DSDP5 User Guide — Software for Semidefinite Programming (for 64-bit Platforms)

Steve Benson, Ray-Hon Sun*, Yinyu Ye†

June 22, 2014

Abstract

DSDP uses a dual-scaling algorithm for the 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. This version is for the users to run DSDP with 64-bit data models (int64) on 64-bit architectures. The features of this 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,
- support of 64-bit data models to run on the 64-bit platforms, 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 this distribution.

*Institute of Computational and Mathematical Engineering, Stanford University, Stanford, CA U.S.A.
†Department of Management Science and Engineering, Stanford University, Stanford, CA U.S.A.

COPYRIGHT NOTIFICATION

rsun@cs.stanford.edu or yinyu-ye@stanford.edu.

Contents

1 Introduction

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

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

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

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 $\mathcal{A} : V \to \mathbb{R}^m$ defined by $(\mathcal{A}X)_i = \langle A_i, X \rangle$; its adjoint $\mathcal{A}^* : \mathbb{R}^m \to V$ is defined by $\mathcal{A}^* y = \sum_{i=1}^{m} y_i A_i$. Equivalent expressions for (P) and (D) can be written

$$\begin{array}{lllll} (P) & \inf \langle C, X \rangle & \text{subject to} & \mathcal{A}X = b, & X \in K, \\ (D) & \sup \; b^T y & \text{subject to} & \mathcal{A}^* y + S = C, & S \in K. \end{array}$$

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.

Assumes that the A_i are linearly independent, there exists $X \in \mathcal{F}^0(P)$, and a starting point $(y, S) \in \mathcal{F}^0(D)$ is known.

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

2 Dual-Scaling Algorithm

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) that 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

$$\begin{aligned} \mathcal{A}(X + \Delta X) &= b, & (1) \\ \mathcal{A}^*(\Delta y) + \Delta S &= 0, & (2) \\ \nu S^{-1} \Delta S S^{-1} + \Delta X &= \nu S^{-1} - X. & (3) \end{aligned}$$

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 matrix on 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 (2),(3), and $\Delta_\nu y$, we get

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

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

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

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.

A complete illustration of the dual scaling algorithm for semidefinite programming is available in the papers [1, 5].

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 $y_i (i = 1, \ldots, m)$ such that $l_i \leq y_i \leq u_i$ where $l_i, u_i \in \mathbb{R}$. 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 $\Gamma = 10^8$. 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:

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

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

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 [6]. When r is less than a very small, positive parameter ϵ_r, DSDP will classify y as feasible. 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. One 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 Data Structures and Parameters

4.1 Data Structures

The DSDP solver computes $\Delta' y$, $\Delta'' y$, $\Delta_\nu y$, $\|P(\nu)\|$, $\langle X(\nu), S\rangle$, and other quantities using operations on vectors, the Schur matrix, and the cones. Vector operations include sums and inner products. Operations on the Schur matrix include inserting elements and factoring the matrix. Objects representing a cone implement routines for computing its dual matrix S from y, evaluating the logarithmic barrier function, computing $\mathcal{A}S^{-1}$, and computing M. The solver object computes M, for example, by calling the corresponding operations on each cone and summing the results. The solver computes $\Delta' y$ through calls to the Schur matrix object that can factor the matrix and solve systems of linear equations. Table 2 shows eight of the primary data structures used in DSDP, the operations they implement, and the other objects required to implement those operations. For dense matrix structures, DSDP uses BLAS and LAPACK to operate on the data.

Table 2: Summary of primary data structures and their functionality.

Solver: Implements an algorithm for linear and semidefinite programming. *Operations:* $\Delta' y$, $\Delta'' y$, $\Delta_\nu y$, $\Delta^c y$, $\|P(\nu)\|$, $\langle X(\nu), S\rangle$, reduce ν. *Implementations:* Dual-Scaling Algorithm. *Instances:* one. *Requires:* Vector, Cone, Schur.
Vector: Represents y, $\mathcal{A}S^{-1}$, b, $\Delta' y$, $\Delta'' y$, $\Delta_\nu y$, and other internal work vectors. *Operations:* sum, inner product, norm. *Implementations:* dense one-dimensional array. *Instances:* about two dozen.
Schur: Represents M, the Schur complement of Newton equations. *Operations:* add to row, add to diagonal, factor, solve, vector-multiply. *Implementations:* sparse, dense, parallel dense. *Instances:* one.
Cone: Represents data C and A_i. *Operations:* check if $S \leftarrow C - \mathcal{A}^* y \in K$, $\ln\det S$, $\mathcal{A}S^{-1}$, M, $X(\nu)$. *Implementations:* SDP Cone, LP Cone, Bounds on y, Variable $r \geq 0$. *Instances:* three or more. *Requires:* Vector, Schur. *SDP Cone requires:* V Matrix, Data Matrices, S Matrix, DS Matrix.
SDP V Matrix: Represents X, $S^{-1}A_i S^{-1}$, and a buffer for $C - \mathcal{A}^* y$. *Operations:* $V \leftarrow 0$, $V \leftarrow V + \gamma w w^T$, get array. *Implementations:* dense. *Instances:* one per block.
SDP Data Matrix: Represents a symmetric data matrix. *Operations:* $V \leftarrow V + \gamma A$, $\langle V, A\rangle$, $w^T A w$, get rank, get eigenvalue/vector. *Implementations:* sparse, dense, identity, low-rank. *Instances:* up to $m+1$ per block.
SDP S Matrix: Represents S and checks whether $X(\nu) \succ 0$. *Operations:* $S \leftarrow V$, Cholesky factor, forward solve, backward solve, determinant. *Implementations:* sparse, dense. *Instances:* two per block.
SDP DS Matrix: Represents ΔS. *Operations:* $\Delta S \leftarrow V$, $w \leftarrow \Delta S v$. *Implementations:* dense, sparse, diagonal. *Instances:* one per block.

