Fast conformal bootstrap and constraints on 3d gravity

Nima Afkhami-Jeddi, Thomas Hartman and Amirhossein Tajdini

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

E-mail: na382@cornell.edu, hartman@cornell.edu, at734@cornell.edu

ABSTRACT: The crossing equations of a conformal field theory can be systematically truncated to a finite, closed system of polynomial equations. In certain cases, solutions of the truncated equations place strict bounds on the space of all unitary CFTs. We describe the conditions under which this holds, and use the results to develop a fast algorithm for modular bootstrap in 2d CFT. We then apply it to compute spectral gaps to very high precision, find scaling dimensions for over a thousand operators, and extend the numerical bootstrap to the regime of large central charge, relevant to holography. This leads to new bounds on the spectrum of black holes in three-dimensional gravity. We provide numerical evidence that the asymptotic bound on the spectral gap from spinless modular bootstrap, at large central charge c , is $\Delta_1 \lesssim c/9.1$.

KEYWORDS: Conformal Field Theory, AdS-CFT Correspondence, Models of Quantum Gravity

ARXIV EPRINT: [1903.06272](https://arxiv.org/abs/1903.06272)

Contents

1	Introduction	1
2	Truncating the primal bootstrap	3
2.1	Setup	3
2.2	Truncation	4
2.3	Comments on monotonicity	5
3	Dual bootstrap	5
3.1	Review	6
3.2	Functionals parameterized by zeroes	6
3.3	Duality	8
3.4	From primal solutions to extremal functionals	9
4	Modular bootstrap algorithm	9
4.1	Setup	9
4.2	Algorithm	11
5	Modular bootstrap results	13
5.1	3d gravity and summary of existing bounds	13
5.2	Bound as a function of c	14
5.3	$c = 12$ and the modular j -function	16
5.4	Algorithm benchmarks	16
6	Discussion	17
A	Direct derivation of optimization duality	19
B	Generating guesses for Newton's method	20
C	High precision bounds with linear programming	21

1 Introduction

The conformal bootstrap is a rapidly growing array of techniques to solve strongly interacting quantum field theories. It had early success in two spacetime dimensions [1], and more recently has proved successful in higher dimensions [2]. An important impetus has been the development of numerical techniques to analyze the crossing equations, starting with the introduction of the functional bootstrap method in 2008 by Rattazzi, Rychkov,

Tonni, and Vichi [3]. This method uses linear (or semidefinite [4]) programming to constrain solutions to the crossing equations, thereby carving out the allowed parameter space of unitary CFTs. It has also been used to solve individual CFTs, including the 3d critical Ising model [5, 6]. The method relies on finding extremal functionals, often numerically, which bound the space of CFTs, and in certain cases encode the spectrum of the target theory [7, 8].

Despite this exciting progress, the current numerical methods are in their infancy, and there are many potential areas for improvement, both technical and conceptual. For example, with semidefinite programming, there is no clear way to leverage analytic results, such as the spectrum at high spin, to improve the numerics [9]. This is a tantalizing hint that much more efficient algorithms are waiting to be discovered. Such algorithms, combined with powerful analytics, will be needed to address certain physics problems, such as solving theories with large global symmetries [10, 11], pushing bootstrap into the hydrodynamic regime [12, 13], or exploring the boundaries of quantum gravity by combining numerical bootstrap with AdS/CFT in the black hole regime.

In this paper, we develop a new algorithm for numerical bootstrap, and apply it to 3d gravity. The method is based on solving truncated crossing equations. It builds on one of the extremal functional methods of El-Showk and Paulos [14], and is very similar in many respects, but differs in some important details that extend the reach by several orders of magnitude. In principle, these methods are general, but in practice, they are limited to a certain class of bootstrap problems, where one can easily generate approximate solutions to the truncated equations. We will show that this is possible for the modular bootstrap problem relevant to 3d gravity.

The modular bootstrap is a variant of the conformal bootstrap introduced by Hellerman to constrain the spectrum of 2d CFTs [15]. Hellerman used it to prove analytically that any theory of 3d gravity has a primary operator with scaling dimension $\Delta \lesssim c/6$. From a holographic point of view, this bound is perhaps weaker than expected, because all known theories of 3d gravity have BTZ black holes [16] starting at $\Delta \sim c/12$. The bound can be systematically improved by numerics, but the difficulty grows rapidly with c [17, 18]. Previous numerical calculations were effective only for $c \lesssim 100$ and did not reach the semiclassical regime — the bound had not converged in c , and it was not possible to estimate the asymptotics. A stronger bound is worth pursuing, because if it could be pushed down to the black hole threshold, then the numerics could also potentially be used to construct or rule out candidate theories of 3d gravity, or to suggest new analytic methods.

Using our algorithm, we extend the modular bootstrap up to $c \sim 1800$. This central charge appears to be in the semiclassical regime, where the bound has nearly converged, and we estimate the asymptotic bound to be $\Delta \lesssim c/9.1$. This estimate relies on an extrapolation in c , which appears to be reliable, but strictly speaking it is impossible to rule out a stronger bound as $c \rightarrow \infty$. Our bound is not at the black hole threshold, so unfortunately cannot be used to construct candidate theories of 3d gravity or to address the questions above. We have not included the constraints from spin; it is possible that incorporating spin would lead to a stronger bound that saturates the black hole threshold.

The functional bootstrap is an expansion in derivatives, truncated at some order N_{deriv} . Our calculations are performed up to derivative order $N_{\text{deriv}} \sim 5000$. For a given c , this takes several CPU-hours on a standard laptop. It is impossible to make a direct comparison to existing methods, but a rough extrapolation suggests that semidefinite programming would require at least 10^7 CPU-hours to perform the same calculation. A basic version of our algorithm, sufficient for modular bootstrap, is straightforward to implement in Mathematica with a few dozen lines of code.

On the down side, our methods do not generalize in any obvious way to bootstrap problems with spin. This has (so far) prevented us from applying the algorithm to the spinning modular bootstrap, or to the 3d critical Ising model.

We begin with a review of the crossing equations and a summary of the truncation in section 2. We explain under what circumstances the truncated equations can be used to place general bounds. In section 3 we review the functional method, introduce a new method to optimize bootstrap functionals, and show that it is Lagrange dual to the truncation method. Section 4 describes the algorithm, and in section 5 it is applied to modular bootstrap and 3d gravity.

Comparing our results to semidefinite programming requires very accurate bounds. In a standalone appendix C, we describe a method to find high precision bounds with linear programming, which has some advantages over the standard bisection method.

Although we have stressed the computational aspects, it seems likely that the truncated crossing equations can also be solving analytically in various limits — lightcone [19, 20], large N [21], short distance [22], etc. This could be used to derive analytic bounds, to jumpstart the numerics, or to develop variational methods that go to much higher orders. We consider it likely that with more refinement, numerical methods can be extended to functionals with millions of derivatives. Further discussion of the outlook, including the challenges presented by problems with spin, is in section 6.

2 Truncating the primal bootstrap

The functional bootstrap is an optimization problem, and as such, it has a *primal* and *dual* formulation. The primal bootstrap refers to the crossing equations themselves. The dual approach uses linear functionals. We will rely on both pictures, starting on the primal side. In this section and section 3, we consider a broad class of spinless bootstrap problems. We will specialize later to modular bootstrap.

2.1 Setup

Consider a crossing equation of the form

$$\sum_{\Delta, \ell} a_{\Delta}^{\ell} W_{\Delta}^{\ell}(z, \bar{z}) = 0 \quad (2.1)$$

where ℓ is spin, and $W_{\Delta, \ell}$ is the difference of a conformal block and a crossed block,

$$W_{\Delta}^{\ell}(z, \bar{z}) = g_{\Delta, \ell}(z, \bar{z}) - g_{\Delta, \ell}(1 - z, 1 - \bar{z}) . \quad (2.2)$$

Unitarity requires $a_\Delta^\ell > 0$, and places lower bounds on Δ . Here we use the notation of the correlator bootstrap, with z and $\bar{z}$ conformal cross ratios, but the equations apply also to modular bootstrap, by replacing $(z, 1-z) \rightarrow (\tau, -1/\tau)$ with τ the torus modulus.

