, since the determinant of the lattice and its dual are inverses. It follows the Poisson summation formula that we wrote down that this is precisely S-invariant:

$$\mathcal{Z}(-1/\tau) = \mathcal{Z}(\tau).$$

- Since it's S-invariant, this looks precisely like the partition function for a CFT with $U(1)$ symmetry, with a spectrum given by

$$\Delta(x) = \frac{x^2}{2}, \quad x \in \Lambda.$$

The modular bootstrap instructs us to look for a bound $\Delta^U(c)$ on the lighest nontrivial primary Δ_0. This automatically gives a bound on the shortest vector in an isodual lattice, $L_{\min}$:

$$L_{\min} < \sqrt{2\Delta^U(c)}.$$

Since in my sphere packing, I'll have a sphere at the origin and a sphere a distance $L_{\min}$ away, the radius of spheres I can put into this lattice is bounded by $L_{\min}/2$. Hence, we get a bound on the density of isodual sphere packings in d-dimensions:

$$\rho^{\text{iso}} \le V_d \left(\frac{L_{\min}}{2}\right)^d \le V_d \left[\frac{\Delta^U(d/2)}{2}\right]^{d/2},$$

using $|\Lambda| = 1$.

4.2 Bounds on density

- Isodual lattices are a sort of warm-up for general sphere packing. Let ω be the optimal linear functional for the modular bootrap for the algebra U, giving a bound $\Delta^U(c)$. Recall that, in that context, we defined $g : \mathbb{R}^{2c} \to \mathbb{R}$ as the action of our functional on the "crossed characters":
$$g(x) = \omega[\Phi^U_{\Delta(x)}] = \omega[\chi^U_{\Delta(x)}] - \omega[S \cdot \chi^U_{\Delta(x)}].$$
Given that ω is optimal, it should *saturate* the conditions that we specified earlier. In particular, it should vanish whenever we feed in the vacuum or any x above the lower bound:
$$g(0) = 0, \quad g(x) = 0 \text{ for } |x| \geq \sqrt{2\Delta^U(c)}.$$
Remember that we also proved that g was Fourier-odd, $\hat{g} = -g$, so this is looking similar to the function g that goes into the Cohn-Elkies conjecture. We haven't checked it has double roots or anything, but let's proceed.

- Suppose there is a complementary radial function $h : \mathbb{R}^{2c} \to \mathbb{R}$ which is Fourir-even and satisfies

  1. $h(0) = 0$.
  2. $h(x) \leq g(x)$ for all $|x| > \sqrt{2\Delta^U(c)}$.
  3. $h(x) + g(x) \geq 0$ for all $x \in \mathbb{R}^{2c}$.

- The Cohn-Elkies theorem states that, if such an h exists, then it translates directly into an LP bound on arbitrary sphere packings:
$$\rho \leq V_d \left[\frac{\Delta^U(d/2)}{2}\right]^{d/2}.$$
An important point is that, if you can find a function h that does this, you have the bound. Until you find it, you conjecture it exists, and that the bound holds.

- This is the main technical point I wanted to reach: LP bounds on arbitrary sphere packings are equal to bounds from the modular bootstrap on the lighest nontrivial primary in a $U(1)$ theory, modulo this technical conjecture about the existence of h.

4.3 Large-d asymptotics

- You might say, great, the mathematicians can bound packing densities. But what do I get out of this, i.e. what do we learn about CFTs? It turns out that, in the large c-limit, we can turn the construction around and use asymptotic bounds on the density of sphere packings to learning about the primary bound, $\Delta^U(c)$, at large c.

- Minkowski proved that the maximum density is bounded below by 2^{-d}. If you plug this into our conjecture, and rewrite the volume of a sphere, you end up with a lower bound
$$\Delta^U(c) \geq \frac{\Gamma(c+1)^{1/c}}{2\pi} \approx \frac{c}{2\pi e} + o(c),$$

where the last expression uses Stirling's approximation, and $2\pi e \approx 17$. Using state of the art lower bounds on density at large c, you can improve this (using the conjecture) to

$$\Delta^U(c) \geq \frac{c}{12.57} + o(c).$$

- Similarly, one can use asymptotic upper bounds on the maximum density given by linear programming:
$$\rho^{\text{LP}}_{\text{max}} \lesssim 2^{-0.6d}.$$
Such an upper bound will naturally be an upper bound on the maximum density of isodual lattice packings, once again from linear programming, and we just showed rigorously that the optimal linear programming bounds are equivalent to the modular bootstrap bounds:
$$\rho^{\text{iso}}_m ax = V_d \left[\frac{\Delta^U(d/2)}{2}\right]^{d/2} \leq \rho^{\text{LP}}_{\text{max}} \lesssim 2^{-0.6d} \quad \Longrightarrow \quad \Delta^U(c) \leq \frac{c}{9.795} + o(c).$$
So, the asymptotic bounds on sphere packing give rise to asymptotic bounds on the lighest primary,
$$\frac{c}{12.57} \lesssim \Delta^U(c) \lesssim \frac{c}{9.795},$$
with the upper bound fully rigorous, and the lower bound depending on this conjecture.

4.4 Virasoro

- This connection between sphere packing and the U(1) modular bootstrap is very beautiful and self-contained. But you might be wondering where quantum gravity comes in.

- I don't have time to talk about it in detail, but the basic insight is that the same extremal functionals that you construct for the U(1) modular bootstrap, and which map to linear programming bounds for sphere packing, also work for the Virasoro algebra, roughly because Virasoro is contained in the bigger algebra. This leads to a new upper bound on the lightest nontrivial primary,
$$\Delta^V(c) \lesssim \frac{c}{8},$$
which for a theory dual to quantum gravity in AdS_3, can be viewed as an upper bound on the lightest black holes appearing in the theory. Neato!