4.2 Parameters

Table 3 summarizes the significant parameters, options, and default values. The most important of these options is the initial variable r. By default, DSDP selects a value much larger than required to make $S \in \hat{K}$. Computational experience indicates that large values are more robust than smaller values. DSDP then sets the initial values of $\bar{z} = 10^1 0$ and $\nu = (\bar{z} - b^T y + \Gamma r)/(n\rho_n)$. Users can manually set $\bar{z}$ and ν, but choices better than the defaults usually require insight into the solution. The number of corrector steps can also significantly improve performance. In some examples, corrector steps can reduce the number of iterations by half—although the impact in total computation time is not as significant. Computational experience suggests that the number of corrector steps should be between 0 and 12, and the time spent in these steps should not exceed 30% [5].

Table 3: Summary of important parameters and initial values.

r: Dual infeasibility. *Default:* heuristic (large) *Suggested Values:* $10^2 - 10^{12}$ *Comments:* Larger values ensure robustness, but smaller values can significantly improve performance.
y: Initial solution. *Default:* 0 *Suggested Values:* Depends on data *Comments:* Initial points that improve performance can be difficult to find.
ρ_n: Bound ρ above by $n \times \rho_n$ and influence the barrier parameter. *Default:* 3.0 *Suggested Values:* $2.0 - 5.0$ *Comments:* Smaller values ensure robustness, but larger values can significantly improve performance.
kk_{max} : Maximum number of corrector steps. *Default:* 4 *Suggested Values:* $0 - 15$ *Comments:* For relatively small block sizes, increase this parameter.
Γ: The penalty parameter r and the bound on the trace of X. *Default:* $1e8$. *Suggested Values:* $10^3 - 10^{15}$ *Comments:* Larger values suitable unless (D) is feasible but has no interior.
l,u: Bounds on the variables y. *Default:* $-10^7, 10^7$ *Suggested Values:* Depends on the data. *Comments:* Tighter bounds do not necessarily improve performance.
$\bar{z}$: Upper bound on (D). *Default:* 10^{10} *Suggested Values:* Depends on the data. *Comments:* A high bound is usually sufficient.
ν: Dual barrier parameter. *Default:* Heuristic *Suggested Values:* Depends on the current solution. *Comments:* The default method sets $\nu = (\bar{z} - b^T y)/\rho$.
k_{max} : Maximum number of dual-scaling iterations. *Default:* 200 *Suggested Value:* $50 - 500$ *Comments:* Iteration counts of 20-60 are common.
η: Terminate when $(\bar{z} - b^T y)/(\lvert b^T y \rvert + 1)$ is less than η. *Default:* 10^{-6} *Suggested Values:* $10^{-2} - 10^{-7}$ *Comments:* Many problems do not require high accuracy.
ρ : Either a dynamic of a fixed value can be used. *Default:* Dynamic. *Suggested Values:* Dynamic *Comments:* The fixed strategy sets $\rho = n \times \rho_n$ and $\nu = (\bar{z} - b^T y)/\rho$, but its performance is usually inferior to the dynamic strategy.
ϵ_r, ϵ_P: Classify solutions as feasible. *Default:* 10^{-8}, 10^{-4} *Suggested Values:* $10^{-2} - 10^{-10}$ *Comments:* Adjust if the scaling of the problem is poor.

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 MEXfunction that reads data from the MATLABenvironment, passes the data to the solver, and returns the solution. The file `readsdpa.c` reads data from a file for data in SDPAformat, 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 solver. The subroutines used in these examples are described in this chapter. Further documentation on the routines and examples can be found in the HTML manual pages in `DSDPROOT/docs/dox/html/`.