Following [3], let us act on (2.1) by linear functionals

$$\mathcal{F} = \sum_{m=0}^M \sum_{\bar{m}=0}^{\bar{M}} \alpha_{m,\bar{m}} \left. \frac{\partial^m}{\partial z^m} \frac{\partial^{\bar{m}}}{\partial \bar{z}^{\bar{m}}} \right|_{z=\bar{z}=1/2} . \quad (2.3)$$

(The methods may also apply to more general bases; see discussion section.) Choose a set $\{\mathcal{F}_i\}_{i=1\dots P}$ of P such functionals, and let

$$f_i^\ell(\Delta) \equiv \mathcal{F}_i[W_\Delta^\ell(z, \bar{z})] , \quad (2.4)$$

where for each ℓ , the f_i^ℓ are linearly independent. Applying $\mathcal{F}_i$ to the crossing equation (2.1) gives

$$\sum_{\Delta,\ell} a_\Delta^\ell f_i^\ell(\Delta) = 0 , \quad (2.5)$$

(assuming the $\mathcal{F}_i$ are sufficiently well behaved to bring the derivatives inside the sum). These are the *primal* bootstrap equations. There are P equations for an infinite number of parameters: the spectrum of (Δ, ℓ) and the coefficients a_Δ^ℓ . After absorbing a positive prefactor into the a_Δ^ℓ , these equations are well approximated by polynomials [2].

We will restrict our attention to bootstrap problems without spin dependence. In this case we set $z = \bar{z}$ (or $\tau = -\bar{\tau}$ in the modular bootstrap), and drop the spin label ℓ . This class of problems includes spinless modular bootstrap, 1d correlators, the $sl(2)$ subsector of correlator bootstrap in higher dimensions.

2.2 Truncation

Let us take P even, and truncate to $P/2+1$ states, including the identity term with $a_0 \equiv 1$, $\Delta_0 \equiv 0$. This results in the *truncated primal equations*:

$$\sum_{\mu=0}^{P/2} a_\mu f_i(\Delta_\mu) = 0 . \quad (2.6)$$

This is now a closed system of P polynomial equations for P unknowns.

This seems a rather brutal truncation of the crossing equations, so at this point, it is not clear that it is useful. Nonetheless we will show that, sometimes, solutions of (2.6) can be used to place bounds on the space of all unitary CFTs.

In the numerical bootstrap, the goal is to maximize some physical quantity, such as the first scaling dimension Δ_1 , over all unitary solutions to the P equations (2.5). In practice, the optimal solution has a finite number of states in the sum. If it happens that the optimal solution has exactly $P/2+1$ states, then it can be found by solving the closed equations (2.6). We will show that this is the case for 2d modular bootstrap at zero angular potential.

This leads to our main analytic result:

In modular bootstrap, at any order P and central charge c , there is a solution to the truncated bootstrap equations (2.6), with scalar gap $\Delta_1^{(c,P)}$, such that all unitary CFTs of central charge c have $\Delta_1 \leq \Delta_1^{(c,P)}$.

It follows that general bounds on all unitary CFTs can be found by solving (2.6). To derive this, we will show that a solution to (2.6) can be used to build an extremal functional, identical to that obtained in the usual numerical bootstrap. Note that this claim does not apply to all solutions of (2.6), but to one special solution, which must be carefully sought among many.

This is based on the same idea used for numerical bootstrap of 1d correlators in [14]. Other approaches to the correlator bootstrap using truncated crossing equations have also been discussed in [23–25].

2.3 Comments on monotonicity

The largest dimension participating in (2.6), $\Delta_{P/2}$, grows with P . Therefore the truncation order P acts as a UV cutoff. Qualitatively, there is a parallel between changing $P \rightarrow P - 2$ and a renormalization group transformation: as P is decreased, the effect of high-dimension operators must be absorbed into the low-dimension spectrum and OPE coefficients, in order to preserve the crossing equations. A corollary of the result stated above is monotonicity:

Under reducing the cutoff $P \rightarrow P - 2$, the gap $\Delta_1^{(c,P)}$ is non-decreasing.

Roughly speaking, we can view a solution to (2.6) as a bootstrap version of an effective field theory, in the following sense. For $P \gg 1$, the low-lying operators Δ_μ are insensitive to the truncation, and take their physical values for the target CFT. For $\mu \sim P/2$, the operators of dimension Δ_μ are not true operators in the full-fledged CFT, but are effective operators in the cutoff theory, at least for the purposes of the crossing equation. Like the higher dimension operators in a Wilsonian effective Lagrangian, they parameterize all of the contributions from the UV, and very good approximations to the low energy physics can be obtained with a finite number of effective operators. This point of view is useful to keep in mind when interpreting solutions to the truncated equations, because we should not expect operators with $\mu \sim P/2$ to have their physical scaling dimensions; they must also encode the effects of the UV states that are removed by the truncation. We will see this effect in the numerics.

3 Dual bootstrap

In this section, we derive the claims of section 2, by relating the primal equations to the functional bootstrap. We will first recast the functional bootstrap as a nonlinear optimization problem, then dualize to the primal equations.

3.1 Review

The functional method [3] relies on the following observation. Consider the action of a linear functional on W_Δ ,

$$f(\Delta) = \sum_{i=1}^P \alpha_i f_i(\Delta) . \quad (3.1)$$

The α_i parameterize the functional. If we find a functional that is positive on the vacuum, and on all dimensions above some gap,

$$f(0) > 0 \quad \text{and} \quad f(\Delta) > 0 \text{ for } \Delta > \Delta_1 \quad (3.2)$$

then every term in the sum (2.1) is positive for dimensions in this range. Therefore, to satisfy crossing, every unitary CFT must have an operator at or below Δ_1 .

The strongest bounds come from minimizing Δ_1 . For our purposes, a useful way to state the optimization procedure is as follows. Fix P , the gap Δ_1 , and the normalization condition $\alpha_P = 1$, and:¹

$$\begin{aligned} &\text{maximize } f(0) \text{ over } \alpha_a \\ &\text{subject to } f(\Delta) \geq 0 \text{ for } \Delta \geq \Delta_1 . \end{aligned} \quad (3.3)$$

Here and below, the indices run over the ranges

$$\begin{aligned} a, b, c, \dots &= 1 \dots P-1 , \\ i, j, k, \dots &= 1 \dots P . \end{aligned} \quad (3.4)$$

Eq. (3.3) is a linear program with an infinite number of constraints, or can be recast as a semidefinite program [4]. If the constraints are feasible, and the objective $f(0)$ is positive at the maximum, then the initial choice of Δ_1 is ruled out. This procedure is repeated, for different values of Δ_1 , until the theory is marginally excluded, i.e. $f(0) = 0$ (to some desired precision). This is the edge of the exclusion region, and the optimal functional at this point is called the extremal functional. The *dual bootstrap* refers to the problem of finding the extremal functional.

3.2 Functionals parameterized by zeroes

The zeroes of the extremal functional correspond to solutions of the primal equations (2.5). This is a general fact of linear programs, which has been used to find the spectrum of various CFTs, including the critical Ising model in two [8] and three [6] dimensions.

In certain spinless bootstrap problems, including modular bootstrap, the extremal functional has a very simple pattern of zeroes. Empirically, for even P , there are single zeroes at $\Delta = 0$ and $\Delta = \Delta_1$, and double zeroes at $P/2 - 1$ additional states. An example is plotted in figure 1. This set of $P/2 + 1$ zeroes corresponds to a solution of the truncated primal equations (2.6).

¹Generally, the optimum may occur either for $\alpha_P > 0$ or $\alpha_P < 0$. We choose the overall sign of the functionals such that it is always positive.

Figure 1. Example of an extremal functional for modular bootstrap, with $c = 12$ and truncation order $P = 6$. There are single zeroes at $\Delta_0 \approx 0$ and $\Delta_1 \approx 2.13$, and double zeroes at $\Delta_2 \approx 3.43$ and $\Delta_3 \approx 5.13$. The additional single zero near $\Delta \approx 1$ plays no role in the discussion. (The root near the origin is slightly shifted due to null state contributions described in section 4.1, but this is too small to be visible in the plot.)

