

DSDP5 User Guide — Software for Semidefinite Programming

Steven J. Benson
Mathematics and Computer Science Division
Argonne National Laboratory
Argonne, IL U.S.A.
<http://www.mcs.anl.gov/~benson>

Yinyu Ye
Department of Management Science and Engineering
Stanford University
Stanford, CA U.S.A.
<http://www.stanford.edu/~yyye>

Technical Report ANL/MCS-TM-277

September 21, 2005

Abstract

DSDP implements the dual-scaling algorithm for semidefinite programming. The source code of this interior-point solver, written entirely in ANSI C, is freely available. The solver can be used as a subroutine library, as a function within the MATLAB environment, or as an executable that reads and writes to files. Initiated in 1997, DSDP has developed into an efficient and robust general purpose solver for semidefinite programming. Although the solver is written with semidefinite programming in mind, it can also be used for linear programming and other constraint cones.

The features of DSDP include:

- a robust algorithm with a convergence proof and polynomially bounded complexity under mild assumptions on the data,
- primal and dual solutions,
- feasible solutions when they exist or approximate certificates of infeasibility,
- initial points that can be feasible or infeasible,
- relatively low memory requirements for an interior-point method,
- sparse and low-rank data structures,
- extensibility that allows applications to customize the solver and improve its performance,
- a subroutine library that enables it to be linked to larger applications,
- scalable performance for large problems on parallel architectures, and
- a well documented interface and examples of its use.

The package has been used in many applications and tested for efficiency, robustness, and ease of use. We welcome and encourage further use under the terms of the license included in the distribution.

Contents

1	Notation	1
2	Dual-Scaling Algorithm	2
3	Feasible Points, Infeasible Points, and Standard Form	4
4	Iteration Monitor	5
5	DSDP Subroutine Library	6
5.1	Creating the Solver	6
5.2	Semidefinite Cone	6
5.3	LP Cone	10
5.4	Applying the Solver	12
5.5	Convergence Criteria	12
5.6	Detecting Infeasibility	13
5.7	Solutions and Statistics	13
5.8	Improving Performance	14
5.9	Iteration Monitor	15
6	PDSDP	16
	Bibliography	19
	Index	19

1 Notation

The DSDP package implements a dual-scaling algorithm to find solutions (X_j, y_i, S_j) to linear and semidefinite optimization problems of the form

$$(P) \quad \inf \sum_{j=1}^p \langle C_j, X_j \rangle \quad \text{subject to} \quad \sum_{j=1}^p \langle A_{i,j}, X_j \rangle = b_i, \quad i = 1, \dots, m, \quad X_j \in K_j,$$

$$(D) \quad \sup \sum_{i=1}^m b_i y_i \quad \text{subject to} \quad \sum_{i=1}^m A_{i,j} y_i + S_j = C_j, \quad j = 1, \dots, p, \quad S_j \in K_j.$$

In this formulation, b_i and y_i are real scalars.

For semidefinite programming, the data $A_{i,j}$ and C_j are symmetric matrices of dimension n_j (\mathbb{S}^{n_j}), and the cone K_j is the set of symmetric positive semidefinite matrices of the same dimension. The inner product $\langle C, X \rangle := C \bullet X := \sum_{k,l} C_{k,l} X_{k,l}$, and the symbol $\succ (\succeq)$ means the matrix is positive (semi)definite. In linear programming, A_i and C are vectors of real scalars, K is the nonnegative orthant, and the inner product $\langle C, X \rangle$ is the usual vector inner product.

More generally, users specify $C_j, A_{i,j}$ from an inner-product space V_j that intersects a cone K_j . Using the notation summarized in Table 1, let the symbol \mathcal{A} denote the linear map $\mathcal{A} : V \rightarrow \mathbb{R}^m$ defined by $(\mathcal{A}X)_i = \langle A_i, X \rangle$; its adjoint $\mathcal{A}^* : \mathbb{R}^m \rightarrow V$ is defined by $\mathcal{A}^*y = \sum_{i=1}^m y_i A_i$. Equivalent expressions for (P) and (D) can be written

$$(P) \quad \inf \langle C, X \rangle \quad \text{subject to} \quad \mathcal{A}X = b, \quad X \in K,$$

$$(D) \quad \sup b^T y \quad \text{subject to} \quad \mathcal{A}^*y + S = C, \quad S \in K.$$

Formulation (P) will be referred to as the *primal* problem, and formulation (D) will be referred to as the *dual* problem. Variables that satisfy the linear equations are called feasible, whereas the others are called infeasible. The interior of the cone will be denoted by \hat{K} , and the interior feasible sets of (P) and (D) will be denoted by $\mathcal{F}^0(P)$ and $\mathcal{F}^0(D)$, respectively.

Table 1: Basic terms and notation for linear (LP), semidefinite (SDP), and conic programming.

Term	LP	SDP	Conic	Notation
Dimension	n	n	$\sum n_j$	n
Data Space ($\ni C, A_i$)	\mathbb{R}^n	\mathbb{S}^n	$V_1 \oplus \dots \oplus V_p$	V
Cone	$x, s \geq 0$	$X, S \succeq 0$	$X, S \in K_1 \oplus \dots \oplus K_p$	$X, S \in K$
Interior of Cone	$x, s > 0$	$X, S \succ 0$	$X, S \in \hat{K}_1 \oplus \dots \oplus \hat{K}_p$	$X, S \in \hat{K}$
Inner Product	$c^T x$	$C \bullet X$	$\sum \langle C_j, X_j \rangle$	$\langle C, X \rangle$
Norm	$\ x\ _2$	$\ X\ _F$	$(\sum \ X_j\ ^2)^{1/2}$	$\ X\ $
Product	$[x_1 s_1 \dots x_n s_n]^T$	XS	$X_1 S_1 \oplus \dots \oplus X_p S_p$	XS
Identity Element	$[1 \dots 1]^T$	I	$I_1 \oplus \dots \oplus I_p$	I
Inverse	$[1/s_1 \dots 1/s_n]^T$	S^{-1}	$S_1^{-1} \oplus \dots \oplus S_p^{-1}$	S^{-1}
Dual Barrier	$\sum \ln s_j$	$\ln \det S$	$\sum \ln \det S_j$	$\ln \det S$

2 Dual-Scaling Algorithm

This section summarizes the dual-scaling algorithm for solving (P) and (D). For simplicity, parts of this discussion assume that the cone is a single semidefinite block, but an extension of the algorithm to multiple blocks and other cones is relatively simple. This discussion also assumes that the A_i s are linearly independent, there exists $X \in \mathcal{F}^0(P)$, and a starting point $(y, S) \in \mathcal{F}^0(D)$ is known. The next section discusses how DSDP generalizes the algorithm to relax these assumptions.

It is well known that under these assumptions, both (P) and (D) have optimal solutions X^* and (y^*, S^*) , which are characterized by the equivalent conditions that the duality gap $\langle X^*, S^* \rangle$ is zero and the product X^*S^* is zero. Moreover, for every $\nu > 0$, there exists a unique primal-dual feasible solution (X_ν, y_ν, S_ν) satisfies the perturbed optimality equation $X_\nu S_\nu = \nu I$. The set of all solutions $\mathcal{C} \equiv \{(X_\nu, y_\nu, S_\nu) : \nu > 0\}$ is known as the central path, and \mathcal{C} serves as the basis for path-following algorithms that solve (P) and (D). These algorithms construct a sequence $\{(X, y, S)\} \subset \mathcal{F}^0(P) \times \mathcal{F}^0(D)$ in a neighborhood of the central path such that the duality gap $\langle X, S \rangle$ goes to zero. A scaled measure of the duality gap that proves useful in the presentation and analysis of path-following algorithms is $\mu(X, S) = \langle X, S \rangle / n$ for all $(X, S) \in K \times K$. Note that for all $(X, S) \in \hat{K} \times \hat{K}$, we have $\mu(X, S) > 0$ unless $XS = 0$. Moreover, $\mu(X_\nu, S_\nu) = \nu$ for all points (X_ν, y_ν, S_ν) on the central path.