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 subroutines, first create a solver object with

```
DSDPCreate(int64 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. The `DSDP` class is actually a pointer to a structure that contains the state of the solver. This subroutine will construct a new structure and point the `DSDP` variable to a new solver. A difference in typeset distinguishes the DSDP software from the `DSDP` class. Objects of this class apply the dual-scaling algorithm to the data.

Specify the objective function associated with these variables using the subroutine

```
DSDPSetDObjective(DSDP dsdp, int64 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,int64 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 require an `SDPCone` object 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, int64 blockj, int64 n)` can be used to specify the dimension of each block and the subroutine `SDPConeSetSparsity( SDPCone sdpcone, int64 blockj, int64 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 belongs to the user, and the performance of the problem depends upon this choice. In each of these subroutines, the first four arguments are an `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 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 `SDPCone` 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:

$$[\ a_{1,1} \quad a_{2,1} \quad a_{2,2} \quad a_{3,1} \quad a_{3,2} \quad a_{3,3} \quad \dots \quad a_{n,n}\ ]. \tag{8}$$

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,int64 blockj, int64 vari,
                       int64 n, double alpha, double val[], int64 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. `SDPCone` 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,int64 blockj, int64 vari, int64 n,
                        double alpha, int64 ishift
                        const int64 ind[], const double val[], int64 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 non-zeros in the lower triangular part of the matrix. The array of integers specifies which elements of the array (8) are included in the array of doubles. For example, the matrix

$$A_{i,j} = \begin{bmatrix} 3 & 2 & 0 \\ 2 & 0 & 6 \\ 0 & 6 & 0 \end{bmatrix} \tag{9}$$

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 `ishift` in the subroutine should be set to 0. In general, the argument `ishift` 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 `ishift` to `ishift` $+ n(n+1)/2 - 1$. Many applications, for instance, prefer to index the array from 1 to $n(n+1)/2$, setting the `index` argument to 1. The matrix (9) 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

$$\begin{array}{ll} \texttt{ind1} = [\ 0 \quad 1 \quad 4\ ] & \texttt{val1} = [\ 3 \quad 2 \quad 6\ ] \\ \texttt{ind2} = [\ 1 \quad 2 \quad 5\ ] & \texttt{val2} = [\ 3 \quad 2 \quad 6\ ] \\ \texttt{ind3} = [\ 7 \quad 3 \quad 5 \quad 4\ ] & \texttt{val3} = [\ 6 \quad 3 \quad 0 \quad 2\ ]. \\ \texttt{ind4} = [\ 0 \quad 1 \quad 4\ ] & \texttt{val4} = [\ 6 \quad 4 \quad 12\ ] \end{array}$$

As these examples suggest, there are many other ways to represent the sparse matrix. The non-zeros in the matrix do not have to be ordered, but ordering them may improve the efficiency of the solver. `SDPCone` 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,int64 blockj, int64 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,int64 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, int64 blockj, double *xmat[], int64 *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, int64 blockj, int64 n, double xmat[], int64 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 Xj are stored internally in `SDPCone` object. The subroutine

```
SDPConeViewX(SDPCone sdpcone, int64 blockj, int64 n, double xmat[], int64 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, int64 blockj, int64 *n);
```

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

The inner product of X_j with C_j, $A_{i,j}$, and I_j can be computed using the routine

```
int64 SDPConeAddADotX(SDPCone sdpcone,int64 blockj, double alpha, double xmat[], int64 nn,
                      double adotx[], int64 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 `alpha` times $\langle C_j, X_j\rangle$ to the initial element of the array, `alpha` times $\langle A_{i,j}, X_j\rangle$ to element i of the array, and `alpha` 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, int64 blockj, double c, double y[], int64 m, double r,
                int64 n, double smat[], int64 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 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, int64 blockj, int64 *dpsdefinite);
SDPConeXVMultiply(SDPCone sdpcone, int64 blockj, double v1[], double v2[], int64 n);
SDPConeAddXVAV(SDPCone sdpcone, int64 blockj, double v[], int64 n, double vav[], int64 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, $A_{1,j}, \ldots, A_{m,j}$, 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,int64 blockj, int64 n, double xmat[], int64 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 `SDPCone` object 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 **s**ymmetric 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

$$[\ a_{1,1} \quad 0 \quad \ldots \quad 0 \quad a_{2,1} \quad a_{2,2} \quad 0 \quad \ldots \quad 0 \quad \ldots \quad a_{n,1} \quad \ldots, \quad a_{n,n}\ ]. \tag{10}$$

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,int64 blockj, char UPLQ);
SDPConeGetStorageFormat(SDPCone sdpcone, int64 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 (9) can be set in the block j and variable i of the semidefinite cone using one of

the routine

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

where

$$\texttt{ind1} = \begin{bmatrix} 0 & 3 & 7 \end{bmatrix} \quad \texttt{val1} = \begin{bmatrix} 3 & 2 & 6 \end{bmatrix}.$$

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` object 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` object, 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 require an `LPCone` object 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, int64 n,
              const int64 nnzin[], const int64 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 non-zeros in column i ($i = 0, \ldots m$) (or row i of A). The length of the second and third array equals the number of nonzero in A and c. The arrays contain the nonzero 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):

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

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:

$$\begin{array}{rl} \texttt{nnzin} & = \begin{bmatrix} 0 & 3 & 5 & 7 \end{bmatrix} \\ \texttt{row} & = \begin{bmatrix} 0 & 1 & 2 & 0 & 1 & 0 & 1 & 2 \end{bmatrix} \\ \texttt{aval} & = \begin{bmatrix} 6.0 & 10.0 & 12.0 & 4.0 & 3.0 & 2.0 & 7.0 & -1.0 \end{bmatrix} \end{array}$$

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,int64 n, const int64 ik[],const int64 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:

$$\begin{array}{rl} \texttt{nnzin} & = \begin{bmatrix} 0 & 2 & 5 & 7 \end{bmatrix} \\ \texttt{row} & = \begin{bmatrix} 0 & 1 & 0 & 1 & 2 & 0 & 1 & 2 \end{bmatrix} \\ \texttt{aval} & = \begin{bmatrix} 4.0 & 3.0 & 2.0 & 7.0 & -1.0 & 6.0 & 10.0 & 12.0 \end{bmatrix} \end{array}$$

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[], int64 *n);
LPConeGetSArray(LPCone lpcone,double *sout[], int64 *n);
```

