

A semidefinite program solver for the conformal bootstrap

David Simmons-Duffin

*School of Natural Sciences, Institute for Advanced Study,
Princeton, New Jersey 08540, U.S.A.*

E-mail: davidsd@gmail.com

ABSTRACT: We introduce SDPB: an open-source, parallelized, arbitrary-precision semidefinite program solver, designed for the conformal bootstrap. SDPB significantly outperforms less specialized solvers and should enable many new computations. As an example application, we compute a new rigorous high-precision bound on operator dimensions in the 3d Ising CFT, $\Delta_\sigma = 0.518151(6)$, $\Delta_\epsilon = 1.41264(6)$.

KEYWORDS: Conformal and W Symmetry, Field Theories in Lower Dimensions

ARXIV EPRINT: [1502.02033](https://arxiv.org/abs/1502.02033)

Contents

1	Introduction	1
2	Design of SDPB	3
2.1	Polynomial matrix programs	3
2.2	Translating PMPs into SDPs	3
2.3	Semidefinite program duality	6
2.4	An interior point method	7
2.4.1	Newton search direction	8
2.4.2	Mehrotra predictor-corrector trick	9
2.4.3	Termination conditions	9
2.4.4	Complete algorithm	10
2.5	Specialization to polynomial matrix programs	11
2.5.1	Block structure of the Schur complement matrix	11
2.5.2	Computing the Schur complement matrix	13
2.5.3	Computing step lengths	14
2.6	Implementation	14
3	Example application: 3d Ising critical exponents	15
3.1	A 3d Ising optimization problem	15
3.2	Approximation as a PMP	16
3.3	Setting up SDPB	17
3.4	Results	18
4	Discussion	22
A	Choices and parameters	23
B	Performance	24
B.1	Complexity comparison	24
B.2	Running time for 3d Ising computations	25

1 Introduction

In [1], Rattazzi, Rychkov, Tonni, and Vichi showed how to bound Conformal Field Theory (CFT) observables using convex optimization. Shockingly, the resulting bounds are sometimes saturated by actual CFTs, allowing high-precision computations of nonperturbative quantities [2–9]. This modern incarnation of the conformal bootstrap [10, 11] has been applied to numerous theories, and the list is growing [12–33].

With recent analytical work [34–40], we now understand in principle how to formulate many interesting bootstrap calculations, like studies of four-point functions of scalars, fermions, conserved currents, stress-tensors, and mixed correlators combining all of these ingredients. Such studies may shed light on the classification of critical points in condensed matter systems, the conformal windows of 4d gauge theories, the landscape of AdS string vacua, and more. Each study inevitably culminates in an optimization problem that only can be solved numerically at present.

SDPB is a custom optimizer with three purposes:

1. To enable the next generation of bootstrap studies involving conserved currents, stress tensors, and other complex ingredients.
2. To help improve predictions for CFTs already isolated with the bootstrap, like the 3d Ising CFT, 3d $O(N)$ vector models, and others.
3. To demonstrate the potential for dramatic improvement in current numerical bootstrap techniques and to encourage researchers in numerical optimization and computer science to contribute ideas and expertise.

Custom optimizers have improved bootstrap calculations in the past. For example in [5], a custom-written linear program solver enabled a new high-precision calculation of critical exponents in the 3d Ising CFT, surpassing what was possible with out-of-the-box solvers like `Mathematica` [41] (used in the original work [1]), `GLPK` [42], and `CPLEX` [43].¹

Unfortunately, linear programming is not applicable to systems of multiple correlators or operators with spin. These more complicated cases can be attacked with semidefinite programming [9, 17], for which previous studies have relied on the solvers `SDPA` [44, 45] and `SDPA-GMP` [46]. The study [9], in particular, pushed `SDPA-GMP` to its limits, with each optimization taking up to 2 weeks. By contrast, `SDPB` can perform the same optimization in 1–3 CPU-hours, or 4–12 minutes on a 16-core machine.

A general bootstrap problem can be approximated as a particular type of semidefinite program called a “polynomial matrix program” (PMP). `SDPB` implements a variant of the well-known primal-dual interior point method for semidefinite programming [47–50], specialized for PMPs. Specialization and parallelization are `SDPB`’s advantages. We are optimistic that better designs and algorithms can be brought to bear in the future.

In section 2, we describe PMPs and the design of `SDPB`. This discussion is mostly independent of the conformal bootstrap and should be comprehensible to readers without a physics background.

In section 3, as an application of `SDPB`, we set a new world-record for precision of critical exponents in the 3d Ising CFT, using multiple correlators as in [9]. Readers interested solely in physics can skip to this section. We conclude in section 4.

`SDPB` is open source and available online at <https://github.com/davidsd/sdpb>.

¹The solver in [5] is more accurately called a semi-infinite program solver.

2 Design of SDPB

2.1 Polynomial matrix programs

SDPB solves the following type of problem, which we call a *polynomial matrix program* (PMP). Consider a collection of symmetric polynomial matrices

$$M_j^n(x) = \begin{pmatrix} P_{j,11}^n(x) & \dots & P_{j,1m_j}^n(x) \\ \vdots & \ddots & \vdots \\ P_{j,m_j1}^n(x) & \dots & P_{j,m_jm_j}^n(x) \end{pmatrix} \quad (2.1)$$

labeled by $0 \leq n \leq N$ and $1 \leq j \leq J$, where each element $P_{j,rs}^n(x)$ is a polynomial. Given $b \in \mathbb{R}^N$, we would like to

$$\begin{aligned} &\text{maximize} && b \cdot y \text{ over } y \in \mathbb{R}^N, \\ &\text{such that} && M_j^0(x) + \sum_{n=1}^N y_n M_j^n(x) \succeq 0 \text{ for all } x \geq 0 \text{ and } 1 \leq j \leq J. \end{aligned} \quad (2.2)$$

The notation $M \succeq 0$ means “ M is positive semidefinite”.

As we review in section 3, a wide class of optimization problems from the confor-
mal bootstrap can be written in this form. Conveniently, PMPs can be translated into
semidefinite programs (SDPs) and solved using interior point methods. In the next few
subsections, we perform this translation and describe an interior point algorithm for solving
general SDPs. Subsequently, we make this algorithm more efficient by exploiting special
structure in PMPs.

2.2 Translating PMPs into SDPs

Let us begin by translating the PMP (2.2) into a more standard semidefinite program of
the following form:

$$\begin{aligned} &\text{maximize} && \text{Tr}(CY) + b \cdot y \text{ over } y \in \mathbb{R}^N, Y \in \mathcal{S}^K, \\ &\text{such that} && \text{Tr}(A_*Y) + By = c, \\ &&& \text{and} && Y \succeq 0, \end{aligned} \quad (2.3)$$

where

$$\begin{aligned} &c \in \mathbb{R}^P, \\ &B \in \mathbb{R}^{P \times N}, \\ &A_1, \dots, A_P, C \in \mathcal{S}^K. \end{aligned} \quad (2.4)$$

Here, \mathcal{S}^K is the space of $K \times K$ symmetric real matrices, and $\text{Tr}(A_*Y)$ denotes the vector
($\text{Tr}(A_1Y), \dots, \text{Tr}(A_PY)$) $\in \mathbb{R}^P$. The SDP (2.3) is similar to those treated by solvers like
SDPA-GMP, except that it includes the variables $y \in \mathbb{R}^N$, called “free variables” because
they are not constrained to be positive. The matrix C will eventually be set to zero, but
we include it for generality.

The first step is to relate positive semidefiniteness of polynomial matrices to positive semidefiniteness of a single matrix Y . Let $q_0(x), q_1(x), \dots$, be a collection of polynomials with degrees $0, 1, \dots$ (for example, monomials $q_m(x) = x^m$), and define the vector $\vec{q}_\delta(x) = (q_0(x), \dots, q_\delta(x))$. We call $q_m(x)$ a “bilinear basis” because products $q_m(x)q_n(x)$ span the space of polynomials. In particular, any polynomial $P(x)$ of degree d can be written

$$P(x) = \text{Tr}_{\mathbb{R}^{\delta_1+1}}(Y_1 Q_{\delta_1}(x)) + x \text{Tr}_{\mathbb{R}^{\delta_2+1}}(Y_2 Q_{\delta_2}(x)), \tag{2.5}$$

where

$$\begin{aligned} Q_\delta(x) &\equiv \vec{q}_\delta(x)\vec{q}_\delta(x)^T, \\ \delta_1 &\equiv \lfloor d/2 \rfloor, \\ \delta_2 &\equiv \lfloor (d-1)/2 \rfloor, \end{aligned} \tag{2.6}$$

and Y_1, Y_2 are (underdetermined) symmetric matrices

$$\begin{aligned} Y_1 &\in \mathcal{S}^{\delta_1+1}, \\ Y_2 &\in \mathcal{S}^{\delta_2+1}. \end{aligned} \tag{2.7}$$

For a symmetric $m \times m$ polynomial matrix $M(x)$ of degree d , we can apply this construction to each element,

$$\begin{aligned} M(x) &= \text{Tr}_{\mathbb{R}^{\delta_1+1}}(Y_1(Q_{\delta_1}(x) \otimes \mathbf{1}_{m \times m})) + x \text{Tr}_{\mathbb{R}^{\delta_2+1}}(Y_2(Q_{\delta_2}(x) \otimes \mathbf{1}_{m \times m})), \\ Y_1 &\in \mathcal{S}^{m(\delta_1+1)}, \\ Y_2 &\in \mathcal{S}^{m(\delta_2+1)}, \end{aligned} \tag{2.8}$$

where now Y_1 acts on $\mathbb{R}^{\delta_1+1} \otimes \mathbb{R}^m$ and Y_2 similarly acts on $\mathbb{R}^{\delta_2+1} \otimes \mathbb{R}^m$, and each trace is over the first tensor factor.

The translation of PMPs into SDPs relies on the following theorem:

Theorem 2.1. *$M(x)$ is positive semidefinite for $x \in \mathbb{R}^+$ if and only if it can be written in the form (2.8) for some positive semidefinite Y_1 and Y_2 .*

Proof. One direction is simple: choose a vector $v \in \mathbb{R}^m$ and consider the pairing $v^T M(x)v = \text{Tr}(Y_1(Q_{\delta_1}(x) \otimes vv^T)) + \text{Tr}(Y_2(xQ_{\delta_2}(x) \otimes vv^T))$. Suppose Y_1, Y_2 are positive semidefinite and $x \geq 0$. Then since $Q_{\delta_1}(x) \otimes vv^T$ and $xQ_{\delta_2}(x) \otimes vv^T$ are both positive semidefinite, it follows that $v^T M(x)v \geq 0$. The other direction is less trivial. It has been proven directly in [51] and also follows as a consequence of the “Biform Theorem” of [52], using the substitution $x = y^2$ and results of [53].² □

Theorem 2.1 lets us rewrite our PMP constraints (2.2) in terms of a collection of positive semidefinite matrices Y_1, \dots, Y_{2J} . We equate each polynomial matrix $M_j^n(x)$ to

²We thank Pablo Parrilo for pointing this out.

an expression of the form (2.8). Individual matrix elements $P_{j,rs}^n$ can be isolated by taking the trace over \mathbb{R}^{m_j} with symmetrized unit matrices

$$(E^{rs})_{ij} \equiv \frac{1}{2}(\delta_i^r \delta_j^s + \delta_i^s \delta_j^r). \quad (2.9)$$

This gives a set of polynomial equalities

$$\begin{aligned} P_{j,rs}^0(x) + \sum_n y_n P_{j,rs}^n(x) &= \text{Tr}\left(Y_{2j-1}(Q_{\delta_{j1}}(x) \otimes E^{rs})\right) + \text{Tr}\left(Y_{2j}(xQ_{\delta_{j2}}(x) \otimes E^{rs})\right), \\ d_j &\equiv \max_{n=0}^N \deg(M_j^n(x)), \\ \delta_{j1} &\equiv \lfloor d_j/2 \rfloor, \\ \delta_{j2} &\equiv \lfloor (d_j - 1)/2 \rfloor. \end{aligned} \quad (2.10)$$

Equality between polynomials of degree d_j is equivalent to equality at $d_j + 1$ points. Thus, evaluating (2.10) at points x_0, \dots, x_{d_j} , we obtain a set of affine relations between the y_n and Y_j ,³

$$\begin{aligned} P_{j,rs}^0(x_k) + \sum_n y_n P_{j,rs}^n(x_k) &= \text{Tr}\left(Y_{2j-1}(Q_{\delta_{j1}}(x_k) \otimes E^{rs})\right) + \text{Tr}\left(Y_{2j}(x_k Q_{\delta_{j2}}(x_k) \otimes E^{rs})\right) \\ 0 \leq r \leq s < m_j, \quad k &= 0, \dots, d_j. \end{aligned} \quad (2.11)$$

Let us group the Y_j 's into a single block-diagonal positive semidefinite matrix

$$Y = \begin{pmatrix} Y_1 & 0 & \cdots & 0 \\ 0 & Y_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & Y_{2J} \end{pmatrix}. \quad (2.12)$$

The constraints (2.11) now take the form

$$\text{Tr}(A_p Y) + (By)_p = c_p, \quad (2.13)$$

where p runs over all tuples (j, r, s, k) satisfying $0 \leq r \leq s < m_j$, and $0 \leq k \leq d_j$. The matrices A_p , B , C and vector c are given by

$$A_{(j,r,s,k)} = \begin{pmatrix} 0 & \cdots & 0 & & 0 & & \cdots & 0 \\ \vdots & \ddots & \vdots & & \vdots & & & \vdots \\ 0 & \cdots & Q_{\delta_{j1}}(x_k) \otimes E^{rs} & & 0 & & \cdots & 0 \\ 0 & \cdots & 0 & & x_k Q_{\delta_{j2}}(x_k) \otimes E^{rs} & & \cdots & 0 \\ \vdots & & \vdots & & \vdots & & \ddots & \vdots \\ 0 & \cdots & 0 & & 0 & & \cdots & 0 \end{pmatrix}, \quad (2.14)$$

$$B_{(j,r,s,k),n} = -P_{j,rs}^n(x_k), \quad (2.15)$$

$$c_{(j,r,s,k)} = P_{j,rs}^0(x_k), \quad (2.16)$$

$$C = 0. \quad (2.17)$$

³We could alternatively match coefficients on both sides of (2.10) as in [7, 9, 17]. We will see in subsection 2.5.2 why pointwise evaluation is preferable.