We will make an ansatz for the functional by assuming that this pattern of zeroes continues at higher P . We will confirm *a posteriori* that this assumption holds for modular bootstrap, at least for all c and P that we consider.

Given this assumption about the pattern of zeroes, we now replace the linear optimization problem (3.3) by a nonlinear optimization. Choose a spectrum Δ_μ for $\mu = 1 \dots P/2$. From this spectrum, construct the unique functional obeying

$$\begin{aligned} f(\Delta_1) &= 0, \\ f(\Delta_\mu) &= f'(\Delta_\mu) = 0 \quad (\mu = 2 \dots P/2). \end{aligned} \tag{3.5}$$

This is $P - 1$ linear equations for the $P - 1$ variables α_a , so the solution is unique. In other words, we can view the functional as parameterized by its roots. Next, the linear optimization (3.3) is replaced by:

$$\text{maximize } f(0) \text{ over } \Delta_\mu, \text{ with } \alpha_a \text{ determined by (3.5)}. \tag{3.6}$$

The optimization step can also be eliminated altogether by solving directly for the extremal functional. This has $f(0) = 0$ and is an extremum with respect to Δ_μ . So we simply append to (3.5) the equations

$$f(0) = 0, \quad \frac{d}{d\Delta_\mu} f(0) = 0, \tag{3.7}$$

where in the latter equation, $\mu = 2 \dots P/2$, and the gradient is taken along solutions to (3.5). That is, $\alpha_a = \alpha_a(\Delta_\mu)$ as determined by (3.5).

This is now a closed system of equations for the extremal functional. There is no guarantee that the positivity conditions will be obeyed in this formulation, but this is easily checked after the functional has been constructed, and holds in the applications we will consider. A similar system of equations, but for the linear version of the problem and involving also the primal variables, appears in [14].

To summarize, we have formulated the problem of finding the extremal functional as $3P/2 - 1$ equations (3.5) and (3.7) for the $P/2$ dimensions Δ_μ and $P - 1$ functional parameters α_a . In the derivative basis, the equations are polynomial (after absorbing a positive Δ -dependent factor into the coefficients a_μ).

This reduces the problem of finding the extremal functional to the problem of finding zeroes of a system of polynomial equations. The degree of the polynomials grows with P , so this is non-trivial. Moreover, we have replaced a convex optimization by a non-convex root-finding problem, where there may be runaways and spurious solutions. In practice, these obstacles can be overcome if we can generate a good guess to use in the first step of Newton's method. We will describe how to do this for modular bootstrap in section 4.

3.3 Duality

As the nomenclature suggests, the truncated versions of the primal and dual bootstrap are, in fact, Lagrange duals. With our nonlinear formulation of the dual problem, this duality holds under the assumption that the double roots Δ_μ of the extremal functional have multiplicity exactly two — not higher.

The duality immediately proves the main results stated in section 2. In the dual bootstrap, monotonicity of the gap under changing P is obvious — if we allow a more general functional, the bound on Δ_1 can only get stronger. When translated into primal language, this monotonicity property becomes more surprising.

The duality that we describe is already very well known in the linear formulation of the problem. We will rederive it in terms of the nonlinear version above.

We will start with the dual bootstrap — i.e. the problem of finding extremal functionals (3.6) — then dualize to the primal bootstrap (2.6). The constrained optimization (3.6) is equivalent to maximizing the Lagrangian

$$L = \alpha_i f_i(0) + a_\mu \alpha_i f_i(\Delta_\mu) + \lambda_\nu \alpha_i f'_i(\Delta_\nu) \quad (3.8)$$

over the independent variables α_a, Δ_μ and the Lagrange multipliers a_μ, λ_ν (with $\mu = 1 \dots P/2$, $\nu = 2 \dots P/2$.) Using the normalization condition $\alpha_P = 1$ and rearranging terms,

$$L = E_P + \alpha_a E_a \quad (3.9)$$

with

$$E_i \equiv f_i(0) + a_\mu f_i(\Delta_\mu) + \lambda_\nu f'_i(\Delta_\nu) . \quad (3.10)$$

To dualize, we reinterpret this Lagrangian, viewing $\Delta_\mu, a_\mu, \lambda_\nu$ as independent variables, and α_a as a Lagrange multiplier. Assuming $f''_i(\Delta_\mu) \neq 0$ so that no roots have multiplicity

higher than two, the equation of motion for Δ_μ imposes $\lambda_\nu = 0$. Therefore, the primal optimization problem is

$$\begin{aligned} &\text{maximize } f_P(0) + a_\mu f_P(\Delta_\mu) \text{ over } \Delta_\mu, a_\mu \\ &\text{subject to } f_a(0) + a_\mu f_a(\Delta_\mu) = 0. \end{aligned} \quad (3.11)$$

These are the same sums as appear in the primal crossing equation (2.6), so the constraints consist of the first $P-1$ primal equations. Finally, let us also impose extremality, $f(0) = 0$, so the assumed gap is marginally excluded. This is equivalent to the P th primal equation.

The conclusion is that the problem of finding the extremal functional is precisely dual to the problem of solving the P truncated crossing equations, (2.6), for the spectrum and coefficients. In appendix A, we derive the same statement by direct analysis of the extremal functional equations (3.5)–(3.7), and in the process, give a formula for the truncated OPE coefficients a_μ in terms of the extremal functional.

3.4 From primal solutions to extremal functionals

Given a solution to the primal equations (2.6) for some P , it is straightforward (and computationally trivial) to construct the dual extremal functional: fix Δ_μ to the primal values, and solve the linear equations (3.5) for the functional α_a . A consistency check is that this functional should vanish on the identity, $f(0) = 0$.

Once a candidate extremal functional has been obtained this way, it can be plotted to confirm the positivity conditions (3.2). If (3.2) is satisfied, then the functional implies rigorous bounds, so in the end, our bounds are not conditioned on any assumptions about the pattern of zeroes.

It is also guaranteed that the bounds we derive are optimal, in the sense that linear programming at the same value of P will always produce exactly the same bounds. It is impossible to derive a stronger bound (at a given P), because doing so would contradict the primal solution that we have found. In other words, the method we have described produces solutions which are primal-dual feasible, and therefore optimal.

4 Modular bootstrap algorithm

So far, our discussion has applied to a general class of spinless bootstrap problems, with a certain pattern of zeroes in the extremal functional. We now focus on the modular bootstrap exclusively. In this section we first review the (spinless) modular bootstrap [15, 17], then describe our algorithm in detail.

4.1 Setup

Consider the partition function of a 2d CFT at zero angular potential,

$$Z(\beta) = \sum_{\text{states}} e^{-2\pi\beta(\Delta - \frac{c}{12})} \quad (4.1)$$

$$= \sum_{\text{primaries}} \chi_\Delta(\beta) \quad (4.2)$$

with $\beta > 0$, and c the central charge. In the second line we have organized the sum into characters of the algebra, $\text{Virasoro} \times \text{Virasoro}$.² Crossing symmetry in this context is modular invariance: $Z(\beta) = Z(1/\beta)$. It is convenient to define the reduced partition function

$$\hat{Z}(\beta) = |\eta(i\beta)|^2 |i\beta|^{1/2} Z(\beta) \quad (4.3)$$

$$= \sum_{\mu=0}^{\infty} \rho_{\mu} G_{\Delta_{\mu}}(\beta) . \quad (4.4)$$

where ρ_{μ} is the degeneracy of primaries with dimension Δ_{μ} . This function is also modular invariant, $\hat{Z}(\beta) = \hat{Z}(1/\beta)$. The vacuum is $\Delta_0 = 0$, $\rho_0 = 1$. The Dedekind η -functions cancel similar factors in χ_{Δ} , such that in (4.4), the blocks for non-vacuum states are

$$G_{\Delta}(\beta) = \beta^{1/2} \exp \left[-2\pi\beta \left(\Delta - \frac{c-1}{12} \right) \right] , \quad (4.5)$$

and, for the vacuum,

$$G_0(\beta) = \beta^{1/2} e^{2\pi\beta \frac{c-1}{12}} (1 - e^{-2\pi\beta})^2 . \quad (4.6)$$