In these subroutines, the second argument sets a pointer to an an array of doubles containing the variables and the integer in the third argument will be set to the length of this array. These array were allocated by the `LPCone` object and the memory will be freed when the `DSDP` 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[], int64 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[], int64 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` object and point the `BCone` variable to the new structure. The bounds on the variables can be set using the subroutines:

```
BConeSetLowerBound(BCone bcone, int64 vari, double yi);
BConeSetUpperBound(BCone bcone, int64 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 bcone, int64 vari);
BConeSetPSlackVariable(BCone bcone, int64 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 bcone, int64 nbounds);` to tell the object how many bounds will be specified. This subroutine is optional, but is may improve the efficiency of the memory allocation. The subroutine `BConeSetXArray( BCone, double xout[], int64 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 bcone);` To retrieve the dual variables x corresponding to these constraints, use the subroutine `BConeCopyX( BCone bcone, double xl[], double xu[], int64 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 constraint64 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[],int64 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 $1, \ldots, m$. Alternatively, a single variable can be set to a value using the subroutine `DSDPFixVariable(DSDP dsdp, int64 vari, double val).`

5.4 Applying the Solver

After setting the data associated with the constraint64 cones, `DSDP` allocates 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` object. 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, int64 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 `1.0e-5` provides roughly five digits of accuracy. The subroutine

```
DSDPSetMaxIts(DSDP dsdp, int64 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 *pdfeasible);
```

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 $\|\mathcal{A}X - b\|_\infty / trace(X) > \epsilon_P$. 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 $\|\mathcal{A}X - b\|_\infty / trace(X) \leq \epsilon_P$. 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[], int64 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[], int64 history);
DSDPGetRHistory(DSDP dsdp, double rhistory[], int64 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, int64 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,int64 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 get 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, int64 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, int64 k);
```

A user can write a customized subroutine of the form

```
int64 (*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, int64 (*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 Installing and Configuring DSDP

The compressed tar file DSDP5.8_64.tar.gz contains an implementation of the dual-scaling algorithm for conic programming optimization problems.

6.1 Installing DSDP

Download all the files into a directory of your choice, e.g. create the DSDP5.8_64 directory structure and enter it. For example,

```
gunzip DSDP5.8_64.tar.gz
tar -xvf DSDP5.8_64.tar
cd DSDP5.8_64
```

DSDPis written in the C programming language. It has been tested using several different compilers, and 32-bit and 64-bit architectures.

6.2 Compiling DSDP

DSDPwas developed using Make – which is available on MacOS, Linux and most Unix systems.

To compile DSDP :

1. Enter the directory DSDP5.8_64 and edit the file `make.include` to define the **D**SDPROOT variable as the full directory name.
2. Verify whether your system has already defined int64. If int64 is not defined, then do the following:
   - Enable the definition of int64 in **D**SDPROOT/include/dsdpbasictypes.h

     ```
     typedef int64_t int64;
     ```
   - You need to check the setting of your compiler associated files, e.g. malloc.h in sys/dsdploginfo.c.
3. In the same file, edit the compiler flags `CC`, `OPTFLAGS`, and `CFLAGS`.

   ```
   CC         = <location of your C compiler>
   OPTFLAGS   = [compiler optimization options]
   DSDPTIMER  = [timing support]
   DSDPCFLAGS = [additional C compiler flags]
   CFLAGS     = [flags for C linker library]
   CLINKER    = [link DSDP library to application]
   ```
4. In the same file, edit the location of the BLAS and LAPACK libraries and include any other libraries required to link to them such as `-lg2c` or `-lm`.

   If you want to execute DSDP5.8 from command line on 64-bit machine, then you need to BLAS and LAPACK with 64-bit integer interface. Refer to `http://www.netlib.org/lapack/WishList/`.
5. If the MATLAB interface is required, also check and edit the `MEX` flag. If the DSDP MATLAB MEX function does not link with the library, then check the setting of MATLABMEX compiler.
6. Compile the source code to create the DSDP library and drivers.

```
make install
```

7. If problems persist, please send a copy of the compilation log to the developers.

6.3 Compiler Flags

1. BLAS and LAPACK : DSDP uses BLAS and LAPACK for many of the underlying operations and must be linked to these libraries. The location of these libraries should be specified in `make.include`.