This completes the translation of our PMP into an SDP (2.3). Note that the matrices A_p are far from generic. Exploiting this fact will help us solve PMPs much more efficiently than a generic SDP. SDPB is specifically designed to solve SDPs with constraint matrices of the form

$$A_{(j,r,s,k)} = \sum_{b \in \text{blocks}_j} \mathcal{B}_b(\vec{v}_{b,k} \vec{v}_{b,k}^T \otimes E^{rs}), \tag{2.18}$$

where $\mathcal{B}_b(M)$ denotes the block-diagonal matrix with M in the b -th block and zeros everywhere else,

$$\mathcal{B}_b(M) \equiv \left(\begin{array}{cccc} 0 & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & & \vdots \\ 0 & \cdots & M & \cdots & 0 \\ \vdots & & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \cdots & 0 \end{array} \right) \Bigg\}^b, \tag{2.19}$$

and the sets blocks_j are disjoint for different j . In the example above, we have

$$\begin{aligned} \text{blocks}_j &= \{2j - 1, 2j\}, \\ \vec{v}_{2j-1,k} &= \vec{q}_{\delta_{j1}}(x_k), \\ \vec{v}_{2j,k} &= x_k^{1/2} \vec{q}_{\delta_{j2}}(x_k). \end{aligned} \tag{2.20}$$

2.3 Semidefinite program duality

Duality plays an important role in our interior point algorithm, so let us briefly review it. The problem (2.3) is related by duality to the following ‘‘primal’’ optimization problem:

$$\begin{aligned} \mathcal{P} : \quad & \text{minimize} \quad c \cdot x \text{ over } x \in \mathbb{R}^P, X \in \mathcal{S}^K, \\ & \text{such that} \quad X = \sum_{p=1}^P A_p x_p - C, \\ & B^T x = b, \\ & X \succeq 0. \end{aligned} \tag{2.21}$$

We refer to the problem (2.21) as \mathcal{P} for ‘‘primal’’ and (2.3) as \mathcal{D} for ‘‘dual’’. We say that \mathcal{P} is ‘‘feasible’’ if there exist x, X satisfying the constraints (2.21). Similarly, \mathcal{D} is feasible if there exist y, Y satisfying the constraints (2.3). The ‘‘duality gap’’ is defined as the difference in primal and dual objective functions, $c \cdot x - \text{Tr}(CY) - b \cdot y$.

For our purposes, the statement of duality is as follows:

Theorem 2.2 (Semidefinite Program Duality). *Given a feasible point (x, X) of \mathcal{P} and a feasible point (y, Y) of \mathcal{D} , the duality gap is nonnegative. If the duality gap vanishes, then (x, X) and (y, Y) are each optimal solutions of \mathcal{P} and \mathcal{D} , and furthermore $XY = 0$.*

Proof. Suppose we have feasible solutions (x, X) and (y, Y) . The duality gap is given by

$$\begin{aligned}
 c \cdot x - \text{Tr}(CY) - b \cdot y &= c \cdot x - \text{Tr}\left(\left(\sum_p A_p x_p - X\right)Y\right) - b \cdot y \\
 &= c \cdot x + \text{Tr}(XY) - x \cdot \text{Tr}(A_* Y) - b \cdot y \\
 &= c \cdot x + \text{Tr}(XY) - x \cdot (\text{Tr}(A_* Y) + B y) \\
 &= \text{Tr}(XY) \geq 0,
 \end{aligned}
 \tag{2.22}$$

where nonnegativity follows because X and Y are positive semidefinite. Now suppose $\text{Tr}(XY)$ vanishes. Clearly this implies $XY = 0$ identically (this condition is called “complementarity”). Because $c \cdot x$ is bounded from below by the dual objective $\text{Tr}(CY) + b \cdot y$ and also equal to the dual objective, the point (x, X) must be optimal. Similarly, since $\text{Tr}(CY) + b \cdot y$ is bounded from above by the primal objective $c \cdot x$ and also equal to the primal objective, the point (y, Y) must be optimal as well. \square

Unlike in linear programming, there is no guarantee that either \mathcal{P} or \mathcal{D} will attain their respective optima, or that the duality gap will vanish. For this, we need additional regularity assumptions. One of them is Slater’s condition, which says that the duality gap vanishes if there exist *strictly* feasible solutions to the primal and dual constraints — i.e. solutions where $X, Y \succ 0$ are positive-definite. Slater’s condition is generic in the sense that a small perturbation of a feasible but not strictly-feasible problem will typically produce a strictly-feasible problem.

2.4 An interior point method

The idea behind primal-dual interior point methods is to solve both the primal and dual equations simultaneously to find an optimal point $q = (x, X, y, Y)$. As we saw in the previous subsection, the optimum is (generically) achieved by a pair of positive semidefinite matrices X, Y satisfying the “complementarity” condition $XY = 0$. Most algorithms work by deforming this condition to

$$XY = \mu I \tag{2.23}$$

for some nonzero μ , where I is the identity matrix. The constraints together with (2.23) then have a unique family of solutions called the “central path”: $q(\mu) = (x(\mu), X(\mu), y(\mu), Y(\mu))$ indexed by $\mu \in \mathbb{R}^+$. By following the central path from positive μ towards $\mu = 0$, we can find the optimum of the original problem.

In practice, instead of moving precisely along the central path, we use the following strategy. Consider an initial point $q = (x, X, y, Y)$ such that X, Y are positive semidefinite. Our goal is to decrease $\text{Tr}(XY)$ and move towards the constraint locus while maintaining positive semidefiniteness.

- Set $\mu = \beta \text{Tr}(XY)/K$ for some $\beta < 1$, where K is the number of rows of X .
- Use Newton’s method to compute a direction $dq = (dx, dX, dy, dY)$ towards the central path with the given μ .

- Take a step along dq , taking care not to violate the positive semidefiniteness of X, Y . This should result in a reduction of $\text{Tr}(XY)$ by roughly a factor of β .
- Repeat.

This is essentially Newton's method with a moving target.

An important advantage of this method is that the initial starting point (x, X, y, Y) need not satisfy any of the equality constraints in (2.3) and (2.21). As long as we start with positive semidefinite X, Y , and the problem is feasible, the above method will converge to a point where the equality constraints are satisfied.

2.4.1 Newton search direction

Let us describe a single Newton step in more detail. The direction dq is defined by replacing $q \rightarrow q + dq$ and solving the constraint equations (2.3), (2.21) and complementarity equation (2.23) at linear order in dq ,

$$\begin{aligned}
 X + dX &= \sum_i A_i(x_i + dx_i) - C, \\
 B^T(x + dx) &= b, \\
 \text{Tr}(A_*(Y + dY)) + B(y + dy) &= c, \\
 XY + XdY + dXY &= \mu I.
 \end{aligned} \tag{2.24}$$

The residues

$$\begin{aligned}
 P &\equiv \sum_i A_i x_i - X - C, \\
 p &\equiv b - B^T x, \\
 d &\equiv c - \text{Tr}(A_* Y) - B y, \\
 R &\equiv \mu I - XY,
 \end{aligned} \tag{2.25}$$

measure the failure of the current point to satisfy the constraints. These residues will decrease with each Newton step. The linearized equations (2.24) can then be written

$$\begin{pmatrix} S & -B \\ B^T & 0 \end{pmatrix} \begin{pmatrix} dx \\ dy \end{pmatrix} = \begin{pmatrix} -d - \text{Tr}(A_* Z) \\ p \end{pmatrix}, \tag{2.26}$$

$$dX = P + \sum_i A_i dx_i, \tag{2.27}$$

$$dY = X^{-1}(R - dXY), \tag{2.28}$$

where $Z = X^{-1}(PY - R)$ and the ‘‘Schur complement’’ matrix S is defined by

$$S_{ij} = \text{Tr}(A_i X^{-1} A_j Y). \tag{2.29}$$

To find the search direction, we first solve (2.26) for dx, dy , and then plug into (2.27) and (2.28) to determine dX, dY . Naively applying (2.28) leads to a dY that is not necessarily

symmetric, which would take us outside of the domain of definition of Y . Several solutions to this problem have been proposed [54]. Our approach, following [49, 55, 56], will be to symmetrize dY by hand, replacing (2.28) with

$$\begin{aligned}\widehat{dY} &= X^{-1}(R - dXY), \\ dY &= \frac{1}{2}(\widehat{dY} + \widehat{dY}^T).\end{aligned}\tag{2.30}$$

2.4.2 Mehrotra predictor-corrector trick

The most expensive operations in the search direction calculation are forming the Schur complement matrix S and solving the Schur complement equation (2.26). We'd like to perform them as rarely as possible. A simple modification to the naïve Newton's method, due to Mehrotra [57], allows us to get closer to the central path while reusing S and most of the work done in solving (2.26).

The rough idea is to solve the constraint and complementarity equations at higher order. This proceeds in two steps, called the “predictor” and “corrector” steps, respectively. For the predictor step, we compute a direction as described above, which we call $dq_p = (dx_p, dX_p, dy_p, dY_p)$. We then replace the linearized complementarity equation (2.24) with

$$XY + XdY + dXY + dX_p dY_p = \mu I\tag{2.31}$$

and re-solve to obtain a corrector direction dq_c . In the corrector step, we may (optionally) use a smaller deformation parameter $\mu \rightarrow \mu_c$ to get closer to $\text{Tr}(XY) = 0$. Note that the replacement of (2.24) with (2.31) does not change the Schur complement matrix S_{ij} , so it can be reused, together with any matrix decompositions performed in solving (2.26). Altogether, the corrector step amounts to simply replacing

$$R \rightarrow \mu_c I - XY - dX_p dY_p\tag{2.32}$$

before solving (2.26), (2.27), and (2.30).

2.4.3 Termination conditions

We say a point q is “primal feasible” if the residues p, P are sufficiently small. Similarly, the solution is “dual feasible” if the residue d is sufficiently small. The precise conditions are

$$\begin{aligned}\text{primal feasible: } \text{primalError} &\equiv \max_{i,j} \{|p_i|, |P_{ij}|\} < \text{primalErrorThreshold}; \\ \text{dual feasible: } \text{dualError} &\equiv \max_i \{|d_i|\} < \text{dualErrorThreshold},\end{aligned}\tag{2.33}$$

where $\text{primalErrorThreshold} \ll 1$ and $\text{dualErrorThreshold} \ll 1$ are parameters chosen by the user.

An optimal point should be both primal and dual feasible, and have (nearly) equal primal and dual objective values. Specifically, let us define dualityGap as the normalized difference between the primal and dual objective functions

$$\begin{aligned}\text{dualityGap} &\equiv \frac{|\text{primalObjective} - \text{dualObjective}|}{\max\{1, |\text{primalObjective} + \text{dualObjective}|\}}, \\ \text{primalObjective} &\equiv c \cdot x, \\ \text{dualObjective} &\equiv \text{Tr}(CY) + b \cdot y.\end{aligned}\tag{2.34}$$

A point is considered “optimal” if

$$\text{dualityGap} < \text{dualityGapThreshold}, \quad (2.35)$$

where $\text{dualityGapThreshold} \ll 1$ is chosen by the user.