(The extra factors in the vacuum block are to account for the null state at level one, $L_{-1}|0\rangle = 0$.) The crossing equation (2.1) is

$$\sum_{\mu=0}^{\infty} \rho_{\mu} W_{\Delta_{\mu}}(\beta) = 0 \quad (4.7)$$

$$W_{\Delta}(\beta) \equiv G_{\Delta}(\beta) - G_{\Delta} \left(\frac{1}{\beta} \right) .$$

We choose the basis of derivative functionals

$$\mathcal{F}_k = \frac{1}{2(2k-1)!} \left[\frac{1}{2} (1 + \beta)^2 \partial_{\beta} \right]^{2k-1} \frac{1 + \beta}{2\sqrt{\beta}} \Big|_{\beta=1} , \quad (4.8)$$

for $k = 1 \dots P$. This is a convenient basis because the resulting f_k are simply odd-index Laguerre polynomials, L_{2k-1} . Indeed, acting with $\mathcal{F}_k$ on the crossing equation (4.7) gives³

$$\sum_{\mu=0}^{\infty} a_{\mu} f_k(\Delta_{\mu}) = 0 , \quad (4.10)$$

where

$$a_{\mu} \equiv \rho_{\mu} e^{-2\pi\Delta_{\mu}} , \quad (4.11)$$

² $\chi_{\Delta}(\beta) \equiv (\chi_{\Delta/2}^{\text{Vir}}(i\beta))^2$, with $\chi_h^{\text{Vir}}(\tau)$ the standard chiral Virasoro character.

³ *Derivation:* the generating function for Laguerre polynomials is

$$\sum_{p \geq 0} t^p L_p(2w) e^{-w} = (1-t)^{-1} \exp \left(-\frac{1+t}{1-t} w \right) \quad (4.9)$$

Therefore acting with $\frac{1}{p!} (\partial_t)^p|_{t=0}$ on the right-hand side gives $L_p(2w) e^{-w}$. With the identification $\beta = \frac{1+t}{1-t}$, this is the same as $\mathcal{F}_k[W_{\Delta}(\beta)]$. The even-index Laguerre's do not appear because the corresponding differential operator vanishes on the modular-odd function $W_{\Delta}(\beta)$.

and, for $\Delta > 0$,

$$f_k(\Delta) = L_{2k-1}(4\pi x), \quad x \equiv \Delta - \frac{c-1}{12}. \quad (4.12)$$

The vacuum state is different due to the extra factors in (4.6),

$$f_k(0) = L_{2k-1}(4\pi x_0) - 2e^{-2\pi} L_{2k-1}(4\pi(x_0 + 1)) + e^{-4\pi} L_{2k-1}(4\pi(x_0 + 2)), \quad (4.13)$$

with $x_0 = -\frac{c-1}{12}$.

At this point, we have written the primal bootstrap equations in the same format as section 2. The truncated version, a closed system of P equations for P unknowns, is

$$\sum_{\mu=0}^{P/2} a_\mu f_k(\Delta_\mu) = 0, \quad (4.14)$$

4.2 Algorithm

We now turn to numerical algorithms based on the reformulation of the crossing problem described in sections 2–3. We focus on the simplest bootstrap question, which is to place an upper bound on the scalar gap Δ_1 , given the central charge c . The algorithms can be modified to accommodate similar problems.

We can consider two main approaches:

- *Dual*: solve the functional equations (3.5) and (3.7), for the spectrum and the functional.
- *Primal*: solve the P truncated primal crossing equations, (2.6), for the spectrum and degeneracies.

Note that standard optimization methods for functional bootstrap require an additional scan over Δ_1 . In our case, Δ_1 is a free parameter, so this step is not necessary.

For modular bootstrap, both the dual and primal algorithms are much faster than semidefinite programming. In the rest of this paper we will consider only the primal algorithm, which was somewhat faster in our preliminary testing, and conceptually simpler.

The purpose of the algorithm is to solve the equations (4.14) for some large value of P . The strategy is to start at small P , solve the equations by Newton’s method, then gradually increase P , using the previous results to generate a good initial guess for Newton’s method. The same strategy was used for 1d correlator bootstrap in [14] (specifically, our algorithm is similar to the ‘upgrading + error correction’ method). Besides the differences between 1d correlator bootstrap and 2d modular bootstrap, our algorithm has two significant improvements: better numerical stability, and a more elaborate, more accurate method to guess the initial points. This will allow us to increase P in large increments.

For numerical stability, we first introduce normalization factors to rescale the $P \times P/2$ matrix $f_k(\Delta_\mu)$. Rewrite (4.14) as

$$\sum_{\mu=0}^{P/2} \tilde{a}_\mu \tilde{f}_k(\Delta_\mu) = 0, \quad (4.15)$$

Figure 2. Spectrum of the $c = 12$ solution at truncation order $P = 110$ (left) and $P = 200$ (right). The horizontal axis is the state number, $\mu = 1, 2, \dots, P/2$. Note that the curves are approximately the same shape. This observation is used to generate very accurate guesses for the initial point in Newton’s method.

where

$$\tilde{a}_\mu = a_\mu \left| \max_k f_k(\Delta_\mu) \right| \quad (4.16)$$

$$\tilde{f}_k(\Delta_\mu) = \frac{f_k(\Delta_\mu)}{f_k(0) \left| \max_k f_k(\Delta_\mu) \right|} . \quad (4.17)$$

The matrix $\tilde{f}_k(\Delta_\mu)$ is scaled such that the vacuum column is all 1’s, and every other column has an absolute maximum entry of ± 1 .

We now apply Newton’s method to the equations (4.15) for the P unknowns $(\tilde{a}_\mu, \Delta_\mu)$. The Laguerre polynomials, and their derivatives, are calculated by 3-term recurrence relations. This is faster and more numerically stable than direct evaluation. Our implementation of Newton’s method is completely standard so we will not review it.

An important element in the algorithm is a guess generator, to provide the initial point in Newton’s method. The method is largely insensitive to the initial guesses for $\tilde{a}_\mu$, so for these coefficients we always guess $\tilde{a}_\mu = 1$. Convergence is very sensitive to the starting Δ_μ , so these must be chosen carefully. We start by solving the equations for $P = 4, 6$. For these initial points, it is easy to find a good guess by hand (or using a linear/semidefinite optimizer such as SDPB [26]). Then we proceed upwards in P recursively, using previous results to generate a guess for Δ_μ .

Our guess generator is based on the observation that the spectrum, at different values of P , is self-similar. That is, the curve Δ_μ as a (discrete) function of μ has almost exactly the same shape for all P , once P is large enough. This is illustrated by the plots in figure 2, which are the exact spectra at $P = 110$ and $P = 200$ for the $c = 12$ modular bootstrap. The shapes in the two plots are almost identical; only the scaling of the axes and the state spacing differs. This motivates a guess where we rescale the lower- P spectrum, and re-sample at the appropriate points. This generates a good guess, and we use precisely this method for $P \lesssim 200$. However, we can do better, and take larger P -steps, by correcting for the fact that the spectrum is not quite self-similar. We do this by parameterizing the shape

curve, fitting the parameters as functions of P , and extrapolating. The full algorithm is given in appendix B.

The resulting guesses are very accurate. For example, in the $c = 100$ modular bootstrap, using data from $P = 120, 122, 124, \dots, 200$, the extrapolation to $P = 280$ produces a spectrum $\Delta_\mu^{\text{guess}}$ that differs from the exact spectrum by under one part in 10^3 – 10^5 . Results at higher P are comparable. This means that we never need to run more than a handful of Newton steps to converge to an accurate spectrum.

This completes the algorithm. The only computationally expensive step is inverting the Jacobian in Newton's method, which scales roughly as P^3 . The computations are performed in Mathematica with high precision arithmetic.