```
LAPACKBLAS = -l<location of LAPACK library>
             -l<location of BLAS library>
```

   Note that these routines in BLAS and LAPACK are called from the C programming language under the assumption that the routine names are lower case and end with an underscore. The most common linking problem occurs when these assumptions are not true.

   Several compiler flags can be defined to change these assumptions. Define

   - `CAPSBLAS` if the names of BLAS and LAPACK routine names use all capital letters.
   - `NOUNDERBLAS` if no underscore is appended to the end of routine names. This flag was used to link DSDP to the reference BLAS available in MATLAB if using the Microsoft Compiler.
   - `__DSDP_NONAMEMANGLING` if a C++ compiler is being used and the BLAS and LAPACK routine names should not be changed.

   See the makefiles in the distribution for examples of using these terms. Note that those compiling in the Microsoft Windows Operating System usually need to define the `NOUNDERBLAS` flag.

   Please check whether your BLAS and LAPACK packages support 64-bit data models, hence the `integer` type in Fortran is the same size as a `long long int` in C, and a `double precision` variable in Fortran is the same size as a `double` in C. If problems persist, the macros and type definitions in the file `dsdplapack.c`, located in the `DSDPROOT/include/` directory will have to be edited.

2. Timing: DSDP can provide time profiles for several important operations. By default, the timing routine is not implemented due to portability issues among architectures. However, the following flags may defined for the compiler to activate timing routines:

   - `DSDP_MS_TIME` activates the timing utility from the Microsoft compiler
   - `DSDP_TIME` activates the timing utility from the GCC and many other compilers.

   The routines that call these two timing utilities can be found in the file `dsdptime.c` in `DSDPROOT/src/sys/` and edited.

3. MATLAB : If the MATLAB interface is be generated,

   several variables for MATLAB macros should be defined, e.g.

```
MATLAB_HOME = <location of MATLAB directory>
MATLAB_BIN  = ${MATLAB_HOME}/bin
MEX_H       = ${MATLAB_HOME}/extern/include/
MEX_LIB     = ${MATLAB_HOME}/extern/lib/maci64/ \% for mac
```

   Note that some options for the MEX compilation for 64-bit MATLAB.

```
mexOPTIMFLAGS = -O -largeArrayDims
mexLIBS       = -lm -L\${MEX_LIB}
mexFLAG       = ${mexOPTIMFLAGS} -I\${MEX_H}
```

You may want to modify the following according to your configuration.

```
MEX           = mex -v -largeArrayDims -lmwlapack -lmwblas -O
DSDPMATLABDIR = $DSDPROOT/matlab
```

Note that for Mac users, please check your Xcode setup. If you encounter the following error:

```
xcodebuild: error: SDK "macosx10.7" cannot be located.
```

, then please follow the "official" solution of

`http://www.mathworks.com/matlabcentral/answers/103904`.

to correct your configuration.

4. Testing:

   - Run the executables by switching to directory `DSDPROOT/exec` and typing

     ```
     > dsdp5  truss1.dat-s
     > maxcut  graph1
     > theta  graph1
     ```

     Compare the output with the files `output.truss1`, `output.maxcut`, and `output.theta`. If the output from any of the tests differs significantly from the files, please report it to the developers. Note that this step is valid only for the 32-bit platforms.

   - DSDP can be called from MATLAB version 6.0 and higher. Run the sample problems by starting MATLAB in the `DSDPROOT/matlab` directory and typing

     ```
     > check;
     ```

     Compare the output with the output in `check.out`. For help using the package, type

     ```
     > help dsdp;
     ```

     Several example MATLAB files have been provided in `DSDPROOT/matlab` that create example problems, read data files, and verify solutions.

7 Applying DSDP to Graph Problems

We illustrate the application of DSDP to solve four types of problems in graph theory using semidefinite cones with rank-one constraint matrices and LP constraints. Sample code is included in the directory `DSDPROOT/examples/`.

A program `maxcut.c` reads a file containing a graph, generates the semidefinite relaxation of a maximum cut problem, and solves the relaxation. For example,

```
> maxcut graph1
```

reads the graph in the file `graph1` and solves this graph problem. The first line of the graph should contain two integers. The first integer states the number of nodes in the graph, and the second integer states the number of edges.

Subsequent lines have two or three entries separated by a space. The first two entries specify the two nodes that an edge connects. The optional third entry specifies the weight of the edge. If no weight is specified, a weight of 1 will be assigned.

The second example is for the Lovász θ problem using `theta.c` that reads graph complement from a file, formulate the Lovász theta problem, and solve it using DSDP. For example,

```
> theta graph1
```

to compute the Lovász theta number for a graph. This number is an upper bound for the maximum clique of a graph, a lower bound for the minimal graph coloring, and serves as a bound for several other combinatorial graph problems. The number is the solution to a semidefinite program.

The third example is to solve maximum stable set problems using `stable.c` that reads graph from a file, formulate the maximum stable (or independent) set problem, and solve it using DSDP.

The fourth example is to solve minimum graph coloring problems using `color.c` that read graph from a file, formulate the semidefinite relaxation of k-coloring problem, solve it using DSDP, and apply randomized algorithm to generate approximate solutions.

8 DSDP with MATLAB

Additional help using the DSDP can be found by typing `help dsdp` in the directory `DSDP5.X`. The command