2.4.4 Complete algorithm

Our complete interior point algorithm is as follows.

1. Choose an initial point $q = (x, X, y, Y) = (0, \Omega_{\mathcal{P}}I, 0, \Omega_{\mathcal{D}}I)$ where $\Omega_{\mathcal{P}}$ and $\Omega_{\mathcal{D}}$ are real and positive. This point probably does not satisfy the constraints.
2. Compute the residues (2.25) and terminate if q is simultaneously primal feasible, dual feasible, and optimal. (Sometimes we may wish to use a different termination criterion, see below.)
3. Let $\mu = \text{Tr}(XY)/K$ and compute the predictor deformation parameter $\mu_p = \beta_p\mu$ where

$$\beta_p = \begin{cases} 0 & \text{if } q \text{ is primal and dual feasible;} \\ \beta_{\text{infeasible}} & \text{otherwise.} \end{cases} \quad (2.36)$$

Here, $\beta_{\text{infeasible}} \in (0, 1)$ is chosen by the user.

4. Compute the predictor search direction dq_p by solving eqs. (2.26), (2.27), (2.30) with $R = \mu_p I - XY$.
5. Compute the corrector deformation parameter $\mu_c = \beta_c\mu$ as follows. Let $r = \text{Tr}((X + dX_p)(Y + dY_p))/(\mu K)$ and set

$$\beta = \begin{cases} r^2 & \text{if } r < 1; \\ r & \text{otherwise,} \end{cases} \quad (2.37)$$

$$\beta_c = \begin{cases} \min(\max(\beta_{\text{feasible}}, \beta), 1) & \text{if } q \text{ is primal and dual feasible;} \\ \max(\beta_{\text{infeasible}}, \beta) & \text{otherwise,} \end{cases} \quad (2.38)$$

where $\beta_{\text{feasible}} \in (0, 1)$ is a parameter chosen by the user. This choice of β_c is modeled after the one in SDPA.

6. Compute the corrector search direction dq_c by solving eqs. (2.26), (2.27), (2.30) with $R = \mu_c I - XY - dX_p dY_p$.
7. Determine the primal and dual step lengths

$$\begin{aligned} \alpha_{\mathcal{P}} &= \min(\gamma\alpha(X, dX_c), 1), \\ \alpha_{\mathcal{D}} &= \min(\gamma\alpha(Y, dY_c), 1), \end{aligned} \quad (2.39)$$

where $\alpha(M, dM)$ is the largest positive real number such that $M + \alpha(M, dM)dM$ is positive semidefinite, and $\gamma \in (0, 1)$ is a parameter chosen by the user.

8. Replace

$$\begin{aligned}
 x &\rightarrow x + \alpha_{\mathcal{P}} dx_c, \\
 X &\rightarrow X + \alpha_{\mathcal{P}} dX_c, \\
 y &\rightarrow y + \alpha_{\mathcal{D}} dy_c, \\
 Y &\rightarrow Y + \alpha_{\mathcal{D}} dY_c,
 \end{aligned}
 \tag{2.40}$$

and go to step 2. Note that the replacement (2.40) is guaranteed to preserve positive semidefiniteness of X and Y .

If the current point is close enough to a primal (or dual) feasible region, the step-length $\alpha_{\mathcal{P}}$ ($\alpha_{\mathcal{D}}$) in (2.39) can be exactly 1. When this occurs, the replacement (2.40) will exactly solve the primal (dual) equality constraints, up to numerical errors. This follows from linearity of the equality constraints, together with the fact that symmetrizing Y in (2.30) has no effect on the constraints. In cases where we only care about primal or dual feasibility, the iteration can be stopped here, see section 3.4.

2.5 Specialization to polynomial matrix programs

As mentioned in section 2.4.2, the most expensive part of the search direction calculation is computing the Schur complement matrix $S_{pq} = \text{Tr}(A_p X^{-1} A_q Y)$ and inverting (2.26) to obtain dx, dy . In this section, we will specialize to PMPs, and study how these calculations can be made more efficient. For similar optimizations, see [58].

2.5.1 Block structure of the Schur complement matrix

Recall that for PMPs, the matrices A_p are given by (2.18) with the index p running over tuples (j, r, s, k) satisfying $0 \leq r \leq s < m_j$, and $0 \leq k \leq d_j$. Since X and Y have the block structure (2.12), the Schur complement matrix $S_{p_1 p_2}$ is block-diagonal: it has nonzero entries only if $j_1 = j_2$,

$$S = \begin{pmatrix} S^{(1)} & 0 & \dots & 0 \\ 0 & S^{(2)} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & S^{(J)} \end{pmatrix}.
 \tag{2.41}$$

The size of the j -th block $\dim S^{(j)}$ is equal to the number of choices for (r, s, k) ,

$$\dim S^{(j)} = \frac{m_j(m_j + 1)}{2} (d_j + 1).
 \tag{2.42}$$

Now consider equation (2.26),

$$T \begin{pmatrix} dx \\ dy \end{pmatrix} = \begin{pmatrix} -d - \text{Tr}(A_* Z) \\ p \end{pmatrix}, \quad \text{where} \quad T \equiv \begin{pmatrix} S & -B \\ B^T & 0 \end{pmatrix}.
 \tag{2.43}$$

We could solve it using an LU (lower triangular \times upper triangular) decomposition of T , but this is extremely expensive, taking cubic time in $\dim T = \sum_j \dim S^{(j)} + N$.

We should use the block structure of T to our advantage. Let $S = LL^T$ be a Cholesky decomposition of S (which can be computed efficiently for each block $S^{(j)} = L^{(j)}L^{(j)T}$), and consider the decomposition

$$\begin{pmatrix} S & -B \\ B^T & 0 \end{pmatrix} = \begin{pmatrix} L & 0 \\ B^T L^{-T} & \mathbf{1} \end{pmatrix} \begin{pmatrix} \mathbf{1} & 0 \\ 0 & B^T L^{-T} L^{-1} B \end{pmatrix} \begin{pmatrix} L^T & -L^{-1} B \\ 0 & \mathbf{1} \end{pmatrix}. \quad (2.44)$$

The outer matrices on the right hand side are triangular, and can be solved efficiently by forward/backward-substitution. Meanwhile, the matrix $Q \equiv B^T L^{-T} L^{-1} B$ typically has a much smaller dimension than S , $\dim Q = N \ll \dim T$, so the middle block-matrix can be easily solved using a Cholesky decomposition.⁴

Unfortunately, the decomposition (2.44) is numerically unstable when S is ill-conditioned. Indeed, suppose S has a very small eigenvalue (so L does too), while the full matrix T does not. Then quantities like $L^{-1}B$ that appear in (2.44) will have large entries which must nearly cancel when recombined into a solution of (2.43). Near-cancellation reduces numerical precision.

These problems stem from the off-diagonal blocks of T , which ultimately come from the free variables y in our semidefinite program. Several authors have considered the problem of efficiently and stably solving semidefinite programs with free variables, with no obvious consensus [59]. For example, [60] suggests eliminating free variables by explicitly solving the primal constraint $B^T x = b$ and taking appropriate linear combinations of the matrices A_p . However this procedure destroys the sparsity structure of S , making it no longer block diagonal and forcing us to use an expensive full Cholesky decomposition.

Preserving the structure of S and T is paramount. The simplest way to stabilize (2.44) is to increase the precision of the underlying arithmetic. In practice, this appears to be necessary anyway for larger-scale bootstrap problems, see appendix A. For additional stabilization, we use an old trick of adding low-rank pieces to S to make it better-conditioned. Suppose u_i are vectors, each with a single nonzero entry, such that $S' \equiv S + \sum_i u_i u_i^T = S + U U^T$ has no small eigenvalues.⁵ Note that S' differs from S in only a few diagonal entries — in particular it has the same block structure. Now let us replace (2.43) with the system

$$T' \begin{pmatrix} dx \\ dy \\ z \end{pmatrix} = \begin{pmatrix} S' & -B & -U \\ B^T & 0 & 0 \\ U^T & 0 & -\mathbf{1} \end{pmatrix} \begin{pmatrix} dx \\ dy \\ z \end{pmatrix} = \begin{pmatrix} -d - \text{Tr}(A_* Z) \\ p \\ 0 \end{pmatrix}. \quad (2.45)$$

By solving for z and substituting back in, it is easy to see that (2.45) is precisely equivalent to (2.43). However, the advantage is that because S' is well-conditioned, a decomposition

⁴The matrix Q , is often called the ‘‘Schur complement’’ in the block matrix decomposition (2.44). We will continue to use the words ‘‘Schur complement’’ to refer to S (which is the Schur complement of a different block matrix system). Hopefully this will not cause confusion.

⁵The u_i can be found as follows. During Cholesky decomposition $S = LL^T$, we keep track of the geometric mean Λ of the diagonal entries. Whenever we encounter a diagonal entry $L_{ii} < \theta \Lambda$, where $\theta \ll 1$ is a parameter, we replace $L_{ii} \rightarrow L_{ii} + \Lambda$, which amounts to choosing $u_i = \Lambda e_i$ where e_i is a unit vector in the i -th direction.

like (2.44) is numerically stable. Indeed, defining $B' = (B \ U)$, we have

$$T' = \begin{pmatrix} L' & 0 \\ B'^T L'^{-T} & \mathbf{1} \end{pmatrix} \begin{pmatrix} \mathbf{1} & 0 \\ 0 & Q' \end{pmatrix} \begin{pmatrix} L'^T & -L'^{-1}B' \\ 0 & \mathbf{1} \end{pmatrix}, \quad (2.46)$$

$$Q' = B'^T L'^{-T} L'^{-1} B' - \begin{pmatrix} 0 & 0 \\ 0 & \mathbf{1} \end{pmatrix}, \quad (2.47)$$

where $S' = L' L'^T$. Because Q' is no longer necessarily positive semidefinite, we are forced to use an LU decomposition to invert the middle matrix in (2.46), which is slightly more expensive than Cholesky decomposition. Fortunately, Q' is usually small enough that this cost is inconsequential. If T itself is ill-conditioned, then this will manifest as instabilities when we try to LU decompose Q' . In this situation, there is little we can do to avoid imprecision.

2.5.2 Computing the Schur complement matrix

Now that we know what to do with the Schur complement matrix S , let us compute it efficiently. The fact that A_p has low rank is helpful. This explains why we chose to evaluate the polynomial equalities (2.10) at discrete points x_k in (2.11). Matching polynomial coefficients on each side, as done in [7, 9, 17], leads to higher-rank matrices A_p . The helpfulness of low-rank constraints in solving SDPs, and their appearance in polynomial optimization, is well known [58].

Recall that X and Y are block diagonal (2.12), with each block X_b acting on a tensor product of the form $\mathbb{R}^{\delta+1} \otimes \mathbb{R}^m$. Let $X_b^{(r,s)} \in \mathbb{R}^{(\delta+1) \times (\delta+1)}$ denote the (r, s) -th block of X_b in the second tensor factor, which acts on $\mathbb{R}^{\delta+1}$.

Since S is block diagonal, we need only compute elements with $j_1 = j_2 = j$. We have

$$\begin{aligned} S_{(j,r_1,s_1,k_1),(j,r_2,s_2,k_2)} &= \sum_{b \in \text{blocks}_j} \text{Tr}((\vec{v}_{b,k_1} \vec{v}_{b,k_1}^T \otimes E^{r_1 s_1}) X_b^{-1} (\vec{v}_{b,k_2} \vec{v}_{b,k_2}^T \otimes E^{r_2 s_2}) Y_b) \\ &= \sum_{b \in \text{blocks}_j} \frac{1}{4} \left((\vec{v}_{b,k_1}^T (X_b^{-1})^{(s_1,r_2)} \vec{v}_{b,k_2}) (\vec{v}_{b,k_2}^T Y_b^{(s_2,r_1)} \vec{v}_{b,k_1}) \right. \\ &\quad \left. + (r_1 \leftrightarrow s_1) + (r_2 \leftrightarrow s_2) + (r_1 \leftrightarrow s_1, r_2 \leftrightarrow s_2) \right). \end{aligned} \quad (2.48)$$

Thus, instead of performing repeated matrix multiplications to calculate $\text{Tr}(A_{p_1} X^{-1} A_{p_2} Y)$, we can precompute the bilinear pairings

$$U_{(d_j+1)s+k_1,(d_j+1)r+k_2}^{(b)} \equiv \vec{v}_{b,k_1}^T (X_b^{-1})^{(s,r)} \vec{v}_{b,k_2}, \quad (2.49)$$

$$V_{(d_j+1)s+k_2,(d_j+1)r+k_1}^{(b)} \equiv \vec{v}_{b,k_2}^T Y_b^{(s,r)} \vec{v}_{b,k_1}, \quad (2.50)$$

and plug them into (2.48). Whereas forming S is often the most expensive operation in less-specialized solvers, the method outlined here makes it subdominant to other computations, see appendix B.1.