The dual-scaling algorithm applies Newton's method to $\mathcal{A}X = b$, $\mathcal{A}^*y + S = C$, and $X = \nu S^{-1}$ to generate

$$\mathcal{A}(X + \Delta X) = b, \quad (1)$$

$$\mathcal{A}^*(\Delta y) + \Delta S = 0, \quad (2)$$

$$\nu S^{-1} \Delta S S^{-1} + \Delta X = \nu S^{-1} - X. \quad (3)$$

Equations (1)-(3) will be referred to as the Newton equations; their Schur complement is

$$\nu \begin{pmatrix} \langle A_1, S^{-1} A_1 S^{-1} \rangle & \cdots & \langle A_1, S^{-1} A_m S^{-1} \rangle \\ \vdots & \ddots & \vdots \\ \langle A_m, S^{-1} A_1 S^{-1} \rangle & \cdots & \langle A_m, S^{-1} A_m S^{-1} \rangle \end{pmatrix} \Delta y = b - \nu \mathcal{A} S^{-1}. \quad (4)$$

The left-hand side of this linear system is positive definite when $S \in \hat{K}$. In this manuscript, it will sometimes be referred to as M . DSDP computes $\Delta'y := M^{-1}b$ and $\Delta''y := M^{-1}\mathcal{A}S^{-1}$. For any ν ,

$$\Delta_\nu y := \frac{1}{\nu} \Delta'y - \Delta''y$$

solves (4). We use the subscript to emphasize that ν can be chosen after computing $\Delta'y$ and $\Delta''y$ and that the value chosen for the primal step may be different from the value chosen for the dual step.

Using $\Delta_\nu y$ and (3), we get

$$X(\nu) := \nu (S^{-1} + S^{-1}(\mathcal{A}^* \Delta_\nu y) S^{-1}), \quad (5)$$

which satisfies $\mathcal{A}X(\nu) = b$. Because $X(\nu) \in \hat{K}$ if and only if

$$C - \mathcal{A}^*(y - \Delta_\nu y) \in \hat{K}, \quad (6)$$

DSDP applies a Cholesky factorization on (6) to test the condition. If $X(\nu) \in \hat{K}$, a new upper bound

$$\bar{z} := \langle C, X(\nu) \rangle = b^T y + \langle X(\nu), S \rangle = b^T y + \nu (\Delta_\nu y^T \mathcal{A} S^{-1} + n) \quad (7)$$

can be obtained without explicitly computing $X(\nu)$. The dual-scaling algorithm does not require $X(\nu)$ to compute the step direction defined by (4), so DSDP does not compute it unless specifically requested. This feature characterizes the algorithm and its performance.

Either (y, S) or X reduces the the dual potential function

$$\psi(y) := \rho \log(\bar{z} - b^T y) - \ln \det S \quad (8)$$

enough at each iteration to achieve linear convergence.

- 1: Setup data structures and factor A_i .
- 2: Choose y such that $S \leftarrow C - \mathcal{A}^*y \in \hat{K}$.
- 3: Choose an upper bound \bar{z} and a barrier parameter ν .
- 4: **for** $k \leftarrow 0, \dots, k_{max}$ **do**
- 5: Monitor solution and check for convergence.
- 6: Compute M and $\mathcal{A}S^{-1}$.
- 7: Solve $M\Delta'y = b$, $M\Delta''y = \mathcal{A}S^{-1}$.
- 8: **if** $C - \mathcal{A}^*(y - \Delta_\nu y) \in \hat{K}$ **then**
- 9: $\bar{z} \leftarrow b^T y + \nu (\Delta_\nu y^T \mathcal{A}S^{-1} + n)$.
- 10: $\bar{y} \leftarrow y$, $\overline{\Delta y} \leftarrow \Delta_\nu y$, $\bar{\mu} \leftarrow \nu$.
- 11: **end if**
- 12: Select ν .
- 13: Find α_d to reduce ψ , and set $y \leftarrow y + \alpha_d \Delta_\nu y$, $S \leftarrow C - \mathcal{A}^*y$.
- 14: **for** $kk = 1, \dots, kk_{max}$ **do**
- 15: Compute $\mathcal{A}S^{-1}$.
- 16: Solve $M\Delta^c y = \mathcal{A}S^{-1}$.
- 17: Select ν .
- 18: Find α_c to reduce ϕ_ν , and set $y \leftarrow y + \alpha_c \Delta_\nu^c y$, $S \leftarrow C - \mathcal{A}^*y$.
- 19: **end for**
- 20: **end for**
- 21: Optional: Compute X using \bar{y} , $\overline{\Delta y}$, $\bar{\mu}$.

3 Feasible Points, Infeasible Points, and Standard Form

The convergence of the algorithm assumes that both (P) and (D) have an interior feasible region and the current solutions are elements of the interior. To satisfy these assumptions, DSDP bounds the variables y such that $l \leq y \leq u$ where $l, u \in \mathbb{R}^m$. By default, $l_i = -10^7$ and $u_i = 10^7$ for each i from 1 through m . Furthermore, DSDP bounds the trace of X by a penalty parameter Γ whose default value is $\Gamma = 10^{10}$. Including these bounds and their associated Lagrange variables $x^l \in \mathbb{R}^m$, $x^u \in \mathbb{R}^m$, and r , DSDP solves following pair of problems:

$$\begin{aligned}
 (PP) \quad & \text{minimize} && \langle C, X \rangle & + && u^T x^u & - & l^T x^l \\
 & \text{subject to} && \mathcal{A}X & + && x^u & - & x^l & = & b, \\
 & && \langle I, X \rangle & && & & & \leq & \Gamma, \\
 & && X \in K, & && x^u \geq 0, & & x^l \geq 0.
 \end{aligned}$$

$$\begin{aligned}
 (DD) \quad & \text{maximize} && b^T y - \Gamma r \\
 & \text{subject to} && C - \mathcal{A}^* y + Ir = S \in K, \\
 & && l \leq y \leq u, & && r \geq 0.
 \end{aligned}$$

The reformulations (PP) and (DD) are bounded and feasible, so the optimal objective values to this pair of problems are equal. Furthermore, (PP) and (DD) can be expressed in the form of (P) and (D).

Unless the user provides a feasible point y , DSDP uses the y values provided by the application (usually all zeros) and increases r until $C - \mathcal{A}^* y + Ir \in \hat{K}$. Large values of r improve robustness, but smaller values often improve performance. In addition to bounding X , the parameter Γ penalizes infeasibility in (D) and forces r toward zero. The nonnegative variable r increases the dimension m by one and adds an inequality to the original problem. The M matrix treats r separately by storing the corresponding row/column as a separate vector and applying the Sherman-Morrison-Woodbury formula. Unlike other inequalities, DSDP allows r to reach the boundary of the cone. Once $r = 0$, it is fixed and effectively removed from the problem.

The bounds on y add $2m$ inequality constraints to the original problem; and, with a single exception, DSDP treats them the same as the constraints on the original model. The lone difference between these bounds and the other constraints is that DSDP explicitly computes the corresponding Lagrangian variables x^l and x^u at each iteration to quantify the infeasibility in (P). The bounds l and u penalize infeasibility in (P), force x^l and x^u toward zero, and prevent numerical difficulties created by variables with large magnitude.

