



MOSEK Rmosek package
Release 9.0.98

MOSEK ApS

15 July 2019

Contents

1	Introduction	1
1.1	Why the Rmosek package?	2
2	Contact Information	3
3	License Agreement	4
4	Installation	6
4.1	System requirements	6
4.2	Installation	7
4.3	Testing the Installation	7
5	Design Overview	8
5.1	Modeling	8
5.2	“Hello World!” in MOSEK	8
6	Optimization Tutorials	10
6.1	Linear Optimization	10
6.2	Quadratic Optimization	12
6.3	Conic Quadratic Optimization	14
6.4	Power Cone Optimization	16
6.5	Conic Exponential Optimization	18
6.6	Semidefinite Optimization	20
6.7	Affine conic constraints (new)	22
6.8	Geometric Programming	25
6.9	Integer Optimization	27
6.10	Problem Modification and Reoptimization	30
7	Solver Interaction Tutorials	35
7.1	Accessing the solution	35
7.2	Errors and exceptions	38
7.3	Input/Output	39
7.4	Setting solver parameters	41
7.5	Retrieving information items	41
8	Debugging Tutorials	43
8.1	Understanding optimizer log	43
8.2	Addressing numerical issues	47
8.3	Debugging infeasibility	49
8.4	Python Console	53
9	Technical guidelines	56
9.1	Multithreading	56
9.2	Parallel optimization Using the The Multicore Package	56
9.3	The license system	57
10	Case Studies	58
10.1	Portfolio Optimization	58

10.2	Least Squares and Other Norm Minimization Problems	70
11	Problem Formulation and Solutions	76
11.1	Linear Optimization	76
11.2	Conic Optimization	79
11.3	Semidefinite Optimization	83
11.4	Quadratic and Quadratically Constrained Optimization	84
11.5	Affine Conic Constraints	85
12	Optimizers	87
12.1	Presolve	87
12.2	Linear Optimization	89
12.3	Conic Optimization - Interior-point optimizer	95
12.4	The Optimizer for Mixed-integer Problems	99
13	Rmosek API Reference	104
13.1	Command Reference	104
13.2	Parameters grouped by topic	108
13.3	Parameters (alphabetical list sorted by type)	119
13.4	Response codes	158
13.5	Enumerations	176
13.6	Nonlinear interfaces (obsolete)	200
14	Supported File Formats	202
14.1	The LP File Format	203
14.2	The MPS File Format	208
14.3	The OPF Format	219
14.4	The CBF Format	228
14.5	The PTF Format	242
14.6	The Task Format	246
14.7	The JSON Format	247
14.8	The Solution File Format	254
15	List of examples	257
16	Interface changes	258
16.1	Backwards compatibility	258
16.2	New API	258
16.3	Parameters	258
16.4	Constants	259
16.5	Response Codes	261
	Bibliography	263
	Symbol Index	264
	Index	276

Chapter 1

Introduction

The **MOSEK** Optimization Suite 9.0.98 is a powerful software package capable of solving large-scale optimization problems of the following kind:

- linear,
- conic:
 - conic quadratic (also known as second-order cone),
 - involving the exponential cone,
 - involving the power cone,
 - semidefinite,
- convex quadratic and quadratically constrained,
- integer.

In order to obtain an overview of features in the **MOSEK** Optimization Suite consult the [product introduction](#) guide.

The most widespread class of optimization problems is *linear optimization problems*, where all relations are linear. The tremendous success of both applications and theory of linear optimization can be ascribed to the following factors:

- The required data are simple, i.e. just matrices and vectors.
- Convexity is guaranteed since the problem is convex by construction.
- Linear functions are trivially differentiable.
- There exist very efficient algorithms and software for solving linear problems.
- Duality properties for linear optimization are nice and simple.