2.5.3 Computing step lengths

To find step lengths $\alpha_{\mathcal{P}}, \alpha_{\mathcal{D}}$, we must be able to compute $\alpha(M, dM)$ where M is a positive semidefinite matrix and $\alpha(M, dM)$ is the largest positive real number such that $M + \alpha(M, dM)dM$ is positive semidefinite. Let $M = LL^T$ be a Cholesky decomposition of M . Since $M + \alpha dM = L(\mathbf{1} + \alpha L^{-1}dML^{-T})L^T$, we have

$$\alpha(M, dM) = -1/\text{min-eigenvalue}(L^{-1}dML^{-T}). \tag{2.51}$$

In SDPB, we compute all the eigenvalues of $L^{-1}dML^{-T}$ using a QR decomposition and then simply take the minimum. Some solvers, like SDPA, implement the more efficient Lanczos method [61] for computing the minimum eigenvalue. In practice, the step-length calculation is a small part of the total running time, so we have neglected this optimization.

2.6 Implementation

SDPB is approximately 3500 lines of C++. It uses the GNU Multiprecision Library (GMP) [62] for arbitrary precision arithmetic, and MPACK [63] for arbitrary precision linear algebra. The relevant MPACK source files are included with SDPB, with some slight modifications:

- The Cholesky decomposition routine `Rpotrf` has been modified to implement the stabilization procedure described in footnote 5.
- Some loops in the LU decomposition routine `Rgetrf` have been parallelized.

SDPB also depends on the Boost C++ libraries [64] and the parsing library `tinyxml2` [65].

Previous experience shows that high-precision arithmetic is important for accurately solving bootstrap optimization problems. It is not fully understood why. The naïve reason is that derivatives $\partial_z^m \partial_{\bar{z}}^n g_{\Delta, \ell}(z, \bar{z})$ of conformal blocks vary by many orders of magnitude relative to each other as Δ varies. It is not possible to scale away this large variation, and answers may depend on near cancellation of large numbers. In practice, the matrix manipulations in our interior point algorithm “leak” precision, so that the search direction (dx, dX, dy, dY) is less precise than the initial point (x, X, y, Y) . By increasing the precision of the underlying arithmetic, the search direction can be made reliable again. This strategy (which we adopt) comes at a cost of increased runtime and memory usage. Better strategies for dealing with numerical instabilities in bootstrap problems could bring enormous gains.

SDPB is parallelized with OpenMP [66]. Because most matrices appearing in the interior point algorithm are block-diagonal, most computations are “embarrassingly parallel”: different blocks can be distributed to different threads. (The most prominent exception is the LU decomposition of Q' , which is why `Rgetrf` was modified.) Consequently, performance scales nearly linearly with the number of cores, as long as the number of matrix blocks is sufficiently large. This is usually the case for interesting bootstrap problems, where J (which sets the number of blocks) is typically much larger than the number of available cores. It should be possible to achieve favorable scaling up to dozens or even hundreds of cores using MPI and more careful memory management. Further scaling should be possible with more fine-grained parallelization.

SDPB is available online at <https://github.com/davidsd/sdpb> under the MIT license. The source code is carefully commented and written for readability (to the extent that C++ code is ever readable). We hope this will encourage customization and improvement.

3 Example application: 3d Ising critical exponents

3.1 A 3d Ising optimization problem

Bootstrap optimization problems can be naturally approximated as PMPs [9, 17]. In this section, we review the PMP for the system of correlators $\{\langle\sigma\sigma\sigma\sigma\rangle, \langle\sigma\sigma\epsilon\epsilon\rangle, \langle\epsilon\epsilon\epsilon\epsilon\rangle\}$ in the 3d Ising CFT. We will be brief. Much more detail is given in [9].

Associativity of the Operator Product Expansion (OPE) for $\{\langle\sigma\sigma\sigma\sigma\rangle, \langle\sigma\sigma\epsilon\epsilon\rangle, \langle\epsilon\epsilon\epsilon\epsilon\rangle\}$ implies the consistency condition

$$(1 \ 1) \vec{V}_{+,0,0} \begin{pmatrix} 1 \\ 1 \end{pmatrix} + \sum_{\mathcal{O}^+} (\lambda_{\sigma\sigma\mathcal{O}} \ \lambda_{\epsilon\epsilon\mathcal{O}}) \vec{V}_{+,\Delta,\ell} \begin{pmatrix} \lambda_{\sigma\sigma\mathcal{O}} \\ \lambda_{\epsilon\epsilon\mathcal{O}} \end{pmatrix} + \sum_{\mathcal{O}^-} \lambda_{\sigma\epsilon\mathcal{O}}^2 \vec{V}_{-,\Delta,\ell} = 0. \quad (3.1)$$

Here, \mathcal{O}^+ runs over \mathbb{Z}_2 -even operators of even spin and \mathcal{O}^- runs over \mathbb{Z}_2 -odd operators of any spin. We have separated out the unit operator. Δ and ℓ are the dimension and spin of \mathcal{O} , respectively. The object $\vec{V}_{-,\Delta,\ell}$ is a 5-vector and $\vec{V}_{+,\Delta,\ell}$ is a 5-vector of 2×2 matrices

$$\vec{V}_{-,\Delta,\ell} = \begin{pmatrix} 0 \\ 0 \\ F_{-,\Delta,\ell}^{\sigma\epsilon,\sigma\epsilon}(u,v) \\ (-1)^\ell F_{-,\Delta,\ell}^{\epsilon\sigma,\sigma\epsilon}(u,v) \\ -(-1)^\ell F_{+,\Delta,\ell}^{\epsilon\sigma,\sigma\epsilon}(u,v) \end{pmatrix}, \quad \vec{V}_{+,\Delta,\ell} = \begin{pmatrix} \begin{pmatrix} F_{-,\Delta,\ell}^{\sigma\sigma,\sigma\sigma}(u,v) & 0 \\ 0 & 0 \end{pmatrix} \\ \begin{pmatrix} 0 & 0 \\ 0 & F_{-,\Delta,\ell}^{\epsilon\epsilon,\epsilon\epsilon}(u,v) \end{pmatrix} \\ \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} \\ \begin{pmatrix} 0 & \frac{1}{2} F_{-,\Delta,\ell}^{\sigma\sigma,\epsilon\epsilon}(u,v) \\ \frac{1}{2} F_{-,\Delta,\ell}^{\sigma\sigma,\epsilon\epsilon}(u,v) & 0 \end{pmatrix} \\ \begin{pmatrix} 0 & \frac{1}{2} F_{+,\Delta,\ell}^{\sigma\sigma,\epsilon\epsilon}(u,v) \\ \frac{1}{2} F_{+,\Delta,\ell}^{\sigma\sigma,\epsilon\epsilon}(u,v) & 0 \end{pmatrix} \end{pmatrix}, \quad (3.2)$$

with entries that are functions of conformal cross-ratios u and v ,

$$F_{\pm,\Delta,\ell}^{ij,kl}(u,v) \equiv v^{\frac{\Delta_k+\Delta_j}{2}} g_{\Delta,\ell}^{\Delta_{ij},\Delta_{kl}}(u,v) \pm u^{\frac{\Delta_k+\Delta_j}{2}} g_{\Delta,\ell}^{\Delta_{ij},\Delta_{kl}}(v,u), \quad \Delta_{ij} \equiv \Delta_i - \Delta_j. \quad (3.3)$$

The $g_{\Delta,\ell}^{\Delta_{ij},\Delta_{kl}}(v,u)$ are conformal blocks, which are known special functions.

The OPE coefficients $\lambda_{\sigma\sigma\mathcal{O}}$, $\lambda_{\sigma\epsilon\mathcal{O}}$, $\lambda_{\epsilon\epsilon\mathcal{O}}$ and dimensions Δ are not known a priori. Nonetheless, we can constrain them by understanding when it is possible for the terms in (3.1) to sum to zero. To do this, consider a 5-vector of functionals $\vec{\alpha} = (\alpha^1, \dots, \alpha^5)$, where each α^i acts on the space of functions of u and v . Acting on (3.1) with $\vec{\alpha}$ gives

$$(1 \ 1) \vec{\alpha} \cdot \vec{V}_{+,0,0} \begin{pmatrix} 1 \\ 1 \end{pmatrix} + \sum_{\mathcal{O}^+} (\lambda_{\sigma\sigma\mathcal{O}} \ \lambda_{\epsilon\epsilon\mathcal{O}}) \vec{\alpha} \cdot \vec{V}_{+,\Delta,\ell} \begin{pmatrix} \lambda_{\sigma\sigma\mathcal{O}} \\ \lambda_{\epsilon\epsilon\mathcal{O}} \end{pmatrix} + \sum_{\mathcal{O}^-} \lambda_{\sigma\epsilon\mathcal{O}}^2 \vec{\alpha} \cdot \vec{V}_{-,\Delta,\ell} = 0. \quad (3.4)$$

The bootstrap logic, in the spirit of [1], is as follows. First we make an assumption about which Δ, ℓ appear in (3.4). We then search for a functional $\vec{\alpha}$ such that

$$\begin{aligned} \begin{pmatrix} 1 & 1 \end{pmatrix} \vec{\alpha} \cdot \vec{V}_{+,0,0} \begin{pmatrix} 1 \\ 1 \end{pmatrix} &> 0, \\ \vec{\alpha} \cdot \vec{V}_{+,\Delta,\ell} &\succeq 0, \quad \text{for all } \mathbb{Z}_2\text{-even operators with even spin,} \\ \vec{\alpha} \cdot \vec{V}_{-,\Delta,\ell} &\succeq 0, \quad \text{for all } \mathbb{Z}_2\text{-odd operators of any spin.} \end{aligned} \quad (3.5)$$

The OPE coefficients λ_{ijk} are real in a unitary CFT. Thus, if $\vec{\alpha}$ exists, then it is impossible to satisfy the consistency condition (3.1), and our assumption is ruled out. By making different assumptions and searching for functionals $\vec{\alpha}$, we can map out the space of allowed Δ .

3.2 Approximation as a PMP

The conditions (3.5) define a feasibility problem with an infinite number of semidefiniteness constraints (one for each Δ, ℓ). To obtain a PMP, we choose a particular type of functional

$$\alpha^i[f] = \sum_{\substack{m \geq n \\ m+n \leq \Lambda}} a_{mn}^i \partial_z^m \partial_{\bar{z}}^n f(z, \bar{z}) \Big|_{z=\bar{z}=\frac{1}{2}}, \quad (3.6)$$

where $u = z\bar{z}$, $v = (1-z)(1-\bar{z})$. Although the bootstrap logic does not depend on the types of functionals considered, only functionals of the form (3.6) lead to a PMP. Other types of functionals require different optimization methods.

