

SDPB 1.0

Contents

1	Introduction	1
1.1	Installation and Requirements	2
2	Polynomial Matrix Programs	2
3	Input to SDPB	2
3.1	Input Format	3
3.2	Mathematica Interface	4
3.3	An Example	6
4	Internal SDP	7
5	Output of SDPB	9
5.1	Terminal Output	9
5.2	Termination	10
5.3	Output File	11
5.4	Checkpoints	12
6	Attribution	13
7	Acknowledgements	13

1 Introduction

SDPB is an arbitrary-precision semidefinite program solver, specialized for “polynomial matrix programs” (defined below). This document describes SDPB’s usage and input/output. Much more detail about its design is given in [1]. The reader is encouraged to look there for a better understanding of SDPB’s parameters and internal operation.

1.1 Installation and Requirements

SDPB requires

- [The Boost C++ Libraries](#) (tested with Boost 1.54).
- [The GNU Multiprecision Library](#).

To install, you must first edit the `Makefile` to define the variables `GMPINCLUDEDIR`, `BOOSTINCLUDEDIR`, and `LIBDIR`. Then type `make` to build the `sdpb` executable.

2 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

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

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

$$\begin{aligned} & \text{maximize } b_0 + b \cdot y \quad \text{over } y \in \mathbb{R}^N, \\ & \text{such that } M_j^0(x) + \sum_{n=1}^N y_n M_j^n(x) \succeq 0 \quad \text{for all } x \geq 0 \text{ and } 1 \leq j \leq J. \end{aligned} \quad (2.2)$$

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

3 Input to SDPB

SDPB takes 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 $b_0 \in \mathbb{R}$ and $b \in \mathbb{R}^N$.

A bilinear basis is a collection of polynomials $q_m^{(j)}(x)$ such that $\deg q_m^{(j)} = m$, for example monomials $q_m^{(j)}(x) = x^m$. (A better choice for numerical stability is usually orthogonal polynomials on the positive real line.) The sample points and sample scalings determine how the PMP is represented internally as an SDP. In principle, they do not affect the solution of the PMP, but in practice they can affect numerical stability. The constant b_0 is completely irrelevant to the solution algorithm, but is included for convenience. See [1] for details.

3.1 Input Format

SDPB reads the data above in the following XML format.

Listing 1: XML input format for SDPB

```

input to SDPB ≡
  <sdp>
    <xml for objective>
    <xml for polynomial vector matrices>
  </sdp>

xml for objective ≡
  <objective>
    <elt>b0</elt>
    ...
    <elt>bN</elt>
  </objective>

xml for polynomial vector matrices ≡
  <polynomialVectorMatrices>
    <xml for polynomial vector matrix M1n(x)>
    ...
    <xml for polynomial vector matrix Mjn(x)>
  </polynomialVectorMatrices>

xml for polynomial vector matrix Mjn(x) ≡
  <polynomialVectorMatrix>
    <rows>mj</rows>
    <cols>mj</cols>
    <elements>
      <xml for polynomial vector Pj,11n(x)>
      ...
      <xml for polynomial vector Pj,mj1n(x)>
      ...
      <xml for polynomial vector Pj,1mjn(x)>
      ...
      <xml for polynomial vector Pj,mjmjn(x)>
    </elements>

```

```

<samplePoints>
  <elt> $x_0^{(j)}$ </elt>
  ...
  <elt> $x_{d_j}^{(j)}$ </elt>
</samplePoints>
<sampleScalings>
  <elt> $s_0^{(j)}$ </elt>
  ...
  <elt> $s_{d_j}^{(j)}$ </elt>
</sampleScalings>
<bilinearBasis>
  <xml for polynomial  $q_0^{(j)}(x)$ >
  ...
  <xml for polynomial  $q_{\lfloor d_j/2 \rfloor}^{(j)}(x)$ >
</bilinearBasis>
</polynomialVectorMatrix>

xml for polynomial vector  $P_{j,rs}^n(x) \equiv$ 
<polynomialVector>
  <xml for polynomial  $P_{j,rs}^0(x)$ >
  ...
  <xml for polynomial  $P_{j,rs}^N(x)$ >
</polynomialVector>

xml for polynomial  $a_0 + a_1x + \dots + a_dx^d \equiv$ 
<polynomial>
  <coeff> $a_0$ </coeff>
  ...
  <coeff> $a_d$ </coeff>
</polynomial>

```

Several aspects of this format are inefficient. Because the matrices are symmetric, `rows` and `cols` are redundant, and most elements are listed twice. Also, XML is extremely verbose. The current choices are in the interest of simplicity and could obviously be changed in a future version.