Even if the linear optimization model is only an approximation to the true problem at hand, the advantages of linear optimization may outweigh the disadvantages. In some cases, however, the problem formulation is inherently nonlinear and a linear approximation is either intractable or inadequate. *Conic optimization* has proved to be a very expressive and powerful way to introduce nonlinearities, while preserving all the nice properties of linear optimization listed above.

The fundamental expression in linear optimization is a linear expression of the form

$$Ax - b \geq 0.$$

In conic optimization this is replaced with a wider class of constraints

$$Ax - b \in \mathcal{K}$$

where \mathcal{K} is a *convex cone*. For example in 3 dimensions \mathcal{K} may correspond to an ice cream cone. The conic optimizer in **MOSEK** supports a number of different types of cones \mathcal{K} , which allows a surprisingly large number of nonlinear relations to be modeled, as described in the **MOSEK Modeling Cookbook**, while preserving the nice algorithmic and theoretical properties of linear optimization.

1.1 Why the Rmosek package?

The Rmosek package provides access to most functionalities of **MOSEK** from an R-language software environment. The package is adjusted for the typical R user.

The Rmosek package provides access to:

- Linear Optimization (LO)
- Conic Quadratic (Second-Order Cone) Optimization (CQO, SOCO)
- Power Cone Optimization
- Conic Exponential Optimization (CEO)
- Convex Quadratic Optimization (QO)
- Semidefinite Optimization (SDO)
- Mixed-Integer Optimization (MIO)

Chapter 2

Contact Information

Phone	+45 7174 9373	
Website	mosek.com	
Email		
	sales@mosek.com	Sales, pricing, and licensing
	support@mosek.com	Technical support, questions and bug reports
	info@mosek.com	Everything else.
Mailing Address		
	MOSEK ApS	
	Fruebjergvej 3	
	Symbion Science Park, Box 16	
	2100 Copenhagen O	
	Denmark	

You can get in touch with **MOSEK** using popular social media as well:

Blogger	https://blog.mosek.com/
Google Group	https://groups.google.com/forum/#!forum/mosek
Twitter	https://twitter.com/mosektw
Google+	https://plus.google.com/+Mosek/posts
Linkedin	https://www.linkedin.com/company/mosek-aps

In particular **Twitter** is used for news, updates and release announcements.

Chapter 3

License Agreement

Before using the **MOSEK** software, please read the license agreement available in the distribution at <MSKHOME>/mosek/9.0/mosek-eula.pdf or on the **MOSEK** website <https://mosek.com/products/license-agreement>.

MOSEK uses some third-party open-source libraries. Their license details follows.

zlib

MOSEK includes the *zlib* library obtained from the [zlib website](#). The license agreement for *zlib* is shown in [Listing 3.1](#).

Listing 3.1: *zlib* license.

```
zlib.h -- interface of the 'zlib' general purpose compression library
version 1.2.7, May 2nd, 2012

Copyright (C) 1995-2012 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied
warranty. In no event will the authors be held liable for any damages
arising from the use of this software.

Permission is granted to anyone to use this software for any purpose,
including commercial applications, and to alter it and redistribute it
freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not
   claim that you wrote the original software. If you use this software
   in a product, an acknowledgment in the product documentation would be
   appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be
   misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly          Mark Adler
jloup@gzip.org            madler@alumni.caltech.edu
```

fplib

MOSEK includes the floating point formatting library developed by David M. Gay obtained from the [netlib website](#). The license agreement for *fplib* is shown in [Listing 3.2](#).