The solution to (PP) and (DD) is a solution to (P) and (D) when the optimal objective values of (P) and (D) exist and are equal, and the bounds are sufficiently large. DSDP identifies unboundedness or infeasibility in (P) and (D) through examination of the solutions to (PP) and (DD). Given parameters ϵ_P and ϵ_D ,

- if $r \leq \epsilon_r$, $\|\mathcal{A}X - b\|_\infty / \langle I, X \rangle > \epsilon_P$, and $b^T y > 0$, it characterizes (D) as unbounded and (P) as infeasible;
- if $r > \epsilon_r$ and $\|\mathcal{A}X - b\|_\infty / \langle I, X \rangle \leq \epsilon_P$, it characterizes (D) as infeasible and (P) as unbounded.

Normalizing unbounded solutions will provide an approximate certificate of infeasibility. Larger bounds may improve the quality of the certificate of infeasibility and permit additional feasible solutions, but they may also create numerical difficulties in the solver.

4 Iteration Monitor

The progress of the DSDP solver can be monitored by using standard output printed to the screen. The data below shows an example of this output.

Iter	PP Objective	DD Objective	PInfeas	DInfeas	Nu	StepLength	Pnrm
0	1.00000000e+02	-1.13743137e+05	2.2e+00	3.8e+02	1.1e+05	0.00 0.00	0.00
1	1.36503342e+06	-6.65779055e+04	5.1e+00	2.2e+02	1.1e+04	1.00 0.33	4.06
2	1.36631922e+05	-6.21604409e+03	5.4e+00	1.9e+01	4.5e+02	1.00 1.00	7.85
3	5.45799174e+03	-3.18292092e+03	1.5e-03	9.1e+00	7.5e+01	1.00 1.00	17.63
4	1.02930559e+03	-5.39166166e+02	1.1e-05	5.3e-01	2.7e+01	1.00 1.00	7.58
5	4.30074471e+02	-3.02460061e+01	3.3e-09	0.0e+00	5.6e+00	1.00 1.00	11.36
...							
11	8.99999824e+00	8.99999617e+00	1.1e-16	0.0e+00	1.7e-08	1.00 1.00	7.03
12	8.99999668e+00	8.99999629e+00	2.9e-19	0.0e+00	3.4e-09	1.00 1.00	14.19

The program will print a variety of statistics for each problem to the screen.

Iter	the iteration number.
PP Objective	the upper bound \bar{z} and objective value in (PP).
DD Objective	the objective value in (DD).
PInfeas	the primal infeasibility in (P) is $\ x^u - x^l\ _\infty$.
DInfeas	the dual infeasibility in (D) is the variable r .
Nu	the barrier parameter ν .
StepLength	the multiple of the step-directions in (P) and (D).
Pnrm	the proximity to the central path: $\ \nabla\psi\ _{M^{-1}}$.

5 DSDP Subroutine Library

DSDP can also be used within a C application through a set of subroutines. There are several examples of applications that use the DSDP application program interface. Within the `DSDPROOT/examples/` directory, the file `dsdp.c` is a mex function that reads data from the MATLAB environment, passes the data to the DSDP solver, and returns the solution. The file `readsdp.c` reads data from a file for data in SDPA format, passes the data to the solver, and prints the solution. The files `maxcut.c` and `theta.c` read a graph, formulate a semidefinite relaxation to a combinatorial problem, and pass the data to a DSDP solver. The subroutines used in these examples are described in this chapter.

Each of these applications includes the header file `DSDPROOT/include/dsdp5.h` and links to the library `DSDPROOT/lib/libdsdp.a`. All DSDP subroutines also return an `int` that represents an error code. A return value of zero indicates success, whereas a nonzero return value indicates that an error has occurred. The documentation of DSDP subroutines in this chapter will not show the return integer, but we highly recommend that applications check for errors after each subroutine.

5.1 Creating the Solver

To use DSDP through subroutine, the solver object must first be created with the command

```
DSDPCreate(int m, DSDP *newsolver);
```

The first argument in this subroutine is the number of variables in the problem. The second argument should be the address of a DSDP variable. This subroutine will allocate memory for the solver and set the address of the DSDP variable to a new solver object. A difference in typeset distinguishes the DSDP package from the DSDP solver object.

The objective function associated with these variables should be specified using the subroutine

```
DSDPSetDObjective(DSDP dsdp, int i, double bi);
```

The first argument is the solver, and the second and third arguments specify a variable number and the objective value b_i associated with it. The variables are numbered 1 through m , where m is the number of variables specified in `DSDPCreate`. The objective associated with each variable must be specified individually. The default value is zero.

The next step is to provide the conic structure and data in the problem. These subroutines will be described in the next sections.

5.2 Semidefinite Cone

To specify an application with a cone of semidefinite constraints, the subroutine

```
DSDPCreateSDPCone(DSDP dsdp, int nblocks, SDPCone *newsdpcone);
```

can be used to create a new object that describes a semidefinite cone with 1 or more blocks. The first argument is an existing semidefinite solver, the second argument is the number of blocks in this cone, and the final argument is an address of a `SDPCone` variable. This subroutine will allocate structures needed to specify the constraints and point the variable to this structure. Multiple cones can be created for the same solver, but it is usually more efficient to group all blocks into the same conic structure.

All subroutines that pass data to the semidefinite cone use this `SDPCone` variable in the first argument. The second argument often refers to a specific block. The blocks will be labeled from 0 to `nblocks-1`. The subroutine `SDPConeSetBlockSize(SDPCone sdpcone, int blockj, int n)` can be used to specify

the dimension of each block and the subroutine `SDPConeSetSparsity(SDPCone sdpcone, int blockj, int nnzmat)` can be used to specify the number of nonzero matrices $A_{i,j}$ in each block. These subroutines are optional, but using them can improve error checking on the data matrices and perform a more efficient allocation of memory.

The data matrices can be specified by any of the following commands. The choice of data structures is up to the user, and the performance of the problem depends upon this choice of data structures. In each of these subroutines, the first four arguments are a pointer to the `SDPCone` object, the block number, and the number of variable associated with it, and the number of rows and columns in the matrix. The blocks must be numbered consecutively, beginning with the number 0. The y variables are numbered consecutively from 1 to m . The objective matrices in (P) are specified as constraint number 0. The data that is passed to the `SDPCone` object will be used in the solver, but not modified. The user is responsible for freeing the arrays of data it passes to DSDP after solving the problem.

The square symmetric data matrices $A_{i,j}$ and C_j can be represented with a single array of numbers. DSDP supports the **symmetric packed** storage format. In symmetric packaged storage format, the elements of a matrix with n rows and columns are ordered as follows:

$$[a_{1,1} \ a_{2,1} \ a_{2,2} \ a_{3,1} \ a_{3,2} \ a_{3,3} \ \dots \ a_{n,n}]. \quad (9)$$

In this array $a_{k,l}$ is the element in row k and column l of the matrix. The length of this array is $n(n+1)/2$, which is the number of distinct elements on or below the diagonal. Several routines described below have an array of this length in its list of arguments. In this storage format, the element in row i and column j , where $i \geq j$, is in element $i(i-1)/2 + j - 1$ of the array.

This array can be passed to the solver using the subroutine

```
SDPConeSetADenseVecMat(SDPCone sdpcone,int blockj, int vari,
                       int n, double alpha, double val[], int nnz);
```

The first argument point to a semidefinite cone object, the second argument specifies the block number j , and the third argument specifies the variable i associated with it. Variables $1, \dots, m$ correspond to matrices $A_{1,j}, \dots, A_{m,j}$ whereas variable 0 corresponds to C_j . The fourth argument is the dimension (number of rows or columns) of the matrix, n . The sixth argument is the array, and seventh argument is the length of the array. The data array will be multiplied by the scalar in the fifth argument. The application is responsible for allocating this array of data and eventually freeing it. DSDP will directly access this array in the course of the solving the problem, so it should not be freed until the solver is finished.