```
> [STAT, y, X] = dsdp(b, AC)
```

attempts to solve the semidefinite program by using a dual-scaling algorithm. The first argument is the objective vector b in (D) and the second argument is a cell array that contains the structure and data for the constraint cones. Most data has a block structure, which should be specified by the user in the second argument. For a problem with p cones of constraints, `AC` is a $p \times 3$ cell array. Each row of the cell array describes a cone. The first element in each row of the cell array is a string that identifies the type of cone. The second element of the cell array specifies the dimension of the cone, and the third element contains the cone data.

8.1 Semidefinite Cones

If the jth cone is a semidefinite cone consisting of a single block with n rows and columns in the matrices, then the first element in this row of the cell array is the string `'SDP'` and the second element is the number `n`. The third element in this row of the cell array is a sparse matrix with $n(n+1)/2$ rows and $m+1$ columns. Columns 1 to m of this matrix represent the constraints $A_{1,j}, \ldots, A_{m,j}$ for this block and column $m+1$ represents C_j.

The square symmetric data matrices $A_{i,j}$ and C_j map to the columns of `AC{j,3}` through the operator **dvec(·)** : $\mathbb{R}^{n\times n} \to \mathbb{R}^{n(n+1)/2}$, which is defined as

$$\mathbf{dvec}(A) \quad = [a_{1,1} \quad a_{1,2} \quad a_{2,2} \quad a_{1,3} \quad a_{2,3} \quad a_{3,3} \quad \ldots \quad a_{n,n}]^T.$$

In this definition, $a_{k,l}$ is the element in row k and column l of A. This ordering is often referred to as symmetric packed storage format. The inverse of `dvec( )` is $\mathbf{d}mat(\cdot) : \mathbb{R}^{n(n+1)/2} \to \mathbb{R}^{n\times n}$, which converts the vector into a square symmetric matrix. Using these operations,

$$A_{i,j} = \texttt{dmat(AC\{j,3\}(:,i))}, \qquad C_j = \texttt{dmat(AC\{j,3\}(:,m+1))}$$

and

$$\texttt{AC\{j,3\}} = [\ \texttt{dvec}(A_{1,j}) \quad \ldots \quad \texttt{dvec}(A_{m,j}) \quad \texttt{dvec}(C_j)\ ];$$

For example, the problem:

$$\begin{array}{lcccccc} \text{Maximize} & & & y_1 & + & & y_2 \\ \text{Subject to} & \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} & & y_1 & + & \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} & y_2 \quad \preceq \begin{bmatrix} 4 & -1 \\ -1 & 5 \end{bmatrix} \end{array}$$

can be solved by:

```
> b = [ 1 1 ]';
> AAC = [ [ 1.0 0 0 ]' [ 0 0 1.0 ]' [ 4.0 -1.0 5.0 ]' ];
> AC{1,1} = 'SDP';
> AC{1,2} = [2];
> AC{1,3} = sparse(AAC);
> [STAT,y,X]=dsdp(b,AC);
> XX=dmat(X{1});
```

The solution y is the column vector `y' = [ 3 4 ]'` and the solution X is a $p \times 1$ cell array. In this case, `X = [ 3 x 1  double ]`, `X{1}'=[ 1.0 1.0 1.0 ]`, and `dmat(X{1})=[ 1 1 ; 1 1 ]`.

Each semidefinite block can be stated in a separate row of the cell array; only the available memory on the machine limits the number of cones that can be specified.

Each semidefinite block may, however, be grouped into a single row in the cell array. To group these blocks together, the second cell entry must be an array of integers stating the dimension of each block. The data from the blocks should be concatenated such that the number of rows in the data matrix increases whereas the number of columns remains constant. The following lines indicate how to group the semidefinite blocks in rows 1 and 2 of cell array `AC1` into a new cell array `AC2`

```
> AC2{1,1} = 'SDP';
> AC2{1,2} = [AC1{1,2}  AC1{2,2}];
> AC2{1,3} = [AC1{1,3}; AC1{2,3}]
```

The new cell array `AC2` can be passed directly into DSDP. The advantage of grouping multiple blocks together is that it uses less memory – especially when there are many blocks and many of the matrices in these blocks are zero. The performance of DSDP, measured by execution time, will change very little.

This distribution contains several examples files in SDPA format. A utility routine called `readsdpa( · )` can read these files and put the problems in DSDP format. They may serve as examples on how to format an application for use by the `DSDP` solver. Another example can be seen in the file `maxcut( · )` , which takes a graph and creates an SDP relaxation of the maximum cut problem from a graph.

8.2 LP Cones

A cone of LP variables can specified separately. For example a randomly generated LP cone $A^T y \leq c$ with 3 variables y and 5 inequality constraints can be specified in the following code.