Listing 3.2: *fplib* license.

```
/*****
 *
```

(continues on next page)

```
* The author of this software is David M. Gay.
*
* Copyright (c) 1991, 2000, 2001 by Lucent Technologies.
*
* Permission to use, copy, modify, and distribute this software for any
* purpose without fee is hereby granted, provided that this entire notice
* is included in all copies of any software which is or includes a copy
* or modification of this software and in all copies of the supporting
* documentation for such software.
*
* THIS SOFTWARE IS BEING PROVIDED "AS IS", WITHOUT ANY EXPRESS OR IMPLIED
* WARRANTY. IN PARTICULAR, NEITHER THE AUTHOR NOR LUCENT MAKES ANY
* REPRESENTATION OR WARRANTY OF ANY KIND CONCERNING THE MERCHANTABILITY
* OF THIS SOFTWARE OR ITS FITNESS FOR ANY PARTICULAR PURPOSE.
*
*****/
```

Zstandard

MOSEK includes the *Zstandard* library developed by Facebook obtained from [github/zstd](https://github.com/facebook/zstd). The license agreement for *Zstandard* is shown in [Listing 3.3](#).

Listing 3.3: *Zstandard* license.

```
BSD License

For Zstandard software

Copyright (c) 2016-present, Facebook, Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification,
are permitted provided that the following conditions are met:

* Redistributions of source code must retain the above copyright notice, this
  list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice,
  this list of conditions and the following disclaimer in the documentation
  and/or other materials provided with the distribution.

* Neither the name Facebook nor the names of its contributors may be used to
  endorse or promote products derived from this software without specific
  prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE
DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR
ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES
(INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON
ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
```

Chapter 4

Installation

In this section we discuss how to install and setup the **MOSEK** Rmosek package.

Important: Before running this **MOSEK** interface please make sure that you:

- Installed **MOSEK** correctly. Some operating systems require extra steps. See the [Installation guide](#) for instructions and common troubleshooting tips.
 - Set up a license. See the [Licensing guide](#) for instructions.
-

Compatibility

The Rmosek package can be used with R version at least 3.5.

Locating files in the MOSEK Optimization Suite

The relevant files for the Rmosek package are organized as reported in [Table 4.1](#).

Table 4.1: Relevant files for the Rmosek package.

Relative Path	Description	Label
<MSKHOME>/mosek/9.0/tools/platform/<PLATFORM>/rmosek	Rmosek install files	<RMOSEKDIR>
<MSKHOME>/mosek/9.0/tools/examples/rmosek	Examples	<EXDIR>
<MSKHOME>/mosek/9.0/tools/examples/data	Additional data	<MISCDIR>

where

- <MSKHOME> is the folder in which the **MOSEK** Optimization Suite has been installed,
- <PLATFORM> is the actual platform among those supported by **MOSEK**, i.e. `win32x86`, `win64x86`, `linux64x86` or `osx64x86`.

4.1 System requirements

There are no pre-compiled binary versions of the Rmosek package available, so the system must be able compile C++ source code and have access to a few command line tools. The summary of relevance to the Rmosek package given below is based on the official guide, [R Installation and Administration](#), for installing packages in source form.

Windows

- Download [Rtools](#) (not an R package). You will need to install the *R toolset*, the *Cygwin DLLs*, and the *toolchain*.
- Make sure you have `mosek` and the executables of Rtools on the `PATH` environment variable.

MacOS

- Make sure you have **Xcode** installed.

Linux

- On Ubuntu (and Debian) you may install the **r-base-dev** package.

4.2 Installation

To install the Rmosek package you may run the `install.rmosek` function of the script `<RMOSEKDIR>/builder.R`. This is a wrapper around the `install.packages` function with recommended default argument values (e.g., to a compatible **MOSEK** repository), cross-platform support of configurations variables, and helpers to resolve Rtools on Windows. As example, if called with no arguments, it will attempt an autoconfigured installation:

```
source("<RMOSEKDIR>/builder.R")
attachbuilder()
install.rmosek()
```

You can inspect the autoconfigured defaults by typing `show(install.rmosek)`. Apart from the typical arguments of `install.packages`, a user might want to change:

- **MSK_BINDIR**: The location of the **MOSEK** library to compile against; e.g., `mosek/<MSKVER>/tools/platform/<PLATFORM>/bin`.
- **using_pkgbuild**: Whether to run the installation in the build-friendly environment of package `pkgbuild` (helps, e.g., to resolve Rtools on Windows).
- **using_sysenv**: Whether to transmit configuration variables via `Sys.setenv()` as opposed to `configure.vars` (depends on platform support).