The options to SDPB are described in detail in the help text, obtained by running “`sdpb --help`.”

3.2 Mathematica Interface

A Mathematica notebook `SDPB.m`, included in the source distribution, generates files of the form in listing 1 starting from Mathematica data. It automatically makes sensible choices for the bilinear bases $q_m^{(j)}(x)$, the sample points $x_k^{(j)}$ and the sample scalings $s_k^{(j)}$.

The `Mathematica` definition of a PMP is slightly different but trivially equivalent to (2.2). It is:

$$\begin{aligned} & \text{maximize} && a \cdot z && \text{over} && z \in \mathbb{R}^{N+1}, \\ & \text{such that} && \sum_{n=0}^N z_n W_j^n(x) \succeq 0 && \text{for all } x \geq 0 \text{ and } 1 \leq j \leq J, \\ & && n \cdot z = 1. \end{aligned} \tag{3.1}$$

where $W_j^n(x)$ are matrix polynomials. The normalization condition $n \cdot z = 1$ can be used to solve for one of the components of z in terms of the others. Calling the remaining components $y \in \mathbb{R}^N$, we arrive at (2.2), where $M_j^n(x)$ are linear combinations of $W_j^n(x)$ and b_0, b_n are linear combinations of the a_n . This difference in convention is for convenient use in the conformal bootstrap.

`SDPB.m` defines a function `WriteBootstrapSDP[file, sdp]`, where `file` is the XML file to be written to, and `sdp` has the following form, where the polynomials $Q_{j,rs}^n(x)$ are the elements of $W_j^n(x)$.

Listing 2: Usage of `WriteBootstrapSDP` in `SDPB.m`

```
function call ≡ WriteBootstrapSDP[file, ⟨sdp⟩]

sdp ≡ SDP[⟨objective⟩, ⟨normalization⟩, ⟨positive matrices with prefactors⟩]

objective ≡ {a0, ..., aN}

normalization ≡ {n0, ..., nN}

positive matrices with prefactors ≡ {
  ⟨positive matrix with prefactor 1⟩,
  ...
  ⟨positive matrix with prefactor J⟩,
}

positive matrix with prefactor j ≡
PositiveMatrixWithPrefactor[⟨prefactor⟩,
{
  {
    {Qj,110(x), ..., Qj,11N(x)}, ..., {Qj,mj10(x), ..., Qj,mj1N(x)}
  },
  ...
  {
    {Qj,1mj0(x), ..., Qj,1mjN(x)}, ..., {Qj,mjmj0(x), ..., Qj,mjmjN(x)}
  },
}
]

prefactor ≡
DampedRational[c, {p1, ..., pk}, b, x]
or
```

const

The prefactor in `PositiveMatrixWithPrefactor` is used for constructing bilinear bases and sample scalings. Specifically, if the prefactor is $\chi(x)$, the bilinear basis is a set of orthogonal polynomials with respect to measure $\chi(x)dx$ on the positive real line, and sample scalings are $\chi(x_k)$, where the x_k are sample points. The notebook `SDPB.m` only deals with damped-rational prefactors because these are relevant to the conformal bootstrap. These stand for

$$\text{DampedRational}[c, \{p_1, \dots, p_k\}, b, x] \rightarrow c \frac{b^x}{\prod_{i=1}^k (x - p_i)}. \quad (3.2)$$

We do not use an exponential-times-rational `Mathematica` function directly because the `DampedRational` data structure makes it easier to extract information needed to construct a bilinear basis. The notebook `SDPB.m` makes a choice of sample points that are reasonable for conformal bootstrap applications.

As an example bootstrap application, the included notebook `Bootstrap2dExample.m` computes a single-correlator dimension bound for 2d CFTs with a \mathbb{Z}_2 symmetry, as in [2].