5 Modular bootstrap results

5.1 3d gravity and summary of existing bounds

Three-dimensional gravity in the semiclassical limit is holographically dual to a 2d CFT with a large central charge, $c \gg 1$. Quantum fields in the bulk, aside from the graviton, are dual to low-dimension operators in the CFT. It follows that the CFT should also have a sparse spectrum of low-dimension operators, in that the number of states below any fixed Δ_* should be finite as $c \rightarrow \infty$.

'Pure' 3d gravity, which consists only of the graviton plus black holes, would have no new primaries between the vacuum and black hole threshold. That is, $\Delta_1 \sim \frac{c}{12}$. It is unknown whether pure gravity exists as a quantum theory. A putative partition function was suggested by Witten [27], but explicit calculations by Maloney and Witten [28] gave a different answer that was incompatible with a holographic interpretation. It is possible that this calculation can be modified to produce a consistent quantum theory [29], but at this point, the realm of possibilities is poorly understood.

Even the simplest question is unanswered: what is the largest gap Δ_1 compatible with diffeomorphism invariance? Large diffeomorphisms in 3d gravity correspond to modular transformations in the dual CFT, so this is a question for the modular bootstrap.

The functional approach to modular bootstrap was initiated in [15], where, working analytically at $P = 2$, it was proved that $\Delta_1 \leq \frac{c}{6} + 0.474$. This is known as the Hellerman bound. At large c , the Hellerman bound is a factor of 2 above the black hole threshold. It has also been shown analytically that, for $\Delta \gtrsim \frac{c}{6}$, 2d CFTs with a large gap have the same entropy as BTZ black holes [30]. However, the range $\frac{c}{12} \leq \Delta \leq \frac{c}{6}$ remains enigmatic. This range certainly does not have a universal spectrum, even in holographic theories, but it is an open question whether there must be any states at all within this window.

Numerical evidence strongly suggests that the asymptotic bootstrap bound can be improved to somewhere between $c/12$ and $c/6$ [17, 18, 31]. (For related work, see [32–46]). For $c \lesssim 150$, numerical bounds are as strong as $\Delta_1 \lesssim c/8.2$. However, $c \sim 150$ is the largest central charge accessible to standard numerical methods, and at this value, it is clear that the bound has not yet converged. Results at larger c are needed to find the bound in the semiclassical, large- c limit.

Figure 3. Upper bound on Δ_1/c , as a function of c . Dots are numerical data for truncation at $P = 270, 310, 350, 390, 470, 550, 630, 710, 870, 1030, 1350, 1670, 1990, 2310$, from top to bottom. The solid red line is the extrapolation to $P = \infty$. The dashed blue line is the asymptotic estimate at large c , $1/9.08$.

The difficulty is that for large c , the bound on Δ_1 converges slowly with P , so large c requires large P . This is a general obstacle to bootstrapping quantum gravity in the regime $\Delta \sim c$, which presents a significant challenge to probing black hole physics with the bootstrap. (On the other hand, many interesting questions in AdS/CFT already arise for states with $\Delta \ll c$, and in these cases numerical bootstrap has been very successful.)

The largest truncation order considered in the literature on modular bootstrap is $P = 92$ [18], using the semidefinite program solver SDPB [26]. We apply our algorithm up to $P \sim 2250$ (derivative order 4500). At small P , we have checked that our bounds agree exactly with those produced by SDPB.

5.2 Bound as a function of c

The upper bound on Δ_1/c is plotted in figure 3. The dots are numerical data, and the solid line is the extrapolation to $P = \infty$ with an exponential fit, $\Delta_1(P) = \Delta_1^\infty + ae^{-bP}$.

To estimate the asymptotic bound, we fit the data in the range $300 \leq c \leq 700$ to a function of c . In this range, the bound has converged very close to its $P \rightarrow \infty$ limit. The estimate depends on what function is used in the fit, but natural choices all give estimates clustered around the same value,

$$\Delta_1 \lesssim c/9.1. \quad (5.1)$$

Of course this is not definitive, because we are working at finite c , but the errors in the fit are very small. Different choices of fitting function, together with the adjusted R^2 of the fit,

Fitting Function	$1 - R_{\text{adj}}^2$	δ_{900}	δ_{1800}	Best fit A
$c/A + B_1$	10^{-7}	9×10^{-4}	3×10^{-3}	9.01
$c/A + B_1 + B_2/c$	10^{-9}	2×10^{-4}	9×10^{-4}	9.04
$c/A + B_1 + B_2 \log c$	6×10^{-11}	4×10^{-6}	2×10^{-5}	9.08

Table 1. Fitting the bound as a function of c .

Figure 4. The numerical spectrum Δ_μ for $c = 500$ and $P = 2310$. In the larger plot we show every 15th scaling dimension to reduce clutter. The smaller inset figure shows the first 30 scaling dimensions with deviation away from linearity. In [47], the linear regime of the spectrum with $\Delta_\mu = \frac{c-4}{8} + \mu$ (dashed line) will be explained analytically.

are reported in table 1. We also give an error estimate δ_c , which is defined as the fractional error of our two extrapolations. For example, δ_{900} compares the extrapolation in c (which was found using data only from $c \leq 700$) to our actual data at $c = 900$, extrapolated up to $P = \infty$. The errors are smallest for the logarithmic fit,

$$\Delta_1^{\max} \approx \frac{c}{9.08} + 0.439 \log c - 0.896. \quad (5.2)$$

The best strict (non-extrapolated) bound on Δ_1/c is $c = 1800, P = 2310$, for which $\Delta_1^{\max} = c/8.956$.

In addition to the gap, Δ_1 , the method also provides a spectrum with over a thousand operators. As an example, the spectrum for $c = 500, P = 2310$ is plotted in figure 4. As discussed in section 2.3, only the operators $\mu \ll P$ should be trusted, but in practice, roughly the lower half of the spectrum with $\mu \lesssim P/4$ is converged.

5.3 $c = 12$ and the modular j -function

The case $c = 12$ is especially interesting, because the numerics appears to converge to a spectrum with integer scaling dimensions $\Delta_\mu = 2, 3, \dots$, and integer degeneracies. (The case $c = 4$ has a similar story.) The bound on the gap at $c = 12$, obtained with $P = 2000$ in a few hours of CPU time, is:

$$\Delta_1 < 2 + 1 \times 10^{-30} \quad (c = 12, P = 2000) . \quad (5.3)$$

The numerical solution of the truncated crossing equations at $P = 2000$ is

$$\begin{aligned} \hat{Z} \approx q^{1/12} \beta^{1/2} & \left[q^{-1} (1 - q)^2 + (196882 + 2 \times 10^{-24}) q \right. \\ & \left. + (21099994 + 4 \times 10^{-22}) q^2 + (821115567 + 3 \times 10^{-20}) q^3 + \dots \right] \end{aligned} \quad (5.4)$$

where $q \equiv e^{-2\pi\beta}$. The degeneracies converge to integers that suggest a link to the modular j -function. Indeed, reintroducing the η -functions removed in (4.3), this corresponds to

$$Z \approx j(\tau) - 744, \quad \tau = i\beta . \quad (5.5)$$

This coincidence will be explained analytically in [47]. Note that we did not impose invariance under $\tau \rightarrow \tau + 1$, so it is surprising to find the j -function appearing.

5.4 Algorithm benchmarks

Runtime is compared to semidefinite programming with SDPB (v1) [26] in figure 5. Extrapolating the runtime of SDPB to large P may be unreliable, but naively fitting to a power law and setting $P = 2000$ gives the estimate $\sim 10^9$ s.