A matrix can be passed to the solver in sparse format using the subroutine

```
SDPConeSetASparseVecMat(SDPCone sdpcone,int blockj, int vari, int n,
                        double alpha, int یشift
                        const int ind[], const double val[], int nnz);
```

In this subroutine, the first five arguments are the same as in the subroutine for dense matrices. The seventh and eighth arguments are an array an integers and an array of double precision variables. The final argument states the length of these two arrays, which should equal the number of nonzeros in the lower triangular part of the matrix. The array of integers specifies which elements of the array (9) are included in the array of doubles. For example, the matrix

$$A_{i,j} = \begin{bmatrix} 3 & 2 & 0 \\ 2 & 0 & 6 \\ 0 & 6 & 0 \end{bmatrix} \quad (10)$$

could be inserted into the cone using one of several ways. If the first element in the `val` array is $a_{1,1}$, the first element in the `ind` array should be 0. If the second element in the `val` array is $a_{3,2}$, then the second element in `ind` array should be 4. When the ordering of elements begins with 0, as just shown, the fifth argument `یشift` in the subroutine should be set to 0. In general, the argument `یشift` specifies the index assigned to $a_{1,1}$. Although the relative ordering of the elements will not change, the indices assigned to them will range from `یشift` to `یشift + n(n+1)/2 - 1`. Many applications, for instance, prefer to index the

array from 1 to $n(n+1)/2$, setting the `index` argument to 1. The matrix (10) can be set in the block j and variable i of the semidefinite cone using one of the routines

```
SDPConeSetASparseVecMat(sdpcone, j, i, 3, 1.0, 0, ind1, val1, 3);
SDPConeSetASparseVecMat(sdpcone, j, i, 3, 1.0, 1, ind2, val2, 3);
SDPConeSetASparseVecMat(sdpcone, j, i, 3, 1.0, 3, ind3, val3, 4);
SDPConeSetASparseVecMat(sdpcone, j, i, 3, 0.5, 0, ind4, val4, 3);
```

where

$$\begin{aligned} \text{ind1} &= \begin{bmatrix} 0 & 1 & 4 \end{bmatrix} & \text{val1} &= \begin{bmatrix} 3 & 2 & 6 \end{bmatrix} \\ \text{ind2} &= \begin{bmatrix} 1 & 2 & 5 \end{bmatrix} & \text{val2} &= \begin{bmatrix} 3 & 2 & 6 \end{bmatrix} \\ \text{ind3} &= \begin{bmatrix} 7 & 3 & 5 & 4 \end{bmatrix} & \text{val3} &= \begin{bmatrix} 6 & 3 & 0 & 2 \end{bmatrix}. \\ \text{ind4} &= \begin{bmatrix} 0 & 1 & 4 \end{bmatrix} & \text{val4} &= \begin{bmatrix} 6 & 4 & 12 \end{bmatrix} \end{aligned}$$

As these examples suggest, there are many other ways to represent the sparse matrix. The nonzeros in the matrix do not have to be ordered, but ordering them may improve the efficiency of the solver. DSDP assumes that all matrices $A_{i,j}$ and C_j that are not explicitly defined and passed to the SDPCone structure will equal the zero matrix. Furthermore, there exist routines `SDPConeAddASparseVecMat` and `SDPConeAddADenseVecMat` that can be used to write a constraint matrix as a sum of multiple matrices. The arguments to these function match those of the corresponding `SDPConeSetASparseVecMat` routines.

To check whether the matrix passed into the cone matches the one intended, the subroutine `SDPConeViewDataMatrix(SDPCone sdpcone, int blockj, int vari)`

can be used to print out the the matrix to the screen. The output prints the row and column numbers, indexed from 0 to $n-1$, of each nonzero element in the matrix. The subroutine `SDPConeView(SDPCone sdpcone, int blockj)` can be used to view all of the matrices in a block.

After the DSDP solver has been applied to the data and the solution matrix X_j have been computed (See `DSDPComputeX`), the matrix can be accessed using the command

```
SDPConeGetXArray(SDPCone sdpcone, int blockj, double *xmat[], int *nn);
```

The third argument is the address of a pointer that will be set to the array containing the solution. The integer whose address is passed in the fourth argument will be set to the length of this array, $n(n+1)/2$, for the packed symmetric storage format. Since the X solutions are usually fully dense, no sparse representation is provided. These arrays were allocated by the SDPCone object during `DSDPSetup` and the memory will be freed by the DSDP solver object when it is destroyed. The array used to store X_j could be overwritten by other operations on the SDPCone object. The command,

```
SDPConeComputeX(SDPCone sdpcone, int blockj, int n, double xmat[], int nn);
```

recomputes the matrix X_j and places it into the array specified in the fourth argument. The length of this array is the fifth argument and the dimension of the block in the third argument. The vectors y and Δy needed to compute the matrices X_j are stored internally in SDPCone object. The subroutine

```
SDPConeViewX(SDPCone sdpcone, int blockj, int n, double xmat[], int nn);
```

can be used to print this matrix to standard output.

The dimension of each block can be found using the routine

```
SDPConeGetBlockSize(SDPCone sdpcone, int blockj, int *n);
```

where the second arguments is the block number and the third arguments are the address of a variable.

The inner product of the matrix X_j with the C_j , $A_{i,j}$, and I_j can be computed using the routine `int SDPConeAddADotX(SDPCone sdpcone, int blockj, double alpha, double xmat[], int nn, double adotx[], int mp2);`

The second argument specifies which block to use, and the third argument is a scalar that will be multiplied by the inner products. The fourth argument is the array containing X_j , and the fifth argument is th length

of the array. The sixth argument is an array of length $m + 2$ and the seventh argument should equal $m + 2$ where m is the number of variables in y . This routine will add α times $\langle C_j, X_j \rangle$ to the initial element of the array, α times $\langle A_{i,j}, X_j \rangle$ to element i of the array, and α times $\langle I_j, X_j \rangle$ to last element of the array.

The matrix S in (D) can be computed and copied into an array using the command

```
SDPConeComputeS(SDPCone sdpcone, int blockj, double c, double y[], int m, double r,
                int n, double smat[], int nn);
```

The second argument specifies which block to use, the fourth argument is an array containing the variables y . The second argument is the multiple of the matrix C to be added, and the sixth argument is the multiple of the identity matrix to be added. The sixth argument is the dimension of the block, and the seventh argument is an array whose length is given in the eighth argument.

Special support for combinatorial applications also exists. In particular, the subroutines

```
SDPConeComputeXV(SDPCone sdpcone, int blockj, int *dpsdefinite);
SDPConeXVMultiply(SDPCone sdpcone, int blockj, double v1[], double v2[], int n);
SDPConeAddXVAV(SDPCone sdpcone, int blockj, double v[], int n, double vav[], int mp2);
```

support the used of randomized algorithms for combinatorial optimization. The first routine computes a matrix V_j such that $X = V_j V_j^T$. The second routine computes the matrix-vector product $w = V_j v$ where w and v are vectors of dimension equal to the dimension of the block. The third routine computes the vector-matrix-vector product $v^T A_{i,j} v$ for $C, A_{1,j}, \dots, A_{m,j}$, and the identity matrix. The length of the array in the fifth argument is $m + 2$ and should be set in the final argument. In these applications, use of the routine `DSDPConeComputeX` may not be necessary if the full matrix is not required.