3.3 An Example

Let's look at an example. Consider the following problem: maximize $-y$ such that

$$1 + x^4 + y \left(\frac{x^4}{12} + x^2 \right) \geq 0 \quad \text{for all } x \geq 0 \quad (3.3)$$

This is an PMP with 1×1 positive-semidefiniteness constraints. We will arbitrarily choose a prefactor of $e^{-x} = \text{DampedRational}[1, \{\}, 1/E, x]$, so that the bilinear basis consists of Laguerre polynomials. The `Mathematica` code for this example is

Listing 3: Mathematica input for the example (3.3)

```
Module[
{
  polys = {
    PositiveMatrixWithPrefactor[
      DampedRational[1, {}, 1/E, x],
      {{{1 + x^4, x^4/12 + x^2}}}
    ]
  },
  norm = {1, 0},
  obj = {0, -1}
},
WriteBootstrapSDP["test.xml", SDP[obj, norm, polys]];
];
```

It produces the following XML file

Listing 4: XML file test.xml produced by listing 3. Decimals are truncated at 12 digits.

```

<sdp>
  <objective><elt>0</elt><elt>-1</elt></objective>
  <polynomialVectorMatrices>
    <polynomialVectorMatrix>
      <rows>1</rows>
      <cols>1</cols>
      <elements>
        <polynomialVector>
          <polynomial>
            <coeff>1</coeff><coeff>0</coeff><coeff>0</coeff>
            <coeff>0</coeff><coeff>1</coeff>
          </polynomial>
          <polynomial>
            <coeff>0</coeff><coeff>0</coeff><coeff>1</coeff>
            <coeff>0</coeff><coeff>0.0833333333333</coeff>
          </polynomial>
        </polynomialVector>
      </elements>
    <samplePoints>
      <elt>0.017496844815</elt><elt>0.157471603340</elt><elt>0.857345395967</elt>
      <elt>2.117118222694</elt><elt>3.936790083523</elt>
    </samplePoints>
    <sampleScalings>
      <elt>0.982655336118</elt><elt>0.854301072560</elt><elt>0.424286902403</elt>
      <elt>0.120378031823</elt><elt>0.019510742190</elt>
    </sampleScalings>
    <bilinearBasis>
      <polynomial><coeff>1</coeff></polynomial>
      <polynomial><coeff>-1</coeff><coeff>1</coeff></polynomial>
      <polynomial><coeff>1</coeff><coeff>-2</coeff><coeff>0.5</coeff></polynomial>
    </bilinearBasis>
  </polynomialVectorMatrix>
</polynomialVectorMatrices>
</sdp>

```

4 Internal SDP

To understand the output of SDPB, we need a rough understanding of its internal representation of the above PMP as a semidefinite program (SDP). Much more detail is given in [1]. The PMP (2.2) is translated into a dual pair of SDPs of the following form:

$$\begin{aligned}
 \mathcal{D} : \quad & \text{maximize} \quad \text{Tr}(CY) + b_0 + b \cdot y \quad \text{over} \quad y \in \mathbb{R}^N, Y \in \mathcal{S}^K, \\
 & \text{such that} \quad \text{Tr}(A_*Y) + By = c, \text{ and} \\
 & Y \succeq 0.
 \end{aligned} \tag{4.1}$$

$$\begin{aligned}
\mathcal{P} : \quad & \text{minimize} \quad b_0 + c \cdot x \quad \text{over} \quad x \in \mathbb{R}^P, X \in \mathcal{S}^K, \\
& \text{such that} \quad X = \sum_{p=1}^P A_p x_p - C, \\
& \quad \quad \quad B^T x = b, \\
& \quad \quad \quad X \succeq 0,
\end{aligned} \tag{4.2}$$

where “ $\succeq 0$ ” means “is positive-semidefinite” and

$$\begin{aligned}
c & \in \mathbb{R}^P, \\
B & \in \mathbb{R}^{P \times N}, \\
A_1, \dots, A_P, C & \in \mathcal{S}^K.
\end{aligned} \tag{4.3}$$

Here, \mathcal{S}^K is the space of $K \times K$ symmetric real matrices, and $\text{Tr}(A_* Y)$ denotes the vector $(\text{Tr}(A_1 Y), \dots, \text{Tr}(A_P Y)) \in \mathbb{R}^P$. An optimal solution to (4.1) and (4.2) is characterized by $XY = 0$ and also equality of the primal and dual objective functions $\text{Tr}(CY) + b_0 + b \cdot y = b_0 + c \cdot x$.