Derivatives of conformal blocks have a systematic approximation in terms of positive functions times polynomials,

$$\partial_z^m \partial_{\bar{z}}^n g_{\Delta,\ell}^{\Delta_{12},\Delta_{34}}(z, \bar{z}) \Big|_{z=\bar{z}=\frac{1}{2}} \approx \chi_\ell(\Delta) p_\ell^{\Delta_{12},\Delta_{34};mn}(\Delta), \quad (3.7)$$

where $p_\ell^{\Delta_{12},\Delta_{34};mn}(\Delta)$ are polynomials and $\chi_\ell(\Delta)$ are functions that are positive for all Δ in a unitary CFT. It follows that

$$\partial_z^m \partial_{\bar{z}}^n F_{\pm,\Delta,\ell}^{ij,kl}(z, \bar{z}) \Big|_{z=\bar{z}=\frac{1}{2}} \approx \chi_\ell(\Delta) P_{\pm,\ell}^{ij,kl;mn}(\Delta), \quad (3.8)$$

where $P_{\pm,\ell}^{ijkl;mn}(\Delta)$ are linear combinations of $p_\ell^{\Delta_{ij},\Delta_{kl};mn}(\Delta)$. Using this approximation, and stripping off the positive factors $\chi_\ell(\Delta)$, (3.5) becomes a PMP:

find a_{mn}^i such that:

$$\begin{aligned} \begin{pmatrix} 1 & 1 \end{pmatrix} Z_0(0) \begin{pmatrix} 1 \\ 1 \end{pmatrix} &> 0, \\ Z_\ell(\Delta) &\succeq 0, \quad \text{for all } \mathbb{Z}_2\text{-even operators with even spin,} \\ Y_\ell(\Delta) &\geq 0, \quad \text{for all } \mathbb{Z}_2\text{-odd operators in the spectrum.} \end{aligned} \quad (3.9)$$

Here $Y_\ell(\Delta)$ are polynomials and $Z_\ell(\Delta)$ are polynomial matrices in Δ defined as

$$\begin{aligned}
 Y_\ell(\Delta) &\equiv \sum_{mn} [a_{mn}^3 P_{-, \ell}^{\sigma\epsilon, \sigma\epsilon; mn}(\Delta) + a_{mn}^4 (-1)^\ell P_{-, \ell}^{\epsilon\sigma, \sigma\epsilon; mn}(\Delta) - a_{mn}^5 (-1)^\ell P_{+, \ell}^{\epsilon\sigma, \sigma\epsilon; mn}(\Delta)], \\
 Z_\ell(\Delta) &\equiv \sum_{mn} \begin{pmatrix} a_{mn}^1 P_{-, \ell}^{\sigma\sigma, \sigma\sigma; mn}(\Delta) & \frac{1}{2}(a_{mn}^4 P_{-, \ell}^{\sigma\sigma, \epsilon\epsilon; mn}(\Delta) + a_{mn}^5 P_{+, \ell}^{\sigma\sigma, \epsilon\epsilon; mn}(\Delta)) \\ \frac{1}{2}(a_{mn}^4 P_{-, \ell}^{\sigma\sigma, \epsilon\epsilon; mn}(\Delta) + a_{mn}^5 P_{+, \ell}^{\sigma\sigma, \epsilon\epsilon; mn}(\Delta)) & a_{mn}^2 P_{-, \ell}^{\epsilon\epsilon, \epsilon\epsilon; mn}(\Delta) \end{pmatrix}.
 \end{aligned}
 \tag{3.10}$$

Typically, we assume that Δ can vary arbitrarily above some minimum value $\Delta_{\min}(\ell)$. Writing $\Delta = \Delta_{\min}(\ell) + x$, we have positive semidefiniteness for all $x \geq 0$.

There are two important differences between (3.9) and our original PMP (2.2):

1. In (2.2) we have a finite number of positive semidefiniteness conditions $j = 1, \dots, J$, whereas here we have an infinite number since ℓ can be any nonnegative integer. In practice, we include spins ℓ up to some large but finite ℓ_{\max} . As long as ℓ_{\max} is large enough, a functional obtained by solving the problem with $\ell \leq \ell_{\max}$ should also satisfy positive semidefiniteness for spins $\ell > \ell_{\max}$. The proper choice of ℓ_{\max} depends on the problem at hand, see appendix A. See [27] for a more careful analysis.
2. In (2.2) we are trying to optimize an objective function $b \cdot y$, whereas here we are only interested in feasibility. To determine feasibility, we can pick the trivial objective function $b = 0$ and run our interior point algorithm until it becomes dual feasible.

3.3 Setting up SDPB

The natural objects entering our calculation are (approximate) derivatives of conformal blocks $\chi_\ell(\Delta) p_\ell^{\Delta_{12}, \Delta_{34}; mn}(\Delta)$, as opposed to just the polynomials $p_\ell^{\Delta_{12}, \Delta_{34}; mn}(\Delta)$. Removing positive factors does not affect positive semidefiniteness, but it does affect the scaling of the resulting SDP. Restoring quantities to their “natural” size can improve numerical stability and performance.⁶ SDPB provides a few different ways to implement this rescaling.

As we saw in section 2.2, translating a PMP into an SDP requires a bilinear basis $q_m(x)$. SDPB allows a choice of bilinear basis $q_m^{(j)}(x)$ for each $j = 1, \dots, J$. For bootstrap problems, we take $q_m^{(j)}(x)$ to be orthogonal polynomials with respect to the norm

$$\langle p, q \rangle^{(j)} = \int_0^\infty dx \chi_\ell(\Delta_{\min}(\ell) + x) p(x) q(x),
 \tag{3.11}$$

where ℓ is the spin corresponding to j . The change of basis between orthogonal polynomials with respect to $\langle \cdot, \cdot \rangle^{(j)}$ and monomials x^m (used in previous bootstrap applications of SDP) is extremely ill-conditioned at high degree. So although the choice of $q_m^{(j)}$ is unimportant in principle, it can have a dramatic effect on numerical stability.⁷

SDPB also requires a set of sample points $x_k^{(j)}$ at which to evaluate polynomials, as well as scaling factors $s_k^{(j)}$ that modify the constraints (2.15), (2.16) as follows:

$$\begin{aligned}
 B_{(j,r,s,k),n} &= -s_k^{(j)} P_{j,rs}^n(x_k^{(j)}), \\
 c_{(j,r,s,k)} &= s_k^{(j)} P_{j,rs}^0(x_k^{(j)}).
 \end{aligned}
 \tag{3.12}$$

⁶A similar observation was made for the algorithm in [5].

⁷We thank Pablo Parrilo for pointing out the usefulness of orthogonal polynomials in improving the numerical stability of polynomial optimization.

Additionally, the $A_{(j,r,s,k)}$ are given by (2.18) with

$$\begin{aligned} v_{2j-1,k} &= (s_k^{(j)})^{1/2} \vec{q}_{\delta 1}^{(j)}(x_k^{(j)}), \\ v_{2j,k} &= (s_k^{(j)})^{1/2} (x_k^{(j)})^{1/2} \vec{q}_{\delta 2}^{(j)}(x_k^{(j)}). \end{aligned} \tag{3.13}$$

This s_k -dependent rescaling gives an isomorphic, but potentially more numerically stable SDP. For bootstrap problems, it is natural to pick $s_k^{(j)} = \chi_\ell(x_k^{(j)})$ where ℓ corresponds to j . The $x_k^{(j)}$ can be any sequence of distinct points. A natural choice are zeros of one of the $q_m^{(j)}$ of sufficiently high degree.

To summarize, SDPB depends on the following input:

- for each $j = 1, \dots, J$:
 - polynomial matrices $M_j^0(x), \dots, M_j^N(x)$ of maximum degree d_j ,
 - bilinear bases $q_m^{(j)}(x)$ ($m = 0, \dots, \lfloor d_j/2 \rfloor$),
 - sample points $x_k^{(j)}$ ($k = 0, \dots, d_j$),
 - sample scalings $s_k^{(j)}$ ($k = 0, \dots, d_j$),
- an objective function $b \in \mathbb{R}^N$.

SDPB reads this data in an XML format described in the manual. A *Mathematica* package that translates the above data from *Mathematica* expressions into XML is included with the source distribution. An example 2d bootstrap computation is also included. More details about SDPB's input and output formats and its various settings can be found in the manual.

3.4 Results

As an application of SDPB, let us improve upon the determinations of 3d Ising critical exponents in [5, 9]. We make an exclusion plot for the operator dimensions $(\Delta_\sigma, \Delta_\epsilon)$ as follows. Fix $(\Delta_\sigma, \Delta_\epsilon)$ and use SDPB to determine if the PMP (3.9) is dual-feasible, i.e. whether there exist (y, Y) satisfying their associated constraints. If the PMP is dual-feasible, then the given $(\Delta_\sigma, \Delta_\epsilon)$ are excluded. If the PMP is not dual-feasible, then we cannot conclude anything about $(\Delta_\sigma, \Delta_\epsilon)$. By scanning over different points, we map out the excluded region in $(\Delta_\sigma, \Delta_\epsilon)$ space.

To determine dual feasibility, we use a vanishing objective function $b = 0$ and run SDPB with the option `--findDualFeasible`. This terminates the solver if (y, Y) are found satisfying their constraints to sufficient precision (i.e. if `dualError < dualErrorThreshold`).⁸ In practice, if SDPB finds a primal feasible solution (x, X) after some number of iterations, then it will never eventually find a dual feasible one.⁹ Thus, we additionally include the option `--findPrimalFeasible`, which terminates the solver whenever

⁸If we kept running the solver, `primalObjective` would converge towards `dualObjective = 0` and an optimum would eventually be reached.

⁹This is an experimental observation based on the computations performed in this work: every time SDPB was allowed to continue after finding a primal feasible solution, it always terminated with `maxComplementarity exceeded`. It would be interesting to understand why this is the case. One can

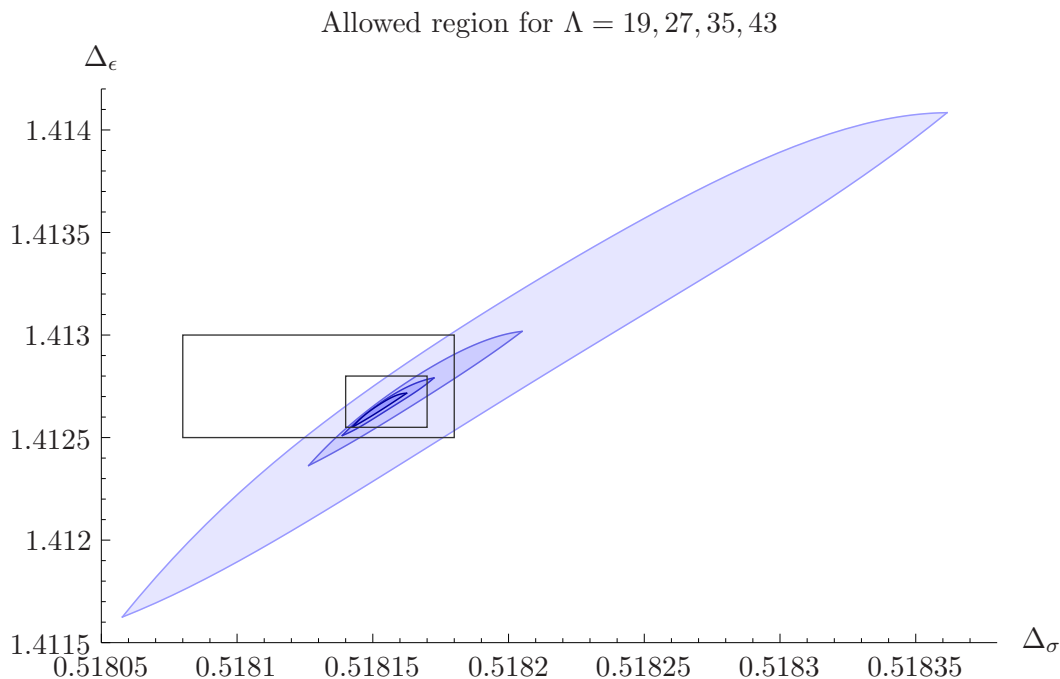


Figure 1. Allowed region for a \mathbb{Z}_2 -symmetric 3d CFT with two relevant scalars, computed using SDPB with the system of correlators $\langle\sigma\sigma\sigma\sigma\rangle$, $\langle\sigma\sigma\epsilon\epsilon\rangle$, and $\langle\epsilon\epsilon\epsilon\epsilon\rangle$. The blue regions correspond to $\Lambda = 19, 27, 35, 43$, in decreasing order of size. The larger black rectangle shows the current most precise Monte Carlo determinations of critical exponents in the 3d Ising CFT [67]. The smaller black rectangle shows the estimate for $(\Delta_\sigma, \Delta_\epsilon)$ using c -minimization at $\Lambda = 41$ for the single correlator $\langle\sigma\sigma\sigma\sigma\rangle$ [5].

`primalError < primalErrorThreshold`. The termination status of SDPB then determines whether a point is allowed or not:

$$\begin{aligned} \text{found dual feasible solution} &\implies (\Delta_\sigma, \Delta_\epsilon) \text{ disallowed,} \\ \text{found primal feasible solution} &\implies (\Delta_\sigma, \Delta_\epsilon) \text{ allowed.} \end{aligned} \tag{3.14}$$

The precise SDPB options used for the computations in this work are described in appendix A.

In figure 1, we plot the allowed regions for different numbers of derivatives labeled by $\Lambda = 19, 27, 35, 43$,¹⁰ corresponding to functionals $\bar{\alpha}$ of dimension 275, 525, 855, and 1265, respectively.¹¹ We focus on $(\Delta_\sigma, \Delta_\epsilon)$ near the 3d Ising CFT, leaving wider exploration to the future. The allowed region is an island that shrinks rapidly with increasing Λ .¹² The largest island, corresponding to $\Lambda = 19$ is the same as the allowed region in figure 5 of [9].

show that the existence of a primal feasible solution with nonsingular X precludes the existence of an optimal solution with nonzero Y . However, it is unclear why a dual feasible solution is not permitted. This may be a property of typical bootstrap problems, and not true in general.