The memory required for the X_j matrix can be significant for large problems. If the application has an array of double precision variables of length $n(n + 1)/2$ available for use by the solver, the subroutine

```
SDPConeSetXArray(SDPCone sdpcone, int blockj, int n, double xmat[], int nn);
```

can be used to pass it to the cone. The second argument specifies the block number whose solution will be placed into the array `xmat`. The third argument is the dimension of the block. The dimension specified in the fifth argument `nn` refers to the length of the array. The DSDP solver will use this array as a buffer for its computations and store the solution X in this array at its termination. The application is responsible for freeing this array after the solution has been found.

DSDP also supports the **symmetric full** storage format. In symmetric full storage format, an $n \times n$ matrix is stored in an array of n^2 elements in row major order. That is, the elements of a matrix with n rows and columns are ordered

$$[a_{1,1} \ 0 \ \dots \ 0 \ a_{2,1} \ a_{2,2} \ 0 \ \dots \ 0 \ \dots \ a_{n,1} \ \dots, \ a_{n,n}]. \quad (11)$$

The length of this array is $n \times n$. Early versions of DSDP5 used the packed format exclusively, but this format has been added because its more convenient for some applications. The routines

```
SDPConeSetStorageFormat(SDPCone sdpcone, int blockj, char UPLQ);
SDPConeGetStorageFormat(SDPCone sdpcone, int blockj, char *UPLQ);
```

set and get the storage format for a block. The second argument specifies the block number and the third argument should be 'P' for packed storage format and 'U' for full storage format. The default value is 'P'. These storage formats correspond the LAPACK storage formats for the upper half matrices in column major ordering. All of the above commands apply to the symmetric full storage format. One difference in their use, however, is that the size of the arrays in the arguments should be $n \times n$ instead of $n \times (n + 1)/2$. Using the symmetric full storage format, if the first element in the `val` array is $a_{1,1}$, the first element in the `ind` array should be 0. If the second element in the `val` array is $a_{3,2}$, then the second element in `ind` array should be 8. The matrix (10) can be set in the block j and variable i of the semidefinite cone using one of the routines

```
SDPConeSetASparseVecMat(sdpcone, j, i, 3, 1.0, 0, ind1, val1, 3);
```

where

$$\text{ind1} = [0 \ 3 \ 7] \quad \text{val1} = [3 \ 2 \ 6] .$$

5.3 LP Cone

To specify an application with a cone of linear scalar inequalities, the subroutine

```
DSDPCreateLPCone( DSDP dsdp, LPCone *newlpcone);
```

can be used to create a new object that describes a cone with 1 or more linear scalar inequalities. The first argument is an existing DSDP solver and the second argument is the address of an LPCone variable. This subroutine will allocate structures needed to specify the constraints and point the variable to this structure. Multiple cones for these inequalities can be created for the same DSDP solver, but it is usually more efficient to group all inequalities of this type into the same structure. All subroutines that pass data to the LP cone use this LPCone variable in the first argument.

A list of n linear inequalities in (D) are passed to the object in sparse column format. The vector $c \in \mathbb{R}^n$ should be considered an additional column of the data. (In the formulation (P), the data A and c is represented in sparse row format.)

Pass the data to the LPCone using the subroutine:

```
LPConeSetData(LPCone lpcone, int n,
              const int nnzin[], const int row[], const double aval[]);
```

In this case, the integer array `nnzin` has length $m + 2$, begins with 0, and `nnzin[i + 1] - nnzin[i]` equals the number of nonzeros in column i ($i = 0, \dots, m$) (or row i of A). The length of the second and third array equals the number of nonzeros in A and c . The arrays contain the nonzeros and the associated row numbers of each element (or column numbers in A). The first column contains the elements of c , the second column contains the elements corresponding to y_1 , and the last column contains elements corresponding to y_m .

For example, consider the following problem in the form of (D):

$$\begin{array}{rll} \text{Maximize} & y_1 & + \quad y_2 \\ \text{Subject to} & 4y_1 & + \quad 2y_2 \leq 6 \\ & 3y_1 & + \quad 7y_2 \leq 10 \\ & & -y_2 \leq 12 \end{array}$$

In this example, there three inequalities, so the dimension of the x vector would be 3 and $n = 3$. The input arrays would be as follows:

$$\begin{array}{l} \text{nnzin} = [0 \ 3 \ 5 \ 7] \\ \text{row} = [0 \ 1 \ 2 \ 0 \ 1 \ 0 \ 1 \ 2] \\ \text{aval} = [6.0 \ 10.0 \ 12.0 \ 4.0 \ 3.0 \ 2.0 \ 7.0 \ -1.0] \end{array}$$

An example of the use of this subroutine can be seen in the `DSDPROOT/examples/readsdpa.c`.

If its more convenient to specify the vector c in the last column, consider using the subroutine:

```
LPConeSetData2(LPCone lpcone, int n, const int ik[], const int cols[], const double vals[]);
```

This input is also sparse column input, but the c column comes last In this form the input arrays would be as follows:

$$\begin{array}{l} \text{nnzin} = [0 \ 2 \ 5 \ 7] \\ \text{row} = [0 \ 1 \ 0 \ 1 \ 2 \ 0 \ 1 \ 2] \\ \text{aval} = [4.0 \ 3.0 \ 2.0 \ 7.0 \ -1.0 \ 6.0 \ 10.0 \ 12.0] \end{array}$$

This subroutine is used in the DSDP Matlab mex function, which can be used as an example.

The subroutines

```
LPConeView(LPCone lpcone);
LPConeView2(LPCone lpcone);
```

can be used to view the data that has been set and verify the correctness of the data. Multiple `LPCone` structures can be used, but for efficiency purposes, it is often better to include all linear inequalities in a single cone.

The variables s in (D) and x in (P) can be found using the subroutines

```
LPConeGetXArray(LPCone lpcone, double *xout[], int *n);
LPConeGetSArray(LPCone lpcone, double *sout[], int *n);
```

In these subroutines, the second argument sets a pointer to an array of doubles containing the variables and the integer in the third argument will be set to the length of this array. These arrays were allocated by the `LPCone` object and the memory will be freed with the `DSDP` solver object is destroyed. Alternatively, the application can give the `LPCone` an array in which to put the solution x of (P). This array should be passed to the cone using the following subroutine

```
LPConeSetXVec(LPCone lpcone, double xout[], int n);
```

At completion of the `DSDP` solver, the solution x will be copied into this array `xout`, which must have length n . The slack variables s may be scaled. To get the unscaled vector, pass an array of appropriate length into the object using `LPConeGetSArray2(LPCone lpcone, double s[], int n)`. This subroutine will copy the slack variables into the array.

A special type of semidefinite and LP cone contains only simple bounds on the variables y . Corresponding to lower and upper bounds on the y variables are surplus and slack variables in (P) with a cost. The subroutine `DSDPCreateBCone(DSDP dsdp, BCone *bcone);`

will create this cone from an existing `DSDP` solver and point the `BCone` variable to the object. The bounds on the variables can be set using the subroutines:

```
BConeSetLowerBound(BCone bccone, int vari, double yi);
BConeSetUpperBound(BCone bccone, int vari, double yi);
```

The first argument is the conic object, the second argument identifies a variable from 1 to m , and the third argument is the bound. Here m is the total number of variables in the vector y . For applications using the formulation (P), the subroutines

```
BConeSetPSurplusVariable(BCone bccone, int vari);
BConeSetPSlackVariable(BCone bccone, int vari);
```

may be more convenient way to represent an inequality in (P). These commands are equivalent to setting a lower or upper bound on a variable y_i to zero. To improve the memory allocation process the application may want to use the subroutines `BConeAllocateBounds(BCone bccone, int nbounds);` to tell the object how many bounds will be specified. This subroutine is optional, but it may improve the efficiency of the memory allocation. The subroutine `BConeGetXArray(BCone bccone, double xout[], int n)` will set an array in the cone where the cone will copy the variables x . To view the bounds, the application may use the subroutine `BConeView(BCone bccone);` To retrieve the dual variables x corresponding to these constraints, use the subroutine `BConeCopyX(BCone bccone, double xl[], double xu[], int m);` The second and third arguments of this routine are an array of length m , the fourth argument. This routine will set the values of this array to the value of the corresponding variable. When no bound is present, the variable will equal zero.