The residues

$$\begin{aligned}
P & \equiv \sum_i A_i x_i - X - C, \\
p & \equiv b - B^T x, \\
d & \equiv c - \text{Tr}(A_* Y) - B y,
\end{aligned} \tag{4.4}$$

measure the failure of x, X, y, Y to satisfy their constraints. We say a point $q = (x, X, y, Y)$ is “primal feasible” or “dual feasible” if the residues are sufficiently small,

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

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

$$\begin{aligned}
\text{dualityGap} & \equiv \frac{|\text{primalObjective} - \text{dualObjective}|}{\max\{1, |\text{primalObjective} + \text{dualObjective}|\}}, \\
\text{primalObjective} & \equiv b_0 + c \cdot x, \\
\text{dualObjective} & \equiv \text{Tr}(CY) + b_0 + b \cdot y.
\end{aligned} \tag{4.5}$$

A point is considered “optimal” if

$$\text{dualityGap} < \text{dualityGapThreshold}, \tag{4.6}$$

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

mu: The value of the complementarity $\text{Tr}(XY)/K$.

P-obj: The primal objective value $b_0 + c \cdot x$.

D-obj: The dual objective value $\text{Tr}(CY) + b_0 + b \cdot y$.

gap: The value of `dualityGap`.

P-err: The primal error $\max_{i,j}\{|p_i|, |P_{ij}|\}$.

D-err: The dual error $\max_i\{|d_i|\}$.

P-step: The primal step length $\alpha_{\mathcal{P}}$ described in [1].

D-step: The dual step length $\alpha_{\mathcal{D}}$ described in [1].

beta: The corrector centering parameter β_c described in [1].

dim/stabilized: N/N' , where N is the dimension of the vector y , and N' is the dimension of the matrix Q' obtained after stabilizing the Schur complement matrix, described in [1]. A large N' will generally cause a big slowdown, so it is best avoided. N' can be reduced by decreasing `choleskyStabilizeThreshold`, though this sometimes requires increasing `precision` to avoid numerical instabilities.

If an optimal solution exists, the primal and dual error will decrease until the problem becomes primal and dual feasible. Then the primal and dual objective functions start to converge, and the complementarity μ decreases until the duality gap becomes smaller than `dualityGapThreshold`.

The terminal output ends with the final values of the primal/dual objectives, primal/dual errors and duality gap, together with the time since the last saved checkpoint and the total solver runtime.

5.2 Termination

The possible termination reasons for SDPB are as follows

`found primal-dual optimal solution`

Found a solution for x, X, y, Y that is simultaneously primal feasible, dual feasible, and optimal.

`found primal feasible solution`

Found a solution for x, X that is primal feasible. SDPB will only terminate with this result if the option `--findPrimalFeasible` is specified.

`found dual feasible solution`

Found a solution for y, Y that is dual feasible. SDPB will only terminate with this result if the option `--findDualFeasible` is specified.

primal feasible jump detected

A Newton step with primal step length α_P just occurred, without resulting in a primal feasible solution. (Usually this means one should increase `precision` and resume from the latest checkpoint.)

dual feasible jump detected

A Newton step with dual step length α_D just occurred, without resulting in a dual feasible solution. (Usually this means one should increase `precision` and resume from the latest checkpoint.)

maxIterations exceeded

SDPB has run for more iterations than specified by the option `--maxIterations`.

maxRuntime exceeded

SDPB has run for longer than specified by the option `--maxRuntime`.

maxComplementarity exceeded

$\mu = \text{Tr}(XY)/\text{dim}(X)$ exceeded the value specified by `--maxComplementarity`. This might indicate that the problem is unbounded and no optimal solution will be found.