¹⁰ $n_{\max} = 10, 14, 18, 22$ in the notation of [9].

¹¹A performance analysis for different values of Λ is given in appendix B.2.

¹²Each allowed region plotted in this work was computed by testing a grid of points and fitting curves to the boundary between allowed and disallowed gridpoints. The raw gridpoint data is available on request.

We can estimate the point towards which the islands shrink as follows. Let (a_Λ, b_Λ) be the bottom-left point of the Λ -allowed island, and similarly let (c_Λ, d_Λ) be the top-right point. Define

$$\begin{aligned} E_x(r) &= \text{stddev}_{\Lambda \in \{19, 27, 35, 43\}}(ra_\Lambda + (1-r)c_\Lambda), \\ E_y(r) &= \text{stddev}_{\Lambda \in \{19, 27, 35, 43\}}(rb_\Lambda + (1-r)d_\Lambda), \end{aligned} \quad (3.15)$$

and let r_x, r_y be the minima of E_x, E_y respectively. Our estimate is¹³

$$\begin{aligned} (\Delta_\sigma, \Delta_\epsilon) &\approx (r_x a_{43} + (1-r_x)c_{43}, r_y b_{43} + (1-r_y)d_{43}) \\ &\approx (0.5181478(5), 1.412617(4)), \end{aligned} \quad (3.16)$$

where the errors are given by $E_x(r_x), E_y(r_y)$.

The dimensions $(\Delta_\sigma, \Delta_\epsilon)$ were estimated in [5] using the conjecture that the 3d Ising CFT minimizes $c \equiv \Delta_\sigma^2 / \lambda_{\sigma\sigma T_{\mu\nu}}$ subject to the constraints of unitarity and crossing symmetry of $\langle \sigma\sigma\sigma\sigma \rangle$. This conjecture, called “ c -minimization”, is expected to be equivalent to the assumption that the 3d Ising CFT lives precisely at the kink in the dimension bound coming from the single correlator $\langle \sigma\sigma\sigma\sigma \rangle$. Although unproven, c -minimization’s advantage is that it allows one to estimate $(\Delta_\sigma, \Delta_\epsilon)$ using a single scalar correlator $\langle \sigma\sigma\sigma\sigma \rangle$. Bootstrap computations for a single scalar correlator can be made relatively efficient with a modified primal simplex algorithm [5, 70].

An advantage of multiple correlators is that it is possible to impose the condition that σ is the only relevant \mathbb{Z}_2 -odd operator, causing the allowed region in $(\Delta_\sigma, \Delta_\epsilon)$ -space to become a closed island, independent of auxiliary assumptions [9]. Because our $\Lambda = 43$ island lies within the error bars of [5], our results verify c -minimization to the precision achieved in [5].¹⁴ Our results also give further evidence for the conjecture that the 3d Ising CFT is the unique \mathbb{Z}_2 -symmetric 3d CFT with two relevant scalars and $\Delta_\sigma \lesssim 0.6$. (The precise condition on $\Delta_\sigma, \Delta_\epsilon$ depends on the shape of the allowed region further away from the 3d Ising point.) It would be very interesting to prove these conjectures analytically, perhaps by showing that the island in figure 2 shrinks to a point as $\Lambda \rightarrow \infty$. An alternative is that the conjectures are still true, but one needs information from other four-point functions to prove them.

The analysis of [9] did not use permutation symmetry of the OPE coefficient $\lambda_{\sigma\sigma\epsilon} = \lambda_{\sigma\epsilon\sigma}$.¹⁵ Including this constraint leads to an additional modest reduction in the allowed region,¹⁶ which we plot in figure 2 at $\Lambda = 43$. The resulting island gives a rigorous determination $\Delta_\sigma = 0.518151(6)$, $\Delta_\epsilon = 1.41264(6)$, which is 5–10 times more precise than the Monte Carlo results of [67]. We summarize the comparison to Monte Carlo in figure 3.

¹³For readers interested in numerology, we recommend [68]. However, see [69].

¹⁴To see this more explicitly, it should be possible to place both upper and lower bounds on c using SDPB and see that it is constrained to lie close to the minimum computed in [5]. We leave this to future work.

¹⁵Note that permutation symmetry holds only when the conformal blocks are correctly normalized as functions of Δ . For scalars, the correct normalization is $g_{\Delta,0}(u, v) = u^{\Delta/2} + \dots$ to leading order in u , up to a Δ -independent constant.

¹⁶This fact was discovered during collaboration with Filip Kos, David Poland, and Alessandro Vichi [71].

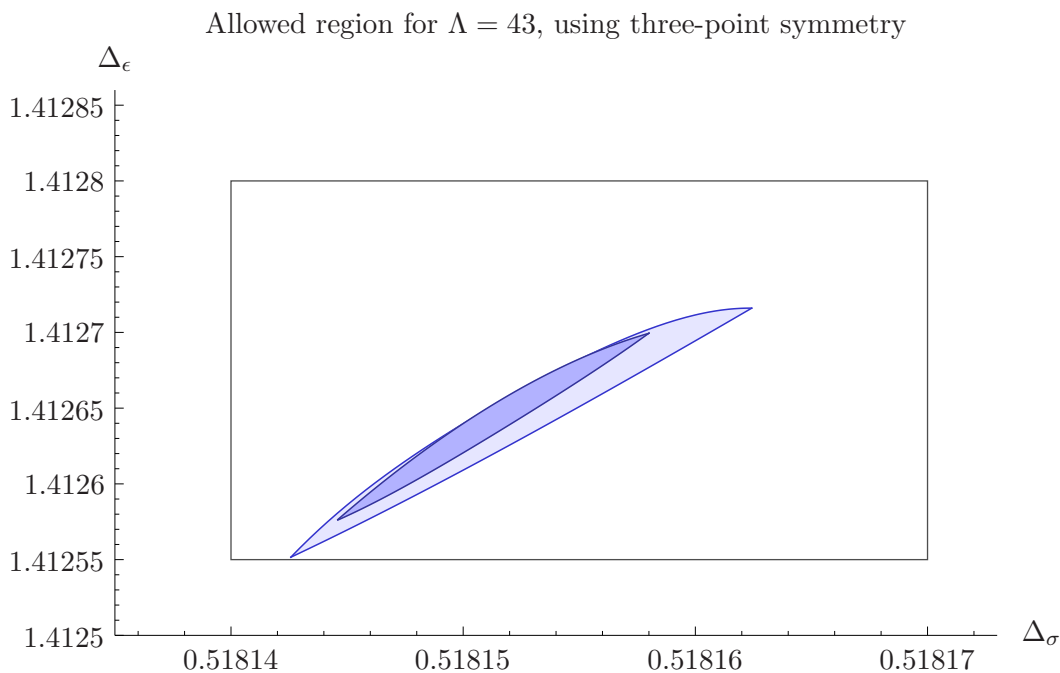


Figure 2. Allowed region for a \mathbb{Z}_2 -symmetric 3d CFT with two relevant operators, computed with SDPB at $\Lambda = 43$. The light-blue region is a zoom of the smallest region in figure 1. The darker-blue region additionally uses symmetry of the OPE coefficients $\lambda_{\sigma\sigma\epsilon} = \lambda_{\sigma\epsilon\sigma}$. The black rectangle shows the estimate for $(\Delta_\sigma, \Delta_\epsilon)$ using c -minimization at $\Lambda = 41$ [5].

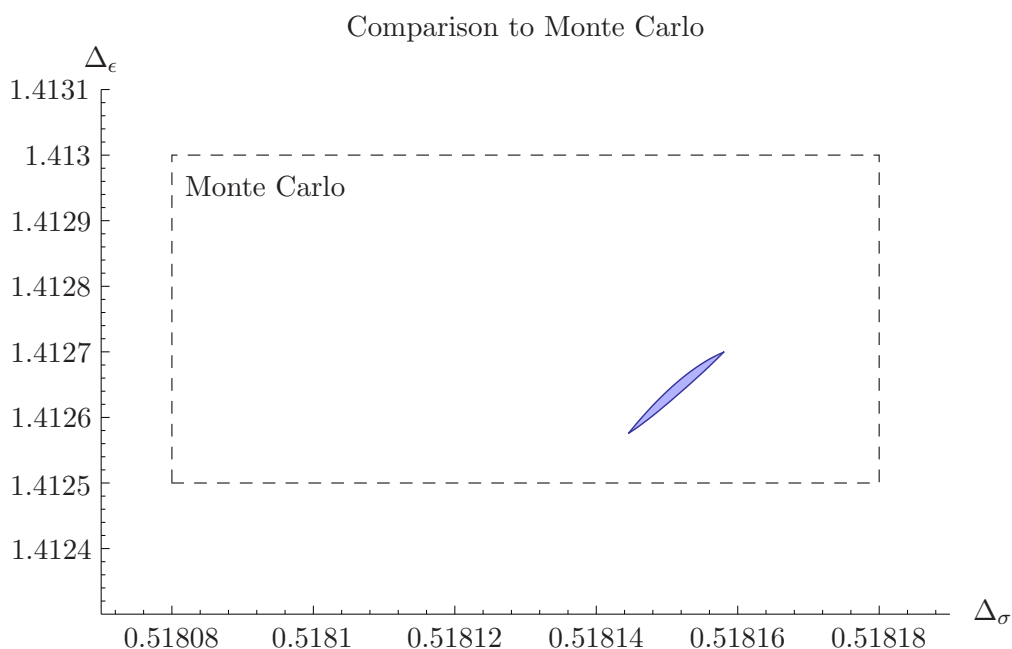


Figure 3. Comparison between the allowed region for the 3d Ising CFT using SDPB with $\Lambda = 43$ (blue) and Monte Carlo determinations of critical exponents (dashed rectangle) [67]. The size of the Monte Carlo rectangle is set by statistical and systematic errors associated with the simulation. By contrast, the blue region is a rigorous bound with sharp edges.

4 Discussion

With SDPB, we have significantly improved the precision of $(\Delta_\sigma, \Delta_\epsilon)$ in the 3d Ising CFT. Our numerics indicate that the window of allowed dimensions may shrink to a point in the limit of infinite computer time. In other words, they suggest that consistency of the correlators $\{\langle\sigma\sigma\sigma\sigma\rangle, \langle\sigma\sigma\epsilon\epsilon\rangle, \langle\epsilon\epsilon\epsilon\epsilon\rangle\}$, together with the assumption that σ and ϵ are the only relevant scalars in the theory, may uniquely fix the dimensions $(\Delta_\sigma, \Delta_\epsilon)$. This conjecture could be more tractable analytically than trying to solve the full CFT consistency conditions.

There are many more 3d Ising observables to explore. For example, the coefficient $f_{\epsilon\epsilon\epsilon}$ should be computable, and it would be interesting to compare with the recent Monte Carlo prediction [72]. It will also be important to consider larger systems of 3d Ising correlators.

However, SDPB should also enable wider exploration of new correlators and diverse theories. SDPB is already being used in several bootstrap studies that would have previously been difficult or impossible [71, 73]. An exciting direction that may now be accessible is studying a four-point function of stress-tensors in 3d CFTs.

In addition to the four-point function bootstrap, semidefinite programming has also recently been applied to the “modular bootstrap” in 2d CFTs [74, 75]. SDPB is equally applicable to modular bootstrap computations, since they too can be phrased in terms of polynomial matrix programs.

From the computing point of view, there are many opportunities for improvement. For example, it should be possible to parallelize SDPB up to hundreds of cores, which could lead to even more precise calculations, and (just as importantly) easier exploration. Very different algorithms, like Second Order Conic Programming (SOCP), cutting plane methods, or constrained nonlinear optimization may also be applicable.

The revival initiated in [1] is still young, and the technology (both analytical and numerical) is evolving rapidly. Current techniques are likely not maximally efficient, and it will be important to consider other methods, from new algorithms and optimization tools to conceptually different approaches. We are optimistic that much more will be possible.

Acknowledgments

I am grateful to Chris Beem, Luca Iliesiu, Silviu Pufu, Slava Rychkov, Balt van Rees, and Ran Yacoby for related discussions, and especially Filip Kos, David Poland, and Alessandro Vichi for discussions, collaboration, assistance testing SDPB, and comments on the draft. Thanks also to Slava Rychkov for comments on the draft. Thanks to Amir Ali Ahmadi, Hande Benson, Pablo Parrilo, and Robert Vanderbei for advice on semidefinite programming and numerical optimization. I am supported by DOE grant number DE-SC0009988 and a William D. Loughlin Membership at the Institute for Advanced Study. The computations in this paper were run on the Hyperion computing cluster supported by the School of Natural Sciences Computing Staff at the Institute for Advanced Study.