4.3 Testing the Installation

First of all, to check that the Rmosek package was properly installed, start R and try

```
require("Rmosek")
```

The installation can further be tested by running some of the enclosed examples. Open a terminal, change folder to `<EXDIR>` and use R to run a selected example, for instance:

```
R -f 1o1.R
```

Chapter 5

Design Overview

5.1 Modeling

Rmosek package is an interface for specifying optimization problems directly in matrix form. It means that an optimization problem such as:

$$\begin{array}{ll}\text{minimize} & c^T x \\ \text{subject to} & Ax \leq b, \\ & x \in \mathcal{K}\end{array}$$

or

$$\begin{array}{ll}\text{minimize} & c^T x \\ \text{subject to} & Ax \leq b, \\ & Fx + g \in \mathcal{K}\end{array}$$

is specified by describing the matrices A , F , vectors b, c, g and a list of cones \mathcal{K} directly.

The main characteristics of this interface are:

- **Simplicity:** once the problem data is assembled in matrix form, it is straightforward to input it into the optimizer.
- **Exploiting sparsity:** data is entered in sparse format, enabling huge, sparse problems to be defined and solved efficiently.
- **Efficiency:** the API incurs almost no overhead between the user's representation of the problem and MOSEK's internal one.

Rmosek package does not aid with modeling. It is the user's responsibility to express the problem in MOSEK's standard form, introducing, if necessary, auxiliary variables and constraints. See [Sec. 11](#) for the precise formulations of problems MOSEK solves.

5.2 “Hello World!” in MOSEK

Here we present the most basic workflow pattern when using Rmosek package.

Create a problem structure

Optimization problems using Rmosek package are specified using a *problem* structure that describes the numerical data of the problem. In most cases it consists of matrices of floating-point numbers.

Retrieving the solutions

When the problem is set up, the optimizer is invoked with the call to *mosek*. The call will return a response and a structure containing the solution to all variables. See further details in [Sec. 7](#).

We refer also to [Sec. 7](#) for information about more advanced mechanisms of interacting with the solver.

Source code example

Below is the most basic code sample that defines and solves a trivial optimization problem

$$\begin{array}{ll}\text{minimize} & x \\ \text{subject to} & 2.0 \leq x \leq 3.0.\end{array}$$

For simplicity the example does not contain any error or status checks.

Listing 5.1: “Hello World!” in MOSEK

```
library("Rmosek")

prob <- list(sense="min")           # Minimization problem
prob$A <- Matrix(nrow=0, ncol=1)   # 0 constraints, 1 variable
prob$bx <- rbind(blx=2.0, bux=3.0) # Bounds on the only variable
prob$c <- c(1.0)                   # The objective coefficient

# Optimize
r <- mosek(prob)

# Print answer
r$sol$itr$xx
```

Chapter 6

Optimization Tutorials

In this section we demonstrate how to set up basic types of optimization problems. Each short tutorial contains a working example of formulating problems, defining variables and constraints and retrieving solutions.

6.1 Linear Optimization

The simplest optimization problem is a purely linear problem. A *linear optimization problem* is a problem of the following form:

Minimize or maximize the objective function

$$\sum_{j=0}^{n-1} c_j x_j + c^f$$

subject to the linear constraints

$$l_k^c \leq \sum_{j=0}^{n-1} a_{kj} x_j \leq u_k^c, \quad k = 0, \dots, m-1,$$

and the bounds

$$l_j^x \leq x_j \leq u_j^x, \quad j = 0, \dots, n-1.$$

The problem description consists of the following elements:

- m and n — the number of constraints and variables, respectively,
- x — the variable vector of length n ,
- c — the coefficient vector of length n

$$c = \begin{bmatrix} c_0 \\ \vdots \\ c_{n-1} \end{bmatrix},$$

- c^f — fixed term in the objective,
- A — an $m \times n$ matrix of coefficients

$$A = \begin{bmatrix} a_{0,0} & \cdots & a_{0,(n-1)} \\ \vdots & \cdots & \vdots \\ a_{(m-1),0} & \cdots & a_{(m-1),(n-1)} \end{bmatrix},$$

- l^c and u^c — the lower and upper bounds on constraints,
- l^x and u^x — the lower and upper bounds on variables.

Please note that we are using 0 as the first index: x_0 is the first element in variable vector x .

6.1.1 Example LO1

The following is an example of a small linear optimization problem:

$$\begin{aligned} & \text{maximize} && 3x_0 &+& 1x_1 &+& 5x_2 &+& 1x_3 \\ & \text{subject to} && 3x_0 &+& 1x_1 &+& 2x_2 && = & 30, \\ & && 2x_0 &+& 1x_1 &+& 3x_2 &+& 1x_3 & \geq & 15, \\ & && && 2x_1 && &+& 3x_3 & \leq & 25, \end{aligned} \tag{6.1}$$

under the bounds

$$\begin{aligned} 0 &\leq x_0 \leq \infty, \\ 0 &\leq x_1 \leq 10, \\ 0 &\leq x_2 \leq \infty, \\ 0 &\leq x_3 \leq \infty. \end{aligned}$$

This is easily programmed in R as shown in [Listing 6.1](#). The first line overwrites any previous definitions of the variable `lo1`, preparing for the new problem description. The problem is then defined and finally solved on the last line.

Listing 6.1: R implementation of problem (6.1).

```
lo1 <- function()
{
  prob <- list()

  # Objective sense (maximize or minimize)
  prob$sense <- "max"

  # Objective coefficients
  prob$c <- c(3, 1, 5, 1)

  # Specify matrix 'A' in sparse format.
  asubi <- c(1, 1, 1, 2, 2, 2, 2, 3, 3)
  asubj <- c(1, 2, 3, 1, 2, 3, 4, 2, 4)
  aval <- c(3, 1, 2, 2, 1, 3, 1, 2, 3)

  prob$A <- sparseMatrix(asubi,asubj,x=aval)

  # Bound values for constraints
  prob$bc <- rbind(blc=c(30, 15, -Inf),
                  buc=c(30, Inf, 25))

  # Bound values for variables
  prob$bx <- rbind(blx=rep(0,4),
                  bux=c(Inf, 10, Inf, Inf))

  # Solve the problem
  r <- mosek(prob)

  # Return the solution
  stopifnot(identical(r$response$code, 0))
  r$sol
}
```

Notice how the R value `Inf` is used in both the constraint bounds (`blc` and `buc`) and the variable upper bound (`bux`), to avoid the specification of an actual bound.

From this example the input arguments for the linear program follows easily.

- **Objective** The string is the objective goal and could be either `minimize`, `min`, `maximize` or `max`. The dense numeric vector specifies the coefficients in front of the variables in the linear objective function, and the optional constant scalar (reads: `c zero`) is a constant in the objective corresponding to c^f , that will be assumed zero if not specified.
- **Constraint Matrix** The sparse matrix is the constraint matrix of the problem with the constraint coefficients written row-wise. Notice that for larger problems it may be more convenient to define an empty sparse matrix and specify the non-zero elements one at a time $A(i, j) = a_{ij}$, rather than writing out the full matrix as done in the example. E.g. `Matrix(0,nrow=30,ncol=50,sparse=TRUE)`.
- **Bounds** The constraint bounds with rows `blc` (constraint lower bound) and `buc` (constraint upper bound), as well as the variable bounds with rows `blx` (variable lower bound) and `bux` (variable upper bound), are both given as dense numeric matrices. These are equivalent to the bounds of problem, namely l^c , u^c , l^x and u^x .