When using SDPB to determine primal or dual feasibility, one can specify the options `--findPrimalFeasible` or `--findDualFeasible`. This will cause the solver to terminate immediately once the primal or dual errors are sufficiently small. This often occurs immediately after the primal or dual step lengths become equal to 1. A step length of 1 means that the solver has found a Newton step that exactly solves the primal or dual constraints, while preserving positive-semidefiniteness of X, Y . Sometimes a step length of 1 does not result in sufficiently small primal/dual errors. This is indicative of numerical instabilities and usually means `precision` should be increased. The options `--detectPrimalFeasibleJump` and `--detectDualFeasibleJump` cause SDPB to terminate if a step length of 1 occurs without resulting in primal/dual feasibility. If desired, one can then restart the solver from the most recent checkpoint, with a higher value of `precision`.

5.3 Output File

Listing 6: Contents of the output file `test.out` corresponding to listing 4. Decimal expansions have been truncated for brevity. `Mathematica` uses `*` instead of the character `e` for scientific notation. Thus, the output format is not quite suitable for import into `Mathematica` without modification. This could be changed in future versions.

```
terminateReason = "found primal-dual optimal solution";
primalObjective = 1.840265763203;
dualObjective   = 1.840265763084;
dualityGap      = 3.221067914086e-11;
primalError     = 4.263251669979e-136;
dualError       = 1.421540011331e-133;
runtime         = 0.16122411191463470458984375;
y = {-1.840265763084};
x = {0.4523538794795, -0.803480855768, 2.460542537885, 0.361240154722, -0.094037214700};
```

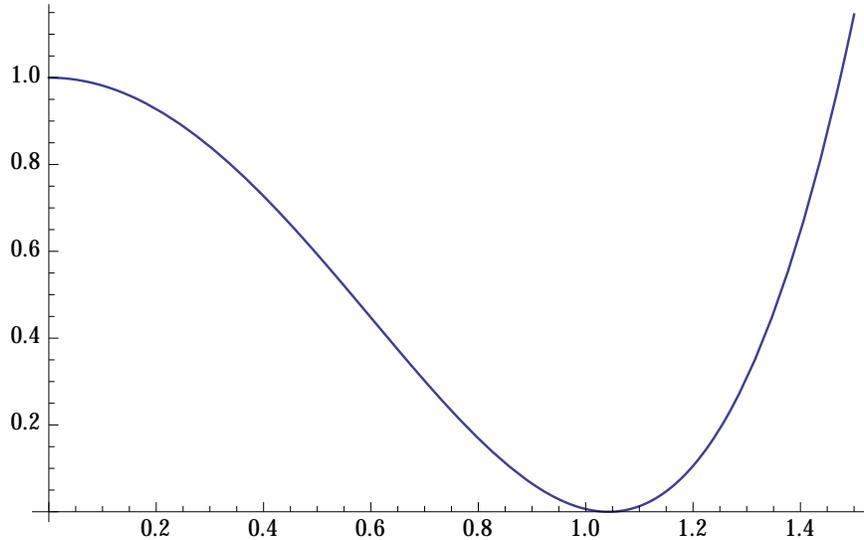


Figure 1: A plot of $1 + x^4 + y \left(\frac{x^4}{12} + x^2 \right)$ with $y = -1.840265763084$ equal to its optimal value. The zero near $x = 1$ shows that $-y$ cannot be further increased without violating the positivity constraint.

The output file `test.out` corresponding to listing 4 is shown in listing 6. It includes the reason for termination, the final primal/dual objective values, the final duality gap, the final primal/dual errors, the total runtime, and the vectors y and x .¹

The value of y gives the solution to our optimization problem. The function

$$1 + x^4 + (-1.840265763084) \left(\frac{x^4}{12} + x^2 \right) \quad (5.1)$$

is plotted in figure 1. The zero near $x = 1$ shows that y is optimal.

5.4 Checkpoints

Every `checkpointInterval`, SDPB saves a new checkpoint file with a `.ck` extension and backs up the old checkpoint file to a `.ck.bk` extension. SDPB also saves a checkpoint after termination, provided the option `--noFinalCheckpoint` is not specified.