A Choices and parameters

The PMP for $\{\langle\sigma\sigma\sigma\sigma\rangle, \langle\sigma\sigma\epsilon\epsilon\rangle, \langle\epsilon\epsilon\epsilon\epsilon\rangle\}$ in the 3d Ising CFT depends on the following choices:

- An integer Λ specifying which derivatives to include in the functional $\vec{\alpha}$. These are given by $\partial_z^m \partial_{\bar{z}}^n$ with $m \geq n$ and $m+n \leq \Lambda$.¹⁷ Because of the symmetry/antisymmetry of $F_{\pm, \Delta, \ell}^{ij,kl}(u, v)$ under $u \leftrightarrow v$, some of these derivatives vanish identically. Including only non-vanishing derivatives, the dimension of $\vec{\alpha}$ is

$$\dim \vec{\alpha} = \frac{\lfloor \frac{\Lambda+2}{2} \rfloor (\lfloor \frac{\Lambda+2}{2} \rfloor + 1)}{2} + 4 \frac{\lfloor \frac{\Lambda+1}{2} \rfloor (\lfloor \frac{\Lambda+1}{2} \rfloor + 1)}{2}. \quad (\text{A.1})$$

- An integer κ controlling the accuracy of the approximation for conformal blocks (3.7). The positive prefactor is

$$\chi_\ell(\Delta) = \frac{r_*^\Delta}{\prod_i (\Delta - \Delta_{*i})} \geq 0, \quad (\text{A.2})$$

where $r_* = 3 - 2\sqrt{2}$ is the radius of the point $z = \bar{z} = \frac{1}{2}$ in the radial coordinates of [76, 77]. The dimensions Δ_{*i} are special values *below* the unitarity bound, so that the product $\prod_i (\Delta - \Delta_{*i})$ is positive for unitary theories. The approximation (3.7) can be systematically improved by including more poles $(\Delta - \Delta_{*i})^{-1}$ and increasing the degree of $p_\ell^{\Delta_{ij}, \Delta_{kl}; mn}(\Delta)$. Our choice of poles is

$$\Delta_{*i} \in \left\{ \begin{array}{l} 1 - \ell - k \quad | \quad k = 1, \dots, \kappa, \\ d/2 - k, \quad | \quad k = 1, \dots, \lfloor \kappa/2 \rfloor, \\ \ell + d - 1 - k \quad | \quad k = 1, \dots, \min(\kappa, \lfloor \ell/2 \rfloor) \end{array} \right\}. \quad (\text{A.3})$$

Smaller κ means smaller-degree polynomials and shorter runtimes. Larger κ is needed to get an accurate approximation for conformal blocks. We choose κ by computing bounds with successively larger values of κ until the results stabilize. Our final values are conservative: smaller κ may still give sufficient accuracy. Derivatives of conformal blocks were computed using the recursion relation in [9] to order r^{90} (far greater accuracy than needed).

- A set of spins $S = \{\ell_1, \dots, \ell_L\}$ to include. If not enough spins are included, the solver may find a functional $\vec{\alpha}$ that violates a positive semidefiniteness constraint for some spin. Because derivatives of conformal blocks converge as a function of spin, in practice it is sufficient to include a finite number of spins to ensure $\vec{\alpha}$ satisfies the constraints for all spins (as can be verified post-hoc by testing $\vec{\alpha}$ on constraints that were not explicitly included). This sufficient number of spins grows with Λ . Our choices are given in (A.4).

$$\begin{aligned} S_{\Lambda=19} &= \{0, \dots, 26\} \cup \{49, 50\} \\ S_{\Lambda=27} &= \{0, \dots, 26\} \cup \{29, 30, 33, 34, 37, 38, 41, 42, 45, 46, 49, 50\} \\ S_{\Lambda=35} &= \{0, \dots, 44\} \cup \{47, 48, 51, 52, 55, 56, 59, 60, 63, 64, 67, 68\} \\ S_{\Lambda=43} &= \{0, \dots, 64\} \cup \{67, 68, 71, 72, 75, 76, 79, 80, 83, 84, 87, 88\}. \end{aligned} \quad (\text{A.4})$$

¹⁷ $\Lambda = 2n_{\max} - 1$, where n_{\max} is the parameter defined in [9].

Λ	19	27	35	43
κ	14	20	30	40
spins	$S_{\Lambda=19}$	$S_{\Lambda=27}$	$S_{\Lambda=35}$	$S_{\Lambda=43}$
precision	448	576	768	960
findPrimalFeasible	True	True	True	True
findDualFeasible	True	True	True	True
detectPrimalFeasibleJump	True	True	True	True
detectDualFeasibleJump	True	True	True	True
dualityGapThreshold	10^{-30}	10^{-30}	10^{-30}	10^{-75}
primalErrorThreshold	10^{-30}	10^{-30}	10^{-40}	10^{-75}
dualErrorThreshold	10^{-30}	10^{-30}	10^{-40}	10^{-75}
initialMatrixScalePrimal ($\Omega_{\mathcal{P}}$)	10^{40}	10^{50}	10^{50}	10^{60}
initialMatrixScaleDual ($\Omega_{\mathcal{D}}$)	10^{40}	10^{50}	10^{50}	10^{60}
feasibleCenteringParameter (β_{feasible})	0.1	0.1	0.1	0.1
infeasibleCenteringParameter ($\beta_{\text{infeasible}}$)	0.3	0.3	0.3	0.3
stepLengthReduction (γ)	0.7	0.7	0.7	0.7
choleskyStabilizeThreshold (θ)	10^{-40}	10^{-40}	10^{-100}	10^{-140}
maxComplementarity	10^{100}	10^{130}	10^{160}	10^{200}

Table 1. Parameters for the computations in this work. Only SDPB parameters that affect the numerics (as opposed to parameters like `maxThreads` and `maxRuntime`) are included. The sets of spins S_{Λ} are given in (A.4). Variables in the interior point algorithm of section 2.4.4 that correspond to SDPB parameters are indicated in parentheses. `precision` is in binary digits.

Once these quantities are fixed, the parameters to SDPB must be chosen to ensure numerical stability, precision, and correctness. Our choices for the computations in this work are summarized in table 1.

B Performance

B.1 Complexity comparison

Let us compare the complexity of SDPB’s algorithm to that of SDPA-GMP for solving PMPs. For simplicity, suppose that each polynomial matrix $M_j^n(x)$ has size $m \times m$ and degree d . The matrices X, Y then have $2J$ blocks, each of dimension $m(d+1)$. We focus on the most expensive parts of each iteration of the interior point algorithm and count the number of multiplications to leading order in J, m , and d .

For SDPA-GMP, we assume the setup described in [9, 17]. There, the free variables y were embedded as 1×1 diagonal blocks in the matrix Y . (By contrast, for SDPB they are treated separately.) The most important contributions to the running time of SDPA-GMP are as follows.

- The Schur complement matrix is dense, so each of its elements must be computed individually. Computing $X^{-1}A_qY$ requires $N + 2J(md)^3$ multiplications, since it involves a block-diagonal dense matrix multiplication $X^{-1} \times (A_qY)$ (the

linear term in N comes from multiplying 1×1 blocks). This must be repeated $\frac{Jm^2d}{2}$ times: once for each q . Now, the matrices A_p are typically sparse, with $O(md)$ entries. Thus, evaluating all the traces $\text{Tr}(A_p X^{-1} A_q Y)$ requires approximately $(N + 2J(md)^3) \frac{Jm^2d}{2} + \frac{1}{2} (\frac{Jm^2d}{2})^2 (md)$ multiplications. These steps dominate the running time for the computations in [9].

- Because the Schur complement matrix is dense, it must be inverted using a full Cholesky decomposition, which requires $\frac{1}{3} (\frac{Jm^2d}{2})^3$ multiplications.

For SDPB, computing S takes negligible time. The most important steps are in solving the Schur complement equation:

- Computing the Cholesky decomposition $S = LL^T$ takes $\frac{1}{3} J (\frac{m^2d}{2})^3$ multiplications, since it can be done block-wise.
- Forming $L^{-1}B$ requires $NJ \frac{m^2d}{2}$ multiplications.
- Forming $Q = (L^{-1}B)^T (L^{-1}B)$ requires $\frac{1}{2} N^2 (J \frac{m^2d}{2})$ multiplications. This step dominates the running time for the computations in this work.
- Computing the LU decomposition of Q requires $\frac{2}{3} N^3$ multiplications.

B.2 Running time for 3d Ising computations

For the 3d Ising computations in this work, we have

$$\begin{aligned}
 d &\approx \Lambda + \frac{5}{2} \kappa \text{ (for large enough } \ell), \\
 m &= 1 \text{ or } 2, \\
 J &\approx \text{number of included spins.}
 \end{aligned}
 \tag{B.1}$$

A full analysis of the complexity in Λ would require determining the correct asymptotics of each quantity (including SDPB parameters like `precision`), which may depend on the computation at hand. In table 2, we simply report average runtimes and approximate multiplications/iteration for the choices given in appendix A.

Solver	Λ	Runtime (dual feasible)	Runtime (dual infeasible)	Mul./iter.
SDPB	19	3.5	0.9	2×10^8
SDPB	27	32	7.6	1×10^9
SDPB	35	190	40	5×10^9
SDPB	43	810	260	2×10^{10}
SDPA-GMP	19	~ 300	~ 300	1×10^{11}
SDPA-GMP	27	–	–	1×10^{12}
SDPA-GMP	35	–	–	8×10^{12}
SDPA-GMP	43	–	–	5×10^{13}

Table 2. Runtimes for a single feasibility computation in the 3d Ising CFT using the correlators $\{\langle\sigma\sigma\sigma\sigma\rangle, \langle\sigma\sigma\epsilon\epsilon\rangle, \langle\epsilon\epsilon\epsilon\epsilon\rangle\}$, as described in [9] and appendix A. Average runtimes are different depending on whether a spectrum is disallowed (dual feasible) or allowed (dual infeasible). All times are in CPU-hours (for SDPB, this means the actual runtime is multiplied by `maxThreads`, which was 16 for most of the computations in this work). Approximate SDPA-GMP times are from [9]. All computations were performed on 3.3 GHz 64-bit Intel Xeon Sandy Bridge processors. The column “mul./iter.” gives the approximate number of multiplications per iteration, calculated according to the discussion in subsection B.1. (To estimate running time from the number of multiplications per iteration, one needs to take into account `precision`, which also increases with Λ .) The SDPA-GMP computations with $\Lambda > 19$ have not been attempted.

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] 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](#)].
- [2] V.S. Rychkov and A. Vichi, *Universal constraints on conformal operator dimensions*, *Phys. Rev. D* **80** (2009) 045006 [[arXiv:0905.2211](https://arxiv.org/abs/0905.2211)] [[INSPIRE](#)].
- [3] S. Rychkov, *Conformal bootstrap in three dimensions?*, [arXiv:1111.2115](https://arxiv.org/abs/1111.2115) [[INSPIRE](#)].
- [4] S. El-Showk et al., *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](#)].
- [5] S. El-Showk et al., *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](#)].
- [6] 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](https://arxiv.org/abs/1211.2810)] [[INSPIRE](#)].
- [7] F. Kos, D. Poland and D. Simmons-Duffin, *Bootstrapping the $O(N)$ vector models*, *JHEP* **06** (2014) 091 [[arXiv:1307.6856](https://arxiv.org/abs/1307.6856)] [[INSPIRE](#)].
- [8] C. Beem, L. Rastelli and B.C. van Rees, *The $\mathcal{N} = 4$ superconformal bootstrap*, *Phys. Rev. Lett.* **111** (2013) 071601 [[arXiv:1304.1803](https://arxiv.org/abs/1304.1803)] [[INSPIRE](#)].
- [9] F. Kos, D. Poland and D. Simmons-Duffin, *Bootstrapping mixed correlators in the 3D Ising model*, *JHEP* **11** (2014) 109 [[arXiv:1406.4858](https://arxiv.org/abs/1406.4858)] [[INSPIRE](#)].