```
> n=5; m=3;
> b = rand(m,1);
> At=rand(n,m);
> c=rand(n,1);
> AC{1,1} = 'LP';
> AC{1,2} = n;
> AC{1,3} = sparse([At c]);
> [STAT,y,X]=dsdp(b,AC);
```

Multiple cones of LP variables may be passed into the DSDP solver, but for efficiency reasons, it is best to group them all together. This cone may also be passed to the DSDP solver as a semidefinite cone, where the matrices A_i and C are diagonal. For efficiency reasons, however, it is best to identify them separately as belonging the the cone of 'LP' variables.

Although y variables that are fixed to a constant can be preprocessed and removed from a model, it is often more convenient to leave them in the model. It is more efficient for to identify fixed variables to DSDP than to model these constraints as a pair of linear inequalities. The following example sets variables 1 and 8 to the values 2.4 and -6.1, respectively.

```
> AC{j,1} = 'FIXED'; AC{j,2} = [ 1 8 ]; AC{j,3} = [ 2.4 -6.1 ];
```

The corresponding variables x to these constraints may be positive or negative.

8.3 Solver Options

There are more ways to call the solver. The command

```
> [STAT,y,X] = DSDP(b,AC,OPTIONS)
```

specifies some options for the solver. The `OPTIONS` structure may contain any of the following fields that may *significantly* affect the performance of the solver. Options that affect the formulation of the problem are:

`r0` initial value for r in (DD). If $r0 < 0$, a heuristic will select a very large number (1e10). IMPORTANT: To improve convergence, use a smaller value. [default -1 (Heuristic)].

`zbar` an upper bound $\bar{z}$ on the objective value at the solution [default 1.0e10].

`penalty` penalty parameter Γ in (DD) that enforces feasibility in (D). IMPORTANT: This parameter must be positive and greater than the trace of the solution X of (P). [default 1e8].

`boundy` determines the bounds l and u on the variables y in (DD). That is, $-boundy = l \leq y_i \leq u = boundy$ for all i. [default: 1e7].

Fields in the `OPTIONS` structure that affect the stopping criteria for the solver are:

`gaptol` tolerance for duality gap as a fraction of the value of the objective functions [default 1e-6].

`maxit` maximum number of iterations allowed [default 1000].

`steptol` tolerance for stopping because of small steps [default 1e-2].

`pnormtol` $\|P(\nu)\|$ of solution should also be less than [default 1e30].

`inftol` the value r in (DD) must be less than this tolerance to classify the final solution of (D) as feasible. [default 1e-8].

`dual_bound` Terminate the solver when it finds a feasible point of (D) with an objective greater than this value. (Helpful in branch-and-bound algorithms.) [default 1e+30].

Fields in the `OPTIONS` structure that affect printing are:

`print` = k to display output at each k iteration, else = 0 [default 10].

`logtime` =1 to profile the performance of DSDP subroutines, else =0. (Assumes proper compilation flags.)

`cc` add this constant the objective value. This parameter is algorithmically irrelevant, but it can make the objective values displayed on the screen more consistent with the underlying application [default 0].

Other fields recognized in `OPTIONS` structure are:

`rho` to set the potential parameter ρ in the function (**??**) to this multiple of the conic dimension n. [default: 3] IMPORTANT! Increasing this parameter to 4 or 5 may significantly improve performance.

`dynamicrho` to use dynamic rho strategy. [default: 1].

`bigM` if > 0, the variable r in (DD) will remain positive (as opposed to nonnegative). [default 0].

`mu0` initial barrier parameter ν. [default -1: use heuristic]

`reuse` sets a maximum on the number of times the Schur complement matrix can be reused. Larger numbers reduce the number of iterations but increase the cost of each iteration. Applications requiring few iterations (¡60) should consider setting this parameter to 0. [default: 4]

For instance, the commands

```
> OPTIONS.gaptol = 0.001;
> OPTIONS.boundy = 1000;
> OPTIONS.rho = 5;
> [STAT,y,X] = DSDP(b,AC,OPTIONS);
```

asks for a solution with approximately three significant digits, bound the y variables by -1000 and $+1000$, and use a potential parameter ρ of 5 times the conic dimension. Some of these fields, especially `rho`, `r0`, and `ybound` can significantly improve performance of the solver.

Using a fourth input argument, the command