A checkpoint file encodes the values of x, X, y, Y . If SDPB detects an existing checkpoint file on startup, it will use those values of x, X, y, Y as initial conditions in the solver. Thus, SDPB can be stopped and started at will without losing progress.

A typical workflow for long-running computations on shared machines is to specify a moderate `checkpointInterval` (e.g. one hour) and a somewhat larger `maxRuntime` (e.g. 12 hours). SDPB will terminate after 12 hours and can then be restarted without losing

¹To include the matrices X, Y as well, uncomment the lines `// ofs << "Y = " << Y << "\n";` and `// ofs << "X = " << X << "\n";` in the source file `SDPSolverIO.cpp`, and recompile SDPB.

progress. If SDPB is killed prematurely, then at most 1 hour of progress will be lost. This pattern of restarting gives other users chances to run their processes. It can be sustained indefinitely, allowing extremely long computations.

6 Attribution

If you use SDPB in work that results in publication, please cite [1]. Depending on how SDPB is used, the following sources might also be relevant:

- The first use of semidefinite programming in the bootstrap [3].
- The generalization of semidefinite programming methods to arbitrary spacetime dimension [4].
- The generalization of semidefinite programming methods to arbitrary systems of correlation functions [5].

7 Acknowledgements

SDPB Makes extensive use of MPACK [9], the multiple precision linear algebra library written by Maho Nakata. Several source files from MPACK are included in the SDPB source tree (see the license at the top of those files). SDPB uses the Boost C++ libraries [10] and Lee Thomason’s `tinyxml2` library [11] for parsing. SDPB was partially based on the solvers SDPA and SDPA-GMP [6–8], which were essential sources of inspiration and examples.

Thanks to Filip Kos, David Poland, and Alessandro Vichi for collaboration in developing semidefinite programming methods for the conformal bootstrap and assistance testing SDPB. Thanks to Amir Ali Ahmadi, Hande Benson, Pablo Parrilo, and Robert Vanderbei for advice and discussions about semidefinite programming.

I am supported by DOE grant number DE-SC0009988 and a William D. Loughlin Membership at the Institute for Advanced Study.

References

- [1] David Simmons-Duffin, “A Semidefinite Program Solver for the Conformal Bootstrap,” [arXiv:1502.02033 \[hep-th\]](#).
- [2] V. S. Rychkov and A. Vichi, “Universal Constraints on Conformal Operator Dimensions,” *Phys. Rev. D* **80**, 045006 (2009) [arXiv:0905.2211 \[hep-th\]](#).
- [3] D. Poland, D. Simmons-Duffin and A. Vichi, “Carving Out the Space of 4D CFTs,” *JHEP* **1205**, 110 (2012) [arXiv:1109.5176 \[hep-th\]](#).

- [4] F. Kos, D. Poland and D. Simmons-Duffin, “Bootstrapping the $O(N)$ vector models,” JHEP **1406**, 091 (2014) [arXiv:1307.6856 \[hep-th\]](#).
- [5] F. Kos, D. Poland and D. Simmons-Duffin, “Bootstrapping Mixed Correlators in the 3D Ising Model,” JHEP **1411**, 109 (2014) [arXiv:1406.4858 \[hep-th\]](#).
- [6] 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).
- [7] M. Yamashita, K. Fujisawa, and M. Kojima, “Implementation and evaluation of SDPA 6.0 (SemiDefinite Programming Algorithm 6.0),” Optimization Methods and Software” 18 491-505 (2003).
- [8] M. Nakata, “A numerical evaluation of highly accurate multiple-precision arithmetic version of semidefinite programming solver: SDPA-GMP, -QD and -DD.,” 2010 IEEE International Symposium on Computer-Aided Control System Design (CACSD), 29-34 Sept 2010.
- [9] M. Nakata, “The MPACK (MBLAS/MLAPACK); a multiple precision arithmetic version of BLAS and LAPACK,” 2010 <http://mplapack.sourceforge.net/>.
- [10] C++ Standards Committee Library Working Group and other contributors, “BOOST C++ Libraries,” <http://www.boost.org>.
- [11] L. Thomason, TinyXML2, <http://www.grinninglizard.com/tinyxml2docs/index.html>