- [10] 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].
- [11] A.M. Polyakov, *Nonhamiltonian approach to conformal quantum field theory*, *Zh. Eksp. Teor. Fiz.* **66** (1974) 23 [INSPIRE].
- [12] F. Caracciolo and V.S. Rychkov, *Rigorous limits on the interaction strength in quantum field theory*, *Phys. Rev.* **D 81** (2010) 085037 [arXiv:0912.2726] [INSPIRE].
- [13] D. Poland and D. Simmons-Duffin, *Bounds on 4D conformal and superconformal field theories*, *JHEP* **05** (2011) 017 [arXiv:1009.2087] [INSPIRE].
- [14] R. Rattazzi, S. Rychkov and A. Vichi, *Central charge bounds in 4D conformal field theory*, *Phys. Rev.* **D 83** (2011) 046011 [arXiv:1009.2725] [INSPIRE].
- [15] R. Rattazzi, S. Rychkov and A. Vichi, *Bounds in 4D conformal field theories with global symmetry*, *J. Phys.* **A 44** (2011) 035402 [arXiv:1009.5985] [INSPIRE].
- [16] A. Vichi, *Improved bounds for CFT's with global symmetries*, *JHEP* **01** (2012) 162 [arXiv:1106.4037] [INSPIRE].
- [17] D. Poland, D. Simmons-Duffin and A. Vichi, *Carving out the space of 4D CFTs*, *JHEP* **05** (2012) 110 [arXiv:1109.5176] [INSPIRE].
- [18] P. Liendo, L. Rastelli and B.C. van Rees, *The bootstrap program for boundary CFT_d*, *JHEP* **07** (2013) 113 [arXiv:1210.4258] [INSPIRE].
- [19] S. El-Showk et al., *Conformal field theories in fractional dimensions*, *Phys. Rev. Lett.* **112** (2014) 141601 [arXiv:1309.5089] [INSPIRE].
- [20] L.F. Alday and A. Bissi, *The superconformal bootstrap for structure constants*, *JHEP* **09** (2014) 144 [arXiv:1310.3757] [INSPIRE].
- [21] D. Gaiotto, D. Mazac and M.F. Paulos, *Bootstrapping the 3d Ising twist defect*, *JHEP* **03** (2014) 100 [arXiv:1310.5078] [INSPIRE].
- [22] D. Bashkirov, *Bootstrapping the $\mathcal{N} = 1$ SCFT in three dimensions*, arXiv:1310.8255 [INSPIRE].
- [23] M. Berkooz, R. Yacoby and A. Zait, *Bounds on $\mathcal{N} = 1$ superconformal theories with global symmetries*, *JHEP* **08** (2014) 008 [arXiv:1402.6068] [INSPIRE].
- [24] Y. Nakayama and T. Ohtsuki, *Approaching the conformal window of $O(n) \times O(m)$ symmetric Landau-Ginzburg models using the conformal bootstrap*, *Phys. Rev.* **D 89** (2014) 126009 [arXiv:1404.0489] [INSPIRE].
- [25] Y. Nakayama and T. Ohtsuki, *Five dimensional $O(N)$ -symmetric CFTs from conformal bootstrap*, *Phys. Lett.* **B 734** (2014) 193 [arXiv:1404.5201] [INSPIRE].
- [26] S.M. Chester, J. Lee, S.S. Pufu and R. Yacoby, *The $\mathcal{N} = 8$ superconformal bootstrap in three dimensions*, *JHEP* **09** (2014) 143 [arXiv:1406.4814] [INSPIRE].
- [27] F. Caracciolo, A.C. Echeverri, B. von Harling and M. Serone, *Bounds on OPE coefficients in 4D conformal field theories*, *JHEP* **10** (2014) 020 [arXiv:1406.7845] [INSPIRE].
- [28] L.F. Alday and A. Bissi, *Generalized bootstrap equations for $\mathcal{N} = 4$ SCFT*, *JHEP* **02** (2015) 101 [arXiv:1404.5864] [INSPIRE].
- [29] Y. Nakayama and T. Ohtsuki, *Bootstrapping phase transitions in QCD and frustrated spin systems*, *Phys. Rev.* **D 91** (2015) 021901 [arXiv:1407.6195] [INSPIRE].

- [30] S.M. Chester, J. Lee, S.S. Pufu and R. Yacoby, *Exact correlators of BPS operators from the 3d superconformal bootstrap*, *JHEP* **03** (2015) 130 [[arXiv:1412.0334](#)] [[INSPIRE](#)].
- [31] J.-B. Bae and S.-J. Rey, *Conformal bootstrap approach to $O(N)$ fixed points in five dimensions*, [arXiv:1412.6549](#) [[INSPIRE](#)].
- [32] C. Beem, M. Lemos, P. Liendo, L. Rastelli and B.C. van Rees, *The $\mathcal{N} = 2$ superconformal bootstrap*, [arXiv:1412.7541](#) [[INSPIRE](#)].
- [33] S.M. Chester, S.S. Pufu and R. Yacoby, *Bootstrapping $O(N)$ vector models in $4 < d < 6$* , *Phys. Rev. D* **91** (2015) 086014 [[arXiv:1412.7746](#)] [[INSPIRE](#)].
- [34] M.S. Costa, J. Penedones, D. Poland and S. Rychkov, *Spinning conformal correlators*, *JHEP* **11** (2011) 071 [[arXiv:1107.3554](#)] [[INSPIRE](#)].
- [35] M.S. Costa, J. Penedones, D. Poland and S. Rychkov, *Spinning conformal blocks*, *JHEP* **11** (2011) 154 [[arXiv:1109.6321](#)] [[INSPIRE](#)].
- [36] D. Simmons-Duffin, *Projectors, shadows and conformal blocks*, *JHEP* **04** (2014) 146 [[arXiv:1204.3894](#)] [[INSPIRE](#)].
- [37] A. Dymarsky, *On the four-point function of the stress-energy tensors in a CFT*, [arXiv:1311.4546](#) [[INSPIRE](#)].
- [38] C. Beem et al., *Infinite chiral symmetry in four dimensions*, *Commun. Math. Phys.* **336** (2015) 1359 [[arXiv:1312.5344](#)] [[INSPIRE](#)].
- [39] A.L. Fitzpatrick et al., *Covariant approaches to superconformal blocks*, *JHEP* **08** (2014) 129 [[arXiv:1402.1167](#)] [[INSPIRE](#)].
- [40] Z.U. Khandker, D. Li, D. Poland and D. Simmons-Duffin, *$\mathcal{N} = 1$ superconformal blocks for general scalar operators*, *JHEP* **08** (2014) 049 [[arXiv:1404.5300](#)] [[INSPIRE](#)].
- [41] Wolfram Research, *Mathematica*, Wolfram Research Inc., Champaign U.S.A. (2014).
- [42] *GNU Linear Programming Kit*, <http://www.gnu.org/software/glpk/glpk.html>.
- [43] *IBM ILOG CPLEX Optimizer* (2010), <http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>.
- [44] M. Yamashita, K. Fujisawa, M. Fukuda, K. Nakata and M. Nakata, *A high-performance software package for semidefinite programs: SDPA 7*, Research Report B-463, Dept. of Mathematical and Computing Science, Tokyo Institute of Technology, Tokyo Japan (2010).
- [45] M. Yamashita, K. Fujisawa and M. Kojima, *Implementation and evaluation of SDPA 6.0 (SemiDefinite Programming Algorithm 6.0)*, *Optim. Method. Softw.* **18** (2003) 491.
- [46] M. Nakata, *A numerical evaluation of highly accurate multiple-precision arithmetic version of semidefinite programming solver: SDPA-GMP, -QD and -DD*, *IEEE Int. Symp. CACSD* (2010) 29.
- [47] Y. Nesterov and A. Nemirovskii, *Interior-point polynomial algorithms in convex programming*, Society for Industrial and Applied Mathematics (1994) [[doi:10.1137/1.9781611970791](#)].
- [48] F. Alizadeh, *Combinatorial optimization with interior point methods and semidefinite matrices*, Ph.D. Thesis, University of Minnesota, Minneapolis U.S.A. (1991).
- [49] C. Helmberg, F. Rendl, R.J. Vanderbei and H. Wolkowicz, *An interior-point method for semidefinite programming*, *SIAM J. Optim.* **6** (1996) 342.

- [50] L. Vandenberghe and S. Boyd, *A primal-dual potential reduction method for problems involving matrix inequalities*, *Math. Program.* **69** (1995) 205.
- [51] C. Hanselka and M. Schweighofer, *Matrix polynomials positive semidefinite on intervals*, to appear.
- [52] M.-D. Choi, T.-Y. Lam and B. Reznick, *Real zeros of positive semidefinite forms. I*, *Math. Z.* **171** (1980) 1.
- [53] E.M. Aylward, S.M. Itani and P.A. Parrilo, *Explicit SOS decompositions of univariate polynomial matrices and the Kalman-Yakubovich-Popov lemma*, *IEEE Conf. Decis. Contr.* (2007) 5660.
- [54] M.J. Todd, *A study of search directions in primal-dual interior-point methods for semidefinite programming*, *Optim. Method. Softw.* **11** (1999) 1.
- [55] M. Kojima, S. Shindoh and S. Hara, *Interior-point methods for the monotone semidefinite linear complementarity problem in symmetric matrices*, *SIAM J. Optim.* **7** (1997) 86.
- [56] R.D. Monteiro, *Primal-dual path-following algorithms for semidefinite programming*, *SIAM J. Optim.* **7** (1997) 663.
- [57] S. Mehrotra, *On the implementation of a primal-dual interior point method*, *SIAM J. Optim.* **2** (1992) 575.
- [58] Z. Liu and L. Vandenberghe, *Low-rank structure in semidefinite programs derived from the KYP lemma*, *IEEE Decis. Contr.* (2007) 5652.
- [59] M.F. Anjos and S. Burer, *On handling free variables in interior-point methods for conic linear optimization*, *SIAM J. Optim.* **18** (2007) 1310.
- [60] K. Kobayashi, K. Nakata and M. Kojima, *A conversion of an SDP having free variables into the standard form SDP*, *Comput. Optim. Appl.* **36** (2007) 289.
- [61] C. Lanczos, *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, *J. Res. Nat. Bur. Std.* **45** (1950) 225.
- [62] T. Granlund and the GMP Development Team, *GNU MP: the GNU Multiple Precision Arithmetic Library*, ed. 5.0.5 (2012), <http://gmplib.org/>.
- [63] M. Nakata, *The MPACK (MBLAS/MLAPACK); a multiple precision arithmetic version of BLAS and LAPACK*, ed. 0.6.7 (2010), <http://mplapack.sourceforge.net/>.
- [64] C++ Standards Committee Library Working Group and other contributors, *BOOST C++ libraries*, <http://www.boost.org>.
- [65] L. Thomason, *TinyXML2*, <http://www.grinninglizard.com/tinyxml2docs/index.html>.
- [66] OpenMP Architecture Review Board, *OpenMP application program interface*, version 3.0 (2008), <http://www.openmp.org/mp-documents/spec30.pdf>.
- [67] M. Hasenbusch, *Finite size scaling study of lattice models in the three-dimensional Ising universality class*, *Phys. Rev. B* **82** (2010) 174433 [[arXiv:1004.4486](https://arxiv.org/abs/1004.4486)].
- [68] P. Borwein et al., *The inverse symbolic calculator* (1995), <http://oldweb.cecm.sfu.ca/projects/ISC/ISCmain.html>.
- [69] F. Lenz, *The ratio of proton and electron masses*, *Phys. Rev.* **82** (1951) 554.
- [70] M.F. Paulos, *JuliBootS: a hands-on guide to the conformal bootstrap*, [arXiv:1412.4127](https://arxiv.org/abs/1412.4127) [[INSPIRE](https://arxiv.org/abs/1412.4127)].

- [71] F. Kos, D. Poland, D. Simmons-Duffin and A. Vichi, *Bootstrapping the $O(N)$ archipelago*, [arXiv:1504.07997](#) [[INSPIRE](#)].
- [72] M. Caselle, G. Costagliola and N. Magnoli, *Numerical determination of the operator-product-expansion coefficients in the 3D Ising model from off-critical correlators*, *Phys. Rev. D* **91** (2015) 061901 [[arXiv:1501.04065](#)] [[INSPIRE](#)].
- [73] L. Iliesiu et al., *Bootstrapping 3D fermions*, to appear.
- [74] S. Hellerman, *A universal inequality for CFT and quantum gravity*, *JHEP* **08** (2011) 130 [[arXiv:0902.2790](#)] [[INSPIRE](#)].
- [75] D. Friedan and C.A. Keller, *Constraints on 2d CFT partition functions*, *JHEP* **10** (2013) 180 [[arXiv:1307.6562](#)] [[INSPIRE](#)].
- [76] D. Pappadopulo, S. Rychkov, J. Espin and R. Rattazzi, *OPE convergence in conformal field theory*, *Phys. Rev. D* **86** (2012) 105043 [[arXiv:1208.6449](#)] [[INSPIRE](#)].
- [77] M. Hogervorst and S. Rychkov, *Radial coordinates for conformal blocks*, *Phys. Rev. D* **87** (2013) 106004 [[arXiv:1303.1111](#)] [[INSPIRE](#)].