In some applications it may be useful to fix a variable to a number. Instead of modeling this constraint as a pair of linear inequalities, fixed variables can be passed directly to the solver using the subroutine

```
DSDPSetFixedVariables(DSDP dsdp, double vars[], double vals[], double x[], int n);
```

In this subroutine, the array of variables in the second argument is set to the values in the array of the third argument. The fourth argument is an optional array in which the solver will put the sensitivities to these fixed variables. The final argument is the length of the arrays. Note, the values in the second argument are

integer numbers from 1 to m represented in double precision. Again, the integers should be one of $1, \dots, m$. Alternatively, a single variable can be set to a value using the subroutine `DSDPFixVariable(DSDP dsdp, int vari, double val)`.

5.4 Applying the Solver

After setting the data associated with the constraint cones, DSDP must allocate internal data structures and factor the data in the subroutine

```
DSDPSetup(DSDP dsdp);
```

This subroutine identifies factors the data, creates a Schur complement matrix with the appropriate sparsity, and allocates additional resources for the solver. This subroutine should be called after setting the data but before solving the problem. Furthermore, it should be called only once for each DSDP solver. On very large problems, insufficient memory on the computer may be encountered in this subroutine, so the error code should be checked. The subroutine

```
DSDPSetStandardMonitor(DSDP dsdp, int k);
```

will tell the solver to print the objective values and other information at each k iteration to standard output. The subroutine `DSDPLogInfoAllow(int,0)`; will print even more information if the first argument is positive.

The subroutine

```
DSDPSolve(DSDP dsdp);
```

attempts to solve the problem. This subroutine can be called more than once. For instance, the user may try solving the problem using different initial points.

After solving the problem, the subroutine

```
DSDPComputeX(DSDP dsdp);
```

can be called to compute the variables X in (P). These computations are not performed within the solver because these variables are not needed to compute the step direction.

Each solver created should be destroyed with the command

```
DSDPDestroy(DSDP dsdp);
```

This subroutine frees the work arrays and data structures allocated by the solver.

5.5 Convergence Criteria

Convergence of the DSDP solver may be defined using several options. The precision of the solution can be set by using the subroutine

```
DSDPSetGapTolerance(DSDP dsdp, double rgaptol);
```

The solver will terminate if there is a sufficiently feasible solution such that the difference between the objective values in (DD) and (PP), divided by the sum of their absolute values, is less than the prescribed number. A tolerance of 0.001 provides roughly three digits of accuracy, whereas a tolerance of $1.0\text{e-}5$ provides roughly five digits of accuracy. The subroutine

```
DSDPSetMaxIts(DSDP dsdp, int maxits);
```

specifies the maximum number of iterations. The subroutine `DSDPSetDualBound(DSDP, double)` specifies an upper bound on the objective value in (D). The algorithm will terminate when it finds a point when the variable r in (DD) is less than the prescribed tolerance and the objective value in (DD) is greater than this number.

5.6 Detecting Infeasibility

Infeasibility in either (P) or (D) can be determined using the subroutine

```
DSDPGetSolutionType(DSDP dsdp, DSDPSolutionType *pdffeasible);
```

This command sets the second argument to an enumerated type. There are four types for `DSDPSolutionType`:

- The type `DSDP_PDFEASIBLE` means that both (D) and (P) have feasible solutions and their objective values are bounded.
- The type `DSDP_UNBOUNDED` means that (D) is unbounded and (P) is infeasible. This type applies when the variable $r \leq \epsilon_r$ and $\|AX - b\|_\infty / \text{trace}(X) > \epsilon_P$. In this case, at least one variable y_i will be near its bound. The subroutine `DSDPSetYBounds` can adjust these bounds if the user thinks that they are not big enough to permit feasibility. Large bounds may create numerical difficulties in the solver, but they may also permit feasible solutions and improve the quality of the certificate of infeasibility. Normalizing the vector y will provide an approximate certificate of infeasibility for (P).
- The type `DSDP_INFEASIBLE` means that (D) is infeasible and (P) is unbounded. This type applies when the variable $r > \epsilon_r$ and $\|AX - b\|_\infty / \text{trace}(X) \leq \epsilon_P$. In this case, the trace of the variables X in (P) will be near the bound Γ . The subroutines `DSDPSetPenaltyParameter` can adjust Γ if the user thinks that it is too small. A larger parameter may create numerical difficulties in the solver, but it may also improve the quality of the certificate of infeasibility. Normalizing these variables so that to have a trace of 1.0 will provide an approximate certificate of infeasibility.
- The type `DSDP_PDUNKNOWN` means DSDP was unable to determine feasibility in either solution. This type applies when the initial point for (DD) was infeasible or if the bounds on y appear to be too small to permit a feasible solution.

The tolerance ϵ_r can be set using the subroutine `DSDPSetRTolerance(DSDP, double)`. The tolerance ϵ_P can be set using the subroutine `DSDPSetPTolerance(DSDP, double)`. The subroutines `DSDPGetRTolerance(DSDP, double*)` and `DSDPGetPTolerance(DSDP, double*)` can be used to get the current tolerances.

5.7 Solutions and Statistics

The objective values in (PP) and (DD) can be retrieved using the commands

```
DSDPGetPPObjective(DSDP dsdp, double *pobj);
DSDPGetDDObjective(DSDP dsdp, double *dobj);
```

The second argument in these routines is the address of a double precision variable.

The solution vector y can be viewed by using the command

```
DSDPGetY(DSDP dsdp, double y[], int m);
```

The user passes an array of size m where m is the number of variables in the problem. This subroutine will copy the solution into this array.

The success of DSDP can be interpreted with the command

```
DSDPStopReason(DSDP dsdp, DSDPTerminationReason *reason);
```

This command sets the second argument to an enumerated type. The various reasons for termination are listed below.

`DSDP_CONVERGED` The solutions to (PP) and (DD) satisfy the convergence criteria.

DSDP_MAX_IT	The solver applied the maximum number of iterations without finding solution.
DSDP_INFEASIBLE_START	The initial point in (DD) was infeasible.
DSDP_INDEFINITE_SCHUR	Numerical issues created an indefinite Schur matrix that prevented the further progress.
DSDP_SMALL_STEPS	Small step sizes prevented further progress.
DSDP_NUMERICAL_ERROR	Numerical issues prevented further progress.

The subroutines

```
DSDPGetBarrierParameter(DSDP dsdp, double *mu);
DSDPGetR(DSDP dsdp, double *r);
DSDPGetStepLengths(DSDP dsdp, double *pstep, double *dstep);
DSDPGetPnorm(DSDP dsdp, double *pnorm);
```

provide more information about the current solution. The subroutines obtain the barrier parameter, the variable r in (DD), the step lengths in (PP) and (DD), and a distance to the central path at the current iteration.

A history of information about the convergence of the solver can be obtained with the commands

```
DSDPGetGapHistory(DSDP dsdp, double gaphistory[], int history);
DSDPGetRHistory(DSDP dsdp, double rhistory[], int history);
```

retrieve the history of the duality gap and the variable r in (DD) for up to 100 iterations. The user passes an array of double precision variables and the length of this array. The subroutine

```
DSDPGetTraceX(DSDP dsdp, double *tracex);
```

gets the trace of the solution X in (P). Recall that the penalty parameter must exceed this quantity in order to return a feasible solution from an infeasible starting point. The subroutine

```
DSDPEventLogSummary(void)
```

will print out a summary of time spent in each cone and many of the primary computational subroutines.