We can also compare to the extremal functional method of El-Showk and Paulos [14]. They solve a different problem (1d correlators) so it is not a direct comparison, but they report runtime of 45 minutes at $P = 150$. For modular bootstrap, our algorithm runs in seconds at the same P . The algorithm is very similar to that described in [14], so we do not know the exact origin of the speedup, but it presumably comes from some combination of improvements to numerical stability, the guess generator described in section 4.2, or inherent differences in modular bootstrap.⁴

Our procedure is numerically stable, in the sense that round-off errors do not accumulate dramatically. This means that the same accuracy as semidefinite programming is achieved with much lower precision at intermediate steps. (This gives a significant speedup, but is not the main factor.) For example, at $c = 12$, machine precision (16 decimal digits) is sufficient up to $P \sim 90$, and 60 digits is sufficient for $P \sim 2000$. Calculations at larger c require higher precision, because more terms participate significantly in the crossing equations. The data in figure 3 was produced with a range of precisions from 16 to 500 digits, chosen to ensure that the matrix inversion step in Newton's method is well behaved.

⁴It is also not clear to us whether the calculations in [14] were done using the primal equations alone, or the primal-dual formulation of the problem. We use only the primal equations.

Figure 5. Runtime of SDPB and our Newton-based algorithm for modular bootstrap at $c=12$, vs the number of polynomials, P . Larger c requires more precision, so runtimes are somewhat longer.

6 Discussion

In summary, we have described a reformulation of the modular bootstrap equations, amenable to both analytic and numerical analysis. On the analytic side, the main logic was to show that (i) extremal functionals can be efficiently parameterized by their roots, and (ii) this is dual to solving the truncated crossing equations.

For modular bootstrap, our algorithm is orders of magnitude faster than semidefinite programming or any other previous method, extending the maximum truncation order from $P \sim 100$ to $P \sim 2500$. Other potential advantages include:

- It is numerically stable, so intermediate steps do not require very high precision.
- The method might apply to non-derivative or non-polynomial functionals. (One choice has been considered [48].) We restricted to derivative functionals only because in this basis, we have confirmed explicitly (in examples) the pattern of zeroes invoked in section 3. The method applies to any basis of functionals with a similar pattern of zeroes.
- The algorithm can be sped up significantly by hot starting, i.e. inputting a guess for the spectrum of the target CFT. This means that analytic results can potentially be leveraged for faster numerics, similar to the way we used results at small P for faster convergence at large P .

- Although the strict bounds we derived apply only to unitary CFTs, the truncated primal equations may also have solutions corresponding to non-unitary CFTs, with potential applications to weakly first order phase transitions [2, 49]. This would be similar to the approach taken in [23, 24].
- The truncated bootstrap equations (2.6) can be studied analytically. This may lend additional insights.

There is also a major drawback to our approach, at least in our current implementation. We considered only bootstrap problems without spin dependence. More generally, for example in the modular bootstrap with angular potential, or the correlator bootstrap of the 3d Ising model, there are more constraints on the extremal functional, labeled by ℓ . This typically leads to a different pattern of zeros, not the simple pattern we observed here. In primal language, the problem is that with spin, the solution of the primal optimization problem at truncation order P typically does not have exactly $P/2 + 1$ states.⁵

This makes it impossible to truncate to a closed, primal-only problem, as we have done here. Instead one can write primal-dual equations for the extremal functional, involving $(\alpha_a, a_\mu, \Delta_\mu)$ simultaneously, as in [14]. In principle, we could use Newton's method to solve this enlarged problem, but in practice, we have not found any efficient way to generate initial guesses. The pattern of states does not change predictably as the truncation order is increased.

As described in the introduction, part of our motivation in this work was to push the numerical bootstrap to the point where it can be used to construct, or rule out, candidate theories of 3d gravity. However, our bound $\Delta_1 \lesssim c/9.1$ does not saturate the black hole threshold, so it appears that spinless modular bootstrap is not powerful enough to attack these problems. It might be possible with spinning modular bootstrap, but spin requires new methods.

Eventually, it may be possible to use numerics to settle longstanding questions in quantum gravity: do all theories of quantum gravity require new states well below the Planck scale, such as strings or Kaluza-Klein modes (see e.g. [52])? Does quantum gravity in the ultraviolet rule out a swampland of effective theories in the infrared [53]? Is there a quantum theory of pure gravity in three dimensions [27]? Versions of all these questions, in anti-de Sitter, can be posed in terms of the space of solutions to the crossing equations. Of course it would be better to solve these problems analytically, but definitive numerical answers would be a good start.

Acknowledgments

We are grateful to David Simmons-Duffin for discussions of semidefinite and simplex methods and to Dalimil Mazáč and Leonardo Rastelli for comments on the draft. We also thank

⁵The simple pattern of zeros in the spinless modular bootstrap may be related to total positivity of the functionals, which was noted in [50]. Totally positive functionals lead to polynomials $f_i(\Delta)$ with interlacing roots [51]. This fact seems to give the polynomials the right shape to be able to produce one new double root with the addition of two new polynomials.

other members of the Simons Bootstrap Collaboration, and the organizers of the bootstrap workshops at ICTS in 2017, and the Azores and Caltech in 2018, as well as the Aspen Center for Physics, where much of this work was done or inspired. This work is supported by Simons Foundation grant 488643.

A Direct derivation of optimization duality

In this section, we re-derive the duality of section 3.3 by direct analysis of the extrema of the functional optimization (3.6). This is useful in practice for translating between dual and primal solutions. Denote

$$f_{i\mu} \equiv f_i(\Delta_\mu), \quad f'_{i\mu} \equiv f'_i(\Delta_\mu), \quad f''_{i\mu} \equiv f''_i(\Delta_\mu) \quad (\text{A.1})$$

and form the block matrix

$$V_{i\beta} = [f_{i\mu} \quad f'_{i\mu}] \quad (\text{A.2})$$

where $i = 1 \dots P$, $\mu = 1 \dots P/2$, $\nu = 2 \dots P/2$, and $\beta = 1 \dots P - 1$. The roots (3.5), using (3.1), impose

$$\alpha_i V_{i\beta} = \alpha_a V_{a\beta} + V_{P\beta} = 0, \quad (\text{A.3})$$

where we have used the normalization condition $\alpha_P = 1$. Therefore, given Δ_μ , the functional with the correct zeroes is

$$\alpha_a = -V_{P\beta} V_{\beta a}^{-1}, \quad (\text{A.4})$$

where V^{-1} denotes the inverse of the square submatrix $V_{a\beta}$ (recall a is restricted to the range $1 \dots P - 1$). The functional evaluated on the identity term is

$$f(0) = \alpha_i f_i(0) = f_P(0) - V_{P\beta} V_{\beta a}^{-1} f_a(0). \quad (\text{A.5})$$

According to (3.6), we seek to maximize this over Δ_μ . Ultimately we are interested in a maximum with $f(0) = 0$, so we can assume the maximum is actually obtained, and therefore occurs at a critical point. The gradient, with $\nabla_\mu \equiv \frac{d}{d\Delta_\mu}$, is

$$\nabla_\mu f(0) = -(\nabla_\mu V_{P\beta}) V_{\beta a}^{-1} f_a(0) + V_{P\beta} V_{\beta b}^{-1} (\nabla_\mu V_{b\gamma}) V_{\gamma a}^{-1} f_a(0) \quad (\text{A.6})$$

$$= -\alpha_i (\nabla_\mu V_{i\beta}) V_{\beta a}^{-1} f_a(0) \quad (\text{A.7})$$

$$= -\alpha_i f'_{i\mu} V_{\mu a}^{-1} f_a(0) - \alpha_i f''_{i\mu} V_{\mu+P/2-1,a}^{-1} f_a(0) \quad (\text{A.8})$$

$$= -\alpha_i f''_{i\mu} V_{\mu+P/2-1,a}^{-1} f_a(0) \quad (\text{A.9})$$

Here $\mu = 2 \dots P/2$ is not summed. The first line uses the matrix identity $\nabla X^{-1} = -X^{-1} \nabla X X^{-1}$; the second line repackages the two terms into $i = 1 \dots P$; the third line uses (A.2); and the last line follows from (A.3).