```
> [STAT,y,X] = DSDP(b,AC,OPTIONS,y0);
```

specifies an initial solution `y0` in (D). The default starting vector is the zero vector.

8.4 Solver Performance and Statistics

The second and third output arguments return objective values for (D) and (P), respectively.

The first output argument is a structure with several fields that describe the solution of the problem:

`stype`	`PDFeasible` if the solutions to both (D) and (P) are feasible, `Infeasible` if (D) in infeasible, and and `Unbounded` if (D) is unbounded.
`obj`	an approximately optimal objective value.
`pobj`	objective value of (PP).
`dobj`	objective value of (DD).
`stopcode`	equals 0 if the solutions to (PP) and (DD) satisfy the prescribed tolerances and equals nonzero if the solver terminated for other reasons.

Additional fields describe characteristics of the solution:

`tracex`	the trace of the solution X of (P).
`r`	the multiple of the identity element added to $C-\mathcal{A}^T(y)$ in the final solution to make S positive definite.
`mu`	the final barrier parameter (ν).
`ynorm`	the largest element of y (infinity norm).
`boundy`	the bounds placed on the magnitude of each variable y.
`penalty`	the penalty parameter Γ used by the solver, which must be greater than the trace of the variables X in (P). (see above).

Additional fields provide statistics from the solver:

`iterations`	number of iterations used by the algorithm.
`pstep`	the final step length.in (PP)
`dstep`	the final step length in (DD).
`pnorm`	the final value $\|P(\nu)\|$.

`rho`	the potential parameter (as a multiple of the total dimension of the cones).
`gaphist`	a history of the duality gap.
`infhist`	a history of the variable r in (DD).
`datanorm`	the Frobenius norm of C, A and b.

DSDP has also provides several utility routines. The utility `derror( · )` verifies that the solution satisfies the constraints and that the objective values (P) and (D) are equal. The errors are computed according to the the standards of the DIMACS Challenge.

9 Reading SDPA files

DSDP can be used if the user has a problem written in sparse SDPA format[8]. These executables have been put in the directory `DSDPROOT/exec/`. The file name should follow the executable. For example,

```
> dsdp5 truss4.dat-s
```

Other options can also be used with DSDP. These should follow the SDPA filename.

`-gaptol <rtol>`	to stop the problem when the relative duality gap is less than this number.
`-mu0 <mu0>`	to specify the initial barrier parameter ν.
`-r0 <r0>`	to specify the initial value of r in (DD).
`-boundy <1e7>`	to bound the magnitude of each variable y in (DD).
`-save <filename>`	to save the solution into a file with a format similar to SDPA.
`-y0 <filename>`	to specify an initial vector y in (D).
`-maxit <iter>`	to stop the problem after a specified number of iterations.
`-rho <3>`	to set the potential parameter ρ to this multiple of the conic dimension n.
`-dobjmin <dd>`	to add a constraint that sets a lower bound on the objective value at the solution.
`-penalty <1e8>`	to set the penalty parameter Γ for infeasibility in (D).
`-print <1>`	print standard output at each k iteration.
`-bigM <0>`	treat the inequality $r \geq 0$ in (DD) as other inequalities and keep it positive.
`-dloginfo <0>`	to print more detailed output. Higher number produce more output.
`-dlogsummary <1>`	to print detailed timing information about each dominant computations.

10 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 infeasiblity 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: $\|\nabla\psi\|_{M^{-1}}$.

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.

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 [1].

1999 DSDP version 2 increased functionality to address semidefinite cones with rank-one constraint matrices and LP constraints. It was used specifically for combinatorial problems such as graph coloring, stable sets, and satisfiability problems.

2000 DSDP version 3 [2] 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.

2002 DSDP version 4 [3] 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.

2004 DSDP version 5 [5] 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. The latest release of version 5 is 5.8 [4].

2014 DSDP version 5.8_64 features a new support 64-bit data models to run DSDP on 64-bit platforms. The interfaces to access 64-bit MATLAB libraries and read SDPAfiles was also updated.

Acknowledgments

We thank for Dr. Ericke DeLage's generous discussion and support of testing on the 64-bit DSDP software. He provides testDSDP.m in the matlab directory to test the DSDP setup with MATLAB interface.

References

[1] Steven J. Benson, Yinyu Ye, and Xiong Zhang, Solving large-scale sparse semidefinite programs for combinatorial optimization. *SIAM Journal on Optimization*, 10(2):443–461, 2000.

[2] Steven J. Benson and Yinyu Ye, DSDP3: Dual-Scaling Algorithm for General Positive Semidefinite Programming, ANL/MCS-P851-1000, Mathematics and Computer Science Division, Argonne National Laboratory, February 2001.

[3] Steven J. Benson and Yinyu Ye, DSDP4: A Software Package Implementing the Dual-Scaling Algorithm for Semidefinite Programming, ANL/MCS-TM-255, Mathematics and Computer Science Division, Argonne National Laboratory, June 2002.

[4] Steven J. Benson and Yinyu Ye, DSDP5 User Guide — Software for Semidefinite Programming. ANL/MCS-TM-277, Mathematics and Computer Science Division, Argonne National Laboratory, July, 2004.

[5] Steven J. Benson and Yinyu Ye, DSDP5: Software for semidefinite programming. ANL/MCS-P1289-0905, Mathematics and Computer Science Division, Argonne National Laboratory, September 2005.

[6] G.H. Golub and C.F. Van Loan, Matrix Computations, second edition, Johns Hopkins University Press, Baltimore, MD, 1989.

[7] Hans D. Mittelmann, Benchmarks for optimization software, 2005. `ftp://plato.la.asu.edu/pub/{sdplib.txt,sparse_sdp.txt,dimacs.txt}`.

[8] K. Fujisawa, M. Kojima, K. Nakata, and M. Yamashita SDPA (SemiDefinite Programming Algorithm) User's Manual - Version 6.00 `http://www.is.titech.ac.jp/~kojima/articles/b-308.ps`.

Index