5.8 Improving Performance

The performance of the DSDP may be *significantly* improved with the proper selection of bounds, parameters and initial point.

The application may specify an initial vector y to (D), a multiple of the identity matrix to make the initial matrix S positive definite, and an initial barrier parameter. The subroutine `DSDPSetY0(DSDP dsdp, int vari, double yi0)` can specify the initial value of the variable y_i . Like the objective function in (D), the variables are labeled from 1 to m . By default the initial values of y equal 0. Since convergence of the algorithm depends on the proximity of the point to the central path, initial points can be difficult to determine. Nonetheless, the subroutine

```
DSDPSetR0(DSDP dsdp, double r0)
```

will set the initial value of r in (DD). If $r0 < 0$, a default value will of $r0$ will be chosen. If S^0 is not positive definite, the solver will terminate will an appropriate termination flag. The default value is usually very large ($1e10$), but smaller values can *significantly* improve performance.

The subroutine `DSDPSetPotentialParameter(DSDP dsdp, double rho)` sets the potential parameter ρ . This parameter must be greater than 1. The default value is 4.0, but larger values such as 5 or 10 can significantly improve performance. Feasibility in (D) is enforced by means of a penalty parameter. By default it is set to $10e8$, but other values can affect the convergence of the algorithm. This parameter can

be set using `DSDPSetPenaltyParameter(DSDP dsdp, double Gamma)`, where `Gamma` is the large positive penalty parameter Γ . This parameter must exceed the trace of the solution X in order to return a feasible solution from an infeasible starting point. The subroutine `DSDPUsePenalty(DSDP dsdp, int yesorno)` is used to modify the algorithm. By default, the value is 0. A positive value means that the variable r in (DD) should be kept positive, treated like other inequalities, and penalized with the parameter `Gamma`. The subroutine `DSDPSetZBar(DSDP dsdp, double zbar)` sets an initial upper bound on the objective value at the solution. This value corresponds to the objective value of any feasible point of (PP). The subroutine `DSDPSetBarrierParameter(DSDP dsdp, double mu0)` sets the initial barrier parameter. The default heuristic is very robust, but performance can generally be improved by providing a smaller value.

DSDP applies the same Schur complement matrix for multiple linear systems. This feature often reduces the number of iterations and improves robustness. The cost of each iteration increases, especially when the dimension of the semidefinite blocks is of similar dimension or larger than the number of variables y . The subroutine

```
DSDPReuseMatrix(DSDP dsdp, int reuse);
```

can set a maximum on the number of times the Schur complement matrix is reused. The default value is 4, although the Matlab mex function and SDPA file reader set this parameter between 0 and 15 depending upon the size of the semidefinite blocks and the number of variables y . Applications whose semidefinite blocks are small relative the the number of variables y should probably use larger values while applications whose semidefinite blocks have size equal to or greater than the number of variables should probably set this parameter to zero.

The convergence of the dual-scaling algorithm assumes the existence of a strict interior in both (P) and (D). The use of a penalty parameter can add an interior to (D). An interior to (P) can be created by bounding the variables y . Default bounds of `-1e7` and `1e7` have been set, but applications may change these bounds using the subroutine

```
DSDPSetYBounds(DSDP dsdp, double minbound, double maxbound);
```

The second argument should be a negative number that is a lower bound of each variable y_i and the third argument is an upper bound of each variable. These bounds should not be tight. If one of the variables nearly equals the bound at the solution, the solver will return a termination code saying (D) is unbounded. To remove these bounds, set both the lower and upper bound to zero.

5.9 Iteration Monitor

A standard monitor that prints out the objective value and other relevant information at the current iterate can be set using the command

```
DSDPSetStandardMonitor(DSDP dsdp, int k);
```

A user can write a customized subroutine of the form

```
int (*monitor)(DSDP dsdp, void* ctx);
```

This subroutine will be called from the DSDP solver each iteration. It is useful for writing a specialized convergence criteria or monitoring the progress of the solver. The objective value and other information can be retrieved from the solver using the commands in the section 5.7. To set this subroutine, use the command `DSDPSetMonitor(DSDP dsdp, int (*monitor)(DSDP, void*), void* ctx);`

In this subroutine, the first argument is the solver, the second argument is the monitoring subroutine, and the third argument will be passed as the second argument in the monitoring subroutine. Examples of two monitors can be found in `DSDPROOT/src/solver/dsdpconverge.c`. The first monitor prints the solver statistics at each iteration and the second monitor determines the convergence of the solver. A monitor can also be used to print the time, duality gap, potential function at each iteration. Monitors have also been used to stop the solver after a specified time limit and change the parameters in the solver.

6 PDSDP

The DSDP package can also be run in parallel using multiple processors. In the parallel version, the Schur complement matrix is computed and solved in parallel. The parallel Cholesky factorization in PDSDP is performed using ScaLAPACK. The parallel Cholesky factorization in ScaLAPACK uses a two-dimensional block cyclic structure to distribute the data. The blocking parameter in ScaLAPACK determines how many rows and columns are in each block. Larger block sizes can be faster and reduce the overhead of passing messages, but smaller block sizes balance the work among the processors more equitably. PDSDP used a blocking parameter of 32 after experimenting with several choices. Since ScaLAPACK uses a dense matrix structure, this version is not appropriate when the Schur complement matrix is sparse.

The following steps should be used to run an application in parallel using PDSDP.

1. Install DSDP. Edit **DSDPROOT/make.include** to set the appropriate compiler flags.
2. Install ScaLAPACK. This package contains parallel linear solvers and is freely available to the public.
3. Go to the directory **DSDPROOT/src/pdsdp/scalapack/** and edit **Makefile** to identify the location of the DSDP and ScaLAPACK libraries.
4. Compile the PDSDP file **pdsdpscalapack.c**, which implements the additional operations. Then compile the executable **readsdpa.c**, which will read an SDPA file.

A PDSDP executable can be used much like serial version of DSDP that reads SDPA files. Given a SDPA file such as **truss1.dat-s**, the command

```
mpirun -np 2 dsdp5 truss1.dat-s -log_summary
```

will solve the problem using two processors. Additional processors may also be used. This implementation is best suited for very large problems.

Use of PDSDP as a subroutine library is also very similar to the use of the serial version of the solver. The application must create the solver and conic object on each processor and provide each processor with a copy of the data matrices, objective vector, and options. At the end of the algorithm, each solver has a copy of the solution. The routines to set the data and retrieve the solution are the same.

The few differences between the serial and parallel version are listed below.

1. All PDSDP programs must include the header file:

```
#include pdsdp5scalapack.h
```

2. Parallel applications should link to the DSDP library, the ScaLAPACK library, and the compiled source code in **dsdpscalapack.c**. Linking to the BLAS and LAPACK libraries are usually included while linking to ScaLAPACK.
3. The application should initialize and finalize MPI.
4. After creating the DSDP solver object, the application should call

```
int PDSDPUseSCALAPACKLinearSolver(DSDP dsdp);
```

5. The monitor should be set on only one processor. For example:

```
MPI_Comm_rank(MPI_COMM_WORLD, &rank);
if (rank==0){ info = DSDPSetStandardMonitor(dsdp); }
```

An example of the usage is provided in `DSDPROOT/pdsdp/ScaLAPACK/readsdpa.c`. Scalability of medium and large-scale problems has been achieved on up to 64 processors. The corrector direction does not scale well, so consider reducing the number of these steps for large groups of processors. See [1] for more details.