Therefore we have two options: $f''_{i\mu} = 0$, so some Δ_μ is actually a triple rather than double root, or

$$V_{\beta a}^{-1} f_a(0) = 0 \quad (\beta = P/2 + 1 \dots P - 1) \quad (\text{A.10})$$

We assume the latter, so that no roots have multiplicity higher than two. To understand the resulting constraint, define $c_\beta = V_{\beta a}^{-1} f_a(0)$. Eq. (A.10) states that the vector c_β takes the block form

$$c_\beta = [-a_\mu \ 0_\nu], \quad (\text{A.11})$$

where $\mu = 1 \dots P/2$ and 0_ν is a vector of $P/2 - 1$ zeros. Inverting gives

$$f_a(0) = -a_\mu f_{a\mu}. \quad (\text{A.12})$$

Therefore, letting $a_0 = 1$ and $\Delta_0 = 0$, we have

$$\sum_{\mu=0}^{P/2} a_\mu f_a(\Delta_\mu) = 0. \quad (\text{A.13})$$

This is precisely the first $P - 1$ equations of the primal bootstrap, (2.6), with a_μ the OPE coefficient. Now let us impose extremality, i.e. that we are at the boundary of the excluded parameter space. That is, we let Δ_1 be an independent variable and impose

$$0 = f(0) = \alpha_a f_a(0) + f_P(0) = -V_{P\beta} V_{\beta a}^{-1} f_a(0) + f_P(0) = -V_{P\beta} c_\beta + f_P(0) = V_{P\mu} a_\mu + f_P(0). \quad (\text{A.14})$$

This is the P th equation of (2.6).

Therefore, the solution of (3.6), at extremality, is equivalent to a solution of the P primal bootstrap equations (2.6) (assuming no roots of multiplicity greater than two). We have also shown that, given the extremal functional, the OPE coefficients in the truncated primal solution are

$$a_\mu = -V_{\mu a}^{-1} f_a(0). \quad (\text{A.15})$$

B Generating guesses for Newton's method

The logic of our guess generator is described in section 4.2. We fit to the curves in figure 2, extrapolate in P to generate a new curve, then sample at the appropriate discrete points. In this appendix we describe the complete algorithm. The choice of parameters and fitting functions that we use is ad hoc, but produces good results.

Suppose we have already found the spectrum for various truncations $P = P_A$, $A = 1, \dots, K - 1$. Denote these spectra by $\Delta_\mu^{P_A}$. Define $\Delta^{P_A}(x)$ to be the continuous function that numerically interpolates this spectrum at integer x , i.e.

$$\Delta^{P_A}(\mu) = \Delta_\mu^{P_A} \quad \text{for } \mu = 1, 2, \dots, P_A/2. \quad (\text{B.1})$$

We use a degree-22 Hermite interpolation (in Mathematica: `Interpolation[... , InterpolationOrder->22]`). Now suppose we want to guess the spectrum at truncation order P_K . Let

$$F_\mu^{P_A} = \Delta^{P_A} \left(\frac{\mu P_A}{P_K} \right). \quad (\text{B.2})$$

This is defined at discrete P_A . Fit this this, as a function of P_A , to the continuous function

$$F_\mu(y) = a_1 + a_2 y + a_3 \log(y) + a_4 \log(y)^2 + a_5 \log(y)^3. \quad (\text{B.3})$$

The parameters a_i are chosen such that $F_\mu(P_A) \approx F_\mu^{P_A}$, with least-squares fitting. Finally, our guess for the spectrum at $P = P_K$ is

$$\Delta_\mu^{\text{guess}} = F_\mu(P_K) . \quad (\text{B.4})$$

C High precision bounds with linear programming

In this appendix we describe a simple method to find high-precision bounds on the gap, Δ_1 , with linear programming. This is used when we compare our results to SDPB, but it is more general. For example, it is faster than the usual method to find bounds on the 3d Ising model (the ‘kink’ in [5]), if the bounds are needed to high accuracy.

The standard method is to run a program solver such as SDPB in feasibility mode, to rule in or out each gap Δ_1 , and to zoom in on the optimal Δ_1 using bisection. Instead, we run the linear program (3.3). Call the objective at the optimum $M(\Delta_1)$. To find the optimal bound, i.e. the gap Δ_1 which is marginally excluded, we use the secant method to solve $M(\Delta_1) = 0$.

In the first few steps, this method is slower, because the program solver takes longer to solve an optimization problem than a feasibility problem. But the secant method converges exponentially faster than bisection, so for high accuracy bounds on Δ_1 , this method is faster.

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] A.A. Belavin, A.M. Polyakov and A.B. Zamolodchikov, *Infinite Conformal Symmetry in Two-Dimensional Quantum Field Theory*, *Nucl. Phys. B* **241** (1984) 333.
- [2] D. Poland, S. Rychkov and A. Vichi, *The Conformal Bootstrap: Theory, Numerical Techniques and Applications*, *Rev. Mod. Phys.* **91** (2019) 15002 [[arXiv:1805.04405](https://arxiv.org/abs/1805.04405)] [[INSPIRE](#)].
- [3] R. Rattazzi, V.S. Rychkov, E. Tonni and A. Vichi, *Bounding scalar operator dimensions in 4D CFT*, *JHEP* **12** (2008) 031 [[arXiv:0807.0004](https://arxiv.org/abs/0807.0004)] [[INSPIRE](#)].
- [4] D. Poland, D. Simmons-Duffin and A. Vichi, *Carving Out the Space of 4D CFTs*, *JHEP* **05** (2012) 110 [[arXiv:1109.5176](https://arxiv.org/abs/1109.5176)] [[INSPIRE](#)].
- [5] 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*, *Phys. Rev. D* **86** (2012) 025022 [[arXiv:1203.6064](https://arxiv.org/abs/1203.6064)] [[INSPIRE](#)].
- [6] 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](https://arxiv.org/abs/1403.4545)] [[INSPIRE](#)].
- [7] D. Poland and D. Simmons-Duffin, *Bounds on 4D Conformal and Superconformal Field Theories*, *JHEP* **05** (2011) 017 [[arXiv:1009.2087](https://arxiv.org/abs/1009.2087)] [[INSPIRE](#)].