6.2 Quadratic Optimization

MOSEK can solve quadratic and quadratically constrained problems, as long as they are convex. This class of problems can be formulated as follows:

$$\begin{aligned} & \text{minimize} && \frac{1}{2}x^T Q^o x + c^T x + c^f \\ & \text{subject to} && \begin{aligned} l_k^c &\leq \frac{1}{2}x^T Q^k x + \sum_{j=0}^{n-1} a_{k,j} x_j &\leq u_k^c, & k = 0, \dots, m-1, \\ l_j^x &\leq x_j &\leq u_j^x, & j = 0, \dots, n-1. \end{aligned} \end{aligned} \quad (6.2)$$

Without loss of generality it is assumed that Q^o and Q^k are all symmetric because

$$x^T Q x = \frac{1}{2}x^T (Q + Q^T)x.$$

This implies that a non-symmetric Q can be replaced by the symmetric matrix $\frac{1}{2}(Q + Q^T)$.

The problem is required to be convex. More precisely, the matrix Q^o must be positive semi-definite and the k th constraint must be of the form

$$l_k^c \leq \frac{1}{2}x^T Q^k x + \sum_{j=0}^{n-1} a_{k,j} x_j \quad (6.3)$$

with a negative semi-definite Q^k or of the form

$$\frac{1}{2}x^T Q^k x + \sum_{j=0}^{n-1} a_{k,j} x_j \leq u_k^c.$$

with a positive semi-definite Q^k . This implies that quadratic equalities are *not* allowed. Specifying a non-convex problem will result in an error when the optimizer is called.

A matrix is positive semidefinite if all the eigenvalues of Q are nonnegative. An alternative statement of the positive semidefinite requirement is

$$x^T Q x \geq 0, \quad \forall x.$$

If the convexity (i.e. semidefiniteness) conditions are not met **MOSEK** will not produce reliable results or work at all.

6.2.1 Example: Quadratic Objective

We look at a small problem with linear constraints and quadratic objective:

$$\begin{aligned} & \text{minimize} && x_1^2 + 0.1x_2^2 + x_3^2 - x_1x_3 - x_2 \\ & \text{subject to} && \begin{aligned} 1 &\leq x_1 + x_2 + x_3 \\ 0 &\leq x. \end{aligned} \end{aligned} \quad (6.4)$$

The matrix formulation of (6.4) has:

$$Q^o = \begin{bmatrix} 2 & 0 & -1 \\ 0 & 0.2 & 0 \\ -1 & 0 & 2 \end{bmatrix}, c = \begin{bmatrix} 0 \\ -1 \\ 0 \end{bmatrix}, A = \begin{bmatrix} 1 & 1 & 1 \end{bmatrix},$$

with the bounds:

$$l^c = 1, u^c = \infty, l^x = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \text{ and } u^x = \begin{bmatrix} \infty \\ \infty \\ \infty \end{bmatrix}$$

Please note the explicit $\frac{1}{2}$ in the objective function of (6.2) which implies that diagonal elements must be doubled in Q , i.e. $Q_{11} = 2$ even though 1 is the coefficient in front of x_1^2 in (6.4).

Setting up the linear part

The linear parts (constraints, variables, objective) are set up using exactly the same methods as for linear problems, and we refer to Sec. 6.1 for all the details. The same applies to technical aspects such as defining an optimization task, retrieving the solution and so on.

Setting up the quadratic objective

The quadratic objective is specified in a list called *qobj* containing the numerical vectors *i*, *j* and *v*. These vectors specify the (row,column,value)-entries for the lower triangular part of the matrix Q^o using an unordered sparse triplet format. In case of example (6.4) we get:

```
prob$qobj$i <- c(1, 3, 2, 3)
prob$qobj$j <- c(1, 1, 2, 3)
prob$qobj$v <- c(2, -1, 0.2, 2)
```

Please note that

- only non-zero elements are specified (any element not specified is 0 by definition),
- the order of the non-zero elements is insignificant, and
- *only* the lower triangular part should be specified.