Source code that uses the parallel conjugate gradient method in PETSc to solve the linear systems is also included in the distribution.

A Brief History

DSDP began as a specialized solver for combinatorial optimization problems. Over the years, improvements in efficiency and design have enabled its use in many applications. Below is a brief history of DSDP.

- 1997 At the University of Iowa the authors release the initial version of DSDP. It solved the semidefinite relaxations of the maximum cut, minimum bisection, s-t cut, and bound constrained quadratic problems [6].
- 1999 DSDP version 2 increased functionality to address semidefinite cones with rank-one constraint matrices and LP constraints [5]. It was used specifically for combinatorial problems such as graph coloring, stable sets [2], and satisfiability problems.
- 2000 DSDP version 3 was a preliminary implementation of a general purpose SDP solver that addressed applications from the Seventh DIMACS Implementation Challenge on Semidefinite and Related Optimization Problems [7]. It ran in serial and parallel [1].
- 2002 DSDP version 4 added new sparse data structures to improve efficiency and precision. A Lanczos based line search and efficient iterative solver were added. It solved all problems in the SDPLIB collection that includes examples from control theory, truss topology design, and relaxations of combinatorial problems [3].
- 2004 DSDP version 5 [4] features a new efficient interface for semidefinite constraints, a corrector direction, and extensibility to structured applications in conic programming. Existence of the central path was ensured by bounding the variables. New applications from computational chemistry, global optimization, and sensor network location motivated the improvements in efficiency in robustness.

Acknowledgments

We thank Xiong Zhang and Cris Choi for their help in developing this code. Xiong Zhang, in particular, was fundamental to the initial version of DSDP. We also thank Hans Mittelmann[8] for his efforts in testing and benchmarking the different versions of the code. Finally, we thank Johan Löfberg, Stefan Ratschan, and all of the users who have commented on previous releases and suggested improvements to the software. Their contributions have made DSDP a more reliable, robust, and efficient package.

This work was supported by the Mathematical, Information, and Computational Sciences Division sub-program of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract W-31-109-ENG-38.

References

- [1] S. J. Benson. Parallel computing on semidefinite programs. Technical Report ANL/MCS-P939-0302, Mathematics and Computer Science Division, Argonne National Laboratory, March 2003.
- [2] S. J. Benson and Y. Ye. Approximating maximum stable set and minimum graph coloring problems with the positive semidefinite relaxation. In *Applications and Algorithms of Complementarity*, volume 50 of *Applied Optimization*, pages 1–18. Kluwer Academic Publishers, 2000.
- [3] S. J. Benson and Y. Ye. DSDP3: Dual-scaling algorithm for general positive semidefinite programming. Technical Report ANL/MCS-P851-1000, Mathematics and Computer Science Division, Argonne National Laboratory, February 2001.
- [4] S. J. Benson and Y. Ye. DSDP5: Software for semidefinite programming. Technical Report ANL/MCS-P1289-0905, Mathematics and Computer Science Division, Argonne National Laboratory, September 2005. Submitted to *ACM Transactions of Mathematical Software*.
- [5] S. J. Benson, Y. Ye, and X. Zhang. Mixed linear and semidefinite programming for combinatorial and quadratic optimization. *Optimization Methods and Software*, 11:515–544, 1999.
- [6] S. J. Benson, Y. Ye, and X. Zhang. Solving large-scale sparse semidefinite programs for combinatorial optimization. *SIAM Journal on Optimization*, 10(2):443–461, 2000.
- [7] DIMACS. The Seventh DIMACS Implementation Challenge: 1999-2000. <http://dimacs.rutgers.edu/Challenges/Seventh/>, 1999.
- [8] Hans D. Mittelmann. Benchmarks for optimization software, 2005. ftp://plato.la.asu.edu/pub/{sdplib.txt,sparse_sdp.txt,dimacs.txt}.

Index

- barrier, 5, 14, 15
- BCone, 11
- BConeAllocateBounds(), 11
- BConeCopyX(), 11
- BConeSetLowerBound(), 11
- BConeSetUpperBound(), 11
- BConeView(), 11
- block structure, 6
- bounds, 4, 15

- certificate, 4, 13
- combinatorial problems, 9
- convergence, 12–15

- DSDP, 6
- DSDPComputeX(), 12
- DSDPCreate(), 6
- DSDPCreateBCone(), 11
- DSDPCreateLPCone(), 10
- DSDPCreateSDPCone(), 6
- DSDPDestroy(), 12
- DSDPEventLogSummary(), 14
- DSDPFixVariable(), 11
- DSDPGetBarrierParameter(), 14
- DSDPGetDDObjective(), 13
- DSDPGetGapHistory(), 14
- DSDPGetPnorm(), 14
- DSDPGetPPObjective(), 13
- DSDPGetR(), 14
- DSDPGetRHistory(), 14
- DSDPGetSolutionType(), 13
- DSDPGetStepLengths(), 14
- DSDPGetTraceX(), 14
- DSDPGetY(), 13
- DSDPReuseMatrix(), 15
- DSDPSetBarrierParameter(), 15
- DSDPSetDObjective(), 6
- DSDPSetDualBound(), 12
- DSDPSetGapTolerance(), 12
- DSDPSetMaxIts(), 12
- DSDPSetMonitor(), 15
- DSDPSetPenaltyParameter(), 13, 15
- DSDPSetPotentialParameter(), 14
- DSDPSetPTolerance(), 13
- DSDPSetR0(), 14
- DSDPSetRTolerance(), 13
- DSDPSetStandardMonitor(), 12, 15
- DSDPSetup(), 12
- DSDPSetY0(), 14
- DSDPSetYBounds(), 13, 15
- DSDPSetZBar, 15

- DSDPSolutionType, 13
- DSDPSolve(), 12
- DSDPStopReason(), 13
- DSDPTerminationReason, 13, 14
- DSDPUsePenalty(), 15

- error check, 6, 13

- graph problems, 18

- header files, 6

- infeasible, 5, 13, 14
- iteration, 5, 12, 15

- LPCone, 10
- LPConeGetSArray(), 11
- LPConeGetXArray(), 11
- LPConeSetData(), 10
- LPConeSetXVec(), 11
- LPConeView(), 11

- maxcut, 6
- monitor, 15
- MPI, 16

- objective function, 6
- objective value, 5, 13

- parallel, 16
- PDSDP, 16
- penalty parameter, 4, 14, 15
- pnorm, 5, 14
- potential, 14
- potential function, 3
- print, 5, 15

- rho, 14

- ScaLAPACK, 16
- SDPA format, 6, 16
- SDPCone, 6
- SDPConeAddASparseVecMat(), 8
- SDPConeAddXVAV, 9
- SDPConeComputeS(), 8, 9
- SDPConeComputeX(), 8
- SDPConeGetBlockSize(), 8
- SDPConeGetFormat(), 9
- SDPConeGetXArray(), 8
- SDPConeSetADenseVecMat(), 7
- SDPConeSetASparseVecMat(), 7, 8, 10
- SDPConeSetBlockSize(), 6
- SDPConeSetSparsity(), 7

SDPConeSetStorageFormat(), 9
SDPConeSetXArray(), 9
SDPConeView(), 8
SDPConeViewDataMatrix(), 8
SDPConeViewX(), 8
SDPConeXVMultiply, 9
slack variable, 11
sparse data, 7, 8, 10
step length, 5, 14
surplus variable, 11
symmetric full storage format, 9
symmetric packed storage format, 7

time, 14
trace, 14

The submitted manuscript has been created by the University of Chicago as Operator of Argonne National Laboratory ("Argonne") under Contract No. W-31-109-ENG-38 with the U.S. Department of Energy. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.