- [8] S. El-Showk and M.F. Paulos, *Bootstrapping Conformal Field Theories with the Extremal Functional Method*, *Phys. Rev. Lett.* **111** (2013) 241601 [[arXiv:1211.2810](#)] [[INSPIRE](#)].
- [9] D. Simmons-Duffin, *The Lightcone Bootstrap and the Spectrum of the 3d Ising CFT*, *JHEP* **03** (2017) 086 [[arXiv:1612.08471](#)] [[INSPIRE](#)].
- [10] F. Kos, D. Poland, D. Simmons-Duffin and A. Vichi, *Precision Islands in the Ising and $O(N)$ Models*, *JHEP* **08** (2016) 036 [[arXiv:1603.04436](#)] [[INSPIRE](#)].
- [11] S.M. Chester and S.S. Pufu, *Towards bootstrapping QED_3* , *JHEP* **08** (2016) 019 [[arXiv:1601.03476](#)] [[INSPIRE](#)].
- [12] L. Iliesiu, M. Kologlu, R. Mahajan, E. Perlmutter and D. Simmons-Duffin, *The Conformal Bootstrap at Finite Temperature*, *JHEP* **10** (2018) 070 [[arXiv:1802.10266](#)] [[INSPIRE](#)].
- [13] L.V. Delacrétaz, T. Hartman, S.A. Hartnoll and A. Lewkowycz, *Thermalization, Viscosity and the Averaged Null Energy Condition*, *JHEP* **10** (2018) 028 [[arXiv:1805.04194](#)] [[INSPIRE](#)].
- [14] S. El-Showk and M.F. Paulos, *Extremal bootstrapping: go with the flow*, *JHEP* **03** (2018) 148 [[arXiv:1605.08087](#)] [[INSPIRE](#)].
- [15] S. Hellerman, *A Universal Inequality for CFT and Quantum Gravity*, *JHEP* **08** (2011) 130 [[arXiv:0902.2790](#)] [[INSPIRE](#)].
- [16] M. Bañados, C. Teitelboim and J. Zanelli, *The Black hole in three-dimensional space-time*, *Phys. Rev. Lett.* **69** (1992) 1849 [[hep-th/9204099](#)] [[INSPIRE](#)].
- [17] D. Friedan and C.A. Keller, *Constraints on 2d CFT partition functions*, *JHEP* **10** (2013) 180 [[arXiv:1307.6562](#)] [[INSPIRE](#)].
- [18] S. Collier, Y.-H. Lin and X. Yin, *Modular Bootstrap Revisited*, *JHEP* **09** (2018) 061 [[arXiv:1608.06241](#)] [[INSPIRE](#)].
- [19] Z. Komargodski and A. Zhiboedov, *Convexity and Liberation at Large Spin*, *JHEP* **11** (2013) 140 [[arXiv:1212.4103](#)] [[INSPIRE](#)].
- [20] A.L. Fitzpatrick, J. Kaplan, D. Poland and D. Simmons-Duffin, *The Analytic Bootstrap and AdS Superhorizon Locality*, *JHEP* **12** (2013) 004 [[arXiv:1212.3616](#)] [[INSPIRE](#)].
- [21] I. Heemskerk, J. Penedones, J. Polchinski and J. Sully, *Holography from Conformal Field Theory*, *JHEP* **10** (2009) 079 [[arXiv:0907.0151](#)] [[INSPIRE](#)].
- [22] B. Mukhametzhanov and A. Zhiboedov, *Analytic Euclidean Bootstrap*, [arXiv:1808.03212](#) [[INSPIRE](#)].
- [23] F. Gliozzi, *More constraining conformal bootstrap*, *Phys. Rev. Lett.* **111** (2013) 161602 [[arXiv:1307.3111](#)] [[INSPIRE](#)].
- [24] F. Gliozzi and A. Rago, *Critical exponents of the 3d Ising and related models from Conformal Bootstrap*, *JHEP* **10** (2014) 042 [[arXiv:1403.6003](#)] [[INSPIRE](#)].
- [25] W. Li, *New method for the conformal bootstrap with OPE truncations*, [arXiv:1711.09075](#) [[INSPIRE](#)].
- [26] D. Simmons-Duffin, *A Semidefinite Program Solver for the Conformal Bootstrap*, *JHEP* **06** (2015) 174 [[arXiv:1502.02033](#)] [[INSPIRE](#)].
- [27] E. Witten, *Three-Dimensional Gravity Revisited*, [arXiv:0706.3359](#) [[INSPIRE](#)].

- [28] A. Maloney and E. Witten, *Quantum Gravity Partition Functions in Three Dimensions*, *JHEP* **02** (2010) 029 [[arXiv:0712.0155](#)] [[INSPIRE](#)].
- [29] C.A. Keller and A. Maloney, *Poincaré Series, 3D Gravity and CFT Spectroscopy*, *JHEP* **02** (2015) 080 [[arXiv:1407.6008](#)] [[INSPIRE](#)].
- [30] T. Hartman, C.A. Keller and B. Stoica, *Universal Spectrum of 2d Conformal Field Theory in the Large c Limit*, *JHEP* **09** (2014) 118 [[arXiv:1405.5137](#)] [[INSPIRE](#)].
- [31] J.D. Qualls and A.D. Shapere, *Bounds on Operator Dimensions in 2D Conformal Field Theories*, *JHEP* **05** (2014) 091 [[arXiv:1312.0038](#)] [[INSPIRE](#)].
- [32] S. Hellerman and C. Schmidt-Colinet, *Bounds for State Degeneracies in 2D Conformal Field Theory*, *JHEP* **08** (2011) 127 [[arXiv:1007.0756](#)] [[INSPIRE](#)].
- [33] C.A. Keller and H. Ooguri, *Modular Constraints on Calabi-Yau Compactifications*, *Commun. Math. Phys.* **324** (2013) 107 [[arXiv:1209.4649](#)] [[INSPIRE](#)].
- [34] C.-M. Chang and Y.-H. Lin, *Bootstrapping 2D CFTs in the Semiclassical Limit*, *JHEP* **08** (2016) 056 [[arXiv:1510.02464](#)] [[INSPIRE](#)].
- [35] E. Shaghoulian, *Modular forms and a generalized Cardy formula in higher dimensions*, *Phys. Rev. D* **93** (2016) 126005 [[arXiv:1508.02728](#)] [[INSPIRE](#)].
- [36] H. Kim, P. Kravchuk and H. Ooguri, *Reflections on Conformal Spectra*, *JHEP* **04** (2016) 184 [[arXiv:1510.08772](#)] [[INSPIRE](#)].
- [37] N. Benjamin, E. Dyer, A.L. Fitzpatrick and S. Kachru, *Universal Bounds on Charged States in 2d CFT and 3d Gravity*, *JHEP* **08** (2016) 041 [[arXiv:1603.09745](#)] [[INSPIRE](#)].
- [38] D. Das, S. Datta and S. Pal, *Charged structure constants from modularity*, *JHEP* **11** (2017) 183 [[arXiv:1706.04612](#)] [[INSPIRE](#)].
- [39] M. Cho, S. Collier and X. Yin, *Genus Two Modular Bootstrap*, *JHEP* **04** (2019) 022 [[arXiv:1705.05865](#)] [[INSPIRE](#)].
- [40] E. Dyer, A.L. Fitzpatrick and Y. Xin, *Constraints on Flavored 2d CFT Partition Functions*, *JHEP* **02** (2018) 148 [[arXiv:1709.01533](#)] [[INSPIRE](#)].
- [41] L. Apolo, *Bounds on CFTs with W_3 algebras and AdS_3 higher spin theories*, *Phys. Rev. D* **96** (2017) 086003 [[arXiv:1705.10402](#)] [[INSPIRE](#)].
- [42] N. Afkhami-Jeddi, K. Colville, T. Hartman, A. Maloney and E. Perlmutter, *Constraints on higher spin CFT_2* , *JHEP* **05** (2018) 092 [[arXiv:1707.07717](#)] [[INSPIRE](#)].
- [43] J.-B. Bae, S. Lee and J. Song, *Modular Constraints on Conformal Field Theories with Currents*, *JHEP* **12** (2017) 045 [[arXiv:1708.08815](#)] [[INSPIRE](#)].
- [44] S. Collier, P. Kravchuk, Y.-H. Lin and X. Yin, *Bootstrapping the Spectral Function: On the Uniqueness of Liouville and the Universality of BTZ*, *JHEP* **09** (2018) 150 [[arXiv:1702.00423](#)] [[INSPIRE](#)].
- [45] J.-B. Bae, S. Lee and J. Song, *Modular Constraints on Superconformal Field Theories*, *JHEP* **01** (2019) 209 [[arXiv:1811.00976](#)] [[INSPIRE](#)].
- [46] T. Anous, R. Mahajan and E. Shaghoulian, *Parity and the modular bootstrap*, *SciPost Phys.* **5** (2018) 022 [[arXiv:1803.04938](#)] [[INSPIRE](#)].
- [47] T. Hartman, D. Mazac and L. Rastelli, to appear.

- [48] A. Castedo Echeverri, B. von Harling and M. Serone, *The Effective Bootstrap*, *JHEP* **09** (2016) 097 [[arXiv:1606.02771](#)] [[INSPIRE](#)].
- [49] V. Gorbenko, S. Rychkov and B. Zan, *Walking, Weak first-order transitions and Complex CFTs*, *JHEP* **10** (2018) 108 [[arXiv:1807.11512](#)] [[INSPIRE](#)].
- [50] N. Arkani-Hamed, Y.-T. Huang and S.-H. Shao, *On the Positive Geometry of Conformal Field Theory*, [arXiv:1812.07739](#) [[INSPIRE](#)].
- [51] A. Pinkus, *Spectral properties of totally positive kernels and matrices*, in *Total positivity and its applications*, Springer, Heidelberg Germany (1996), pg. 477.
- [52] A. Belin, C.A. Keller and A. Maloney, *String Universality for Permutation Orbifolds*, *Phys. Rev. D* **91** (2015) 106005 [[arXiv:1412.7159](#)] [[INSPIRE](#)].
- [53] C. Vafa, *The String landscape and the swampland*, [hep-th/0509212](#) [[INSPIRE](#)].