Source code

Listing 6.2: Script implementing problem (6.4)

```
##
# Copyright : Copyright (c) MOSEK ApS, Denmark. All rights reserved.
#
# File :      qo1.R
#
# Purpose :   To demonstrate how to solve a small quadratic
#             optimization problem using the Rmosek package.
##
library("Rmosek")

qo1 <- function()
{
  # Specify the non-quadratic part of the problem.
  prob <- list(sense="min")
  prob$c <- c(0, -1, 0)
  prob$A <- Matrix(c(1, 1, 1), nrow=1, sparse=TRUE)
  prob$bc <- rbind(blc=1,
                  buc=Inf)
```

(continues on next page)

```

prob$bx <- rbind(blx=rep(0,3),
                 bux=rep(Inf,3))

# Specify the quadratic objective matrix in triplet form.
prob$qobj$i <- c(1, 3, 2, 3)
prob$qobj$j <- c(1, 1, 2, 3)
prob$qobj$v <- c(2, -1, 0.2, 2)

# Solve the problem
r <- mosek(prob)

# Return the solution
stopifnot(identical(r$response$code, 0))
r$sol
}

qo1()

```

6.2.2 Quadratic constraints

Quadratic constraints are not currently supported in Rmosek. Your options are as follows:

- Use another **MOSEK** interface.
- Use a conic reformulation of the quadratic constraints (see the [Modeling Cookbook](#)).
- Edit the open-source Rmosek interface to add the features you want.

6.3 Conic Quadratic Optimization

Conic optimization is a generalization of linear optimization, allowing constraints of the type

$$x^t \in \mathcal{K}_t,$$

where x^t is a subset of the problem variables and \mathcal{K}_t is a convex cone. Since the set \mathbb{R}^n of real numbers is also a convex cone, we can simply write a compound conic constraint $x \in \mathcal{K}$ where $\mathcal{K} = \mathcal{K}_1 \times \dots \times \mathcal{K}_l$ is a product of smaller cones and x is the full problem variable.

MOSEK can solve conic quadratic optimization problems of the form

$$\begin{aligned}
 & \text{minimize} && c^T x + c^f \\
 & \text{subject to} && l^c \leq Ax \leq u^c, \\
 & && l^x \leq x \leq u^x, \\
 & && x \in \mathcal{K},
 \end{aligned}$$

where the domain restriction, $x \in \mathcal{K}$, implies that all variables are partitioned into convex cones

$$x = (x^0, x^1, \dots, x^{p-1}), \quad \text{with } x^t \in \mathcal{K}_t \subseteq \mathbb{R}^{n_t}.$$

In this tutorial we describe how to use the two types of quadratic cones defined as:

- Quadratic cone:

$$\mathcal{Q}^n = \left\{ x \in \mathbb{R}^n : x_0 \geq \sqrt{\sum_{j=1}^{n-1} x_j^2} \right\}.$$

- Rotated quadratic cone:

$$\mathcal{Q}_r^n = \left\{ x \in \mathbb{R}^n : 2x_0x_1 \geq \sum_{j=2}^{n-1} x_j^2, \quad x_0 \geq 0, \quad x_1 \geq 0 \right\}.$$

For other types of cones supported by **MOSEK** see [Sec. 6.4](#), [Sec. 6.5](#), [Sec. 6.6](#). Different cone types can appear together in one optimization problem.

For example, the following constraint:

$$(x_4, x_0, x_2) \in \mathcal{Q}^3$$

describes a convex cone in \mathbb{R}^3 given by the inequality:

$$x_4 \geq \sqrt{x_0^2 + x_2^2}.$$

Furthermore, each variable may belong to one cone at most. The constraint $x_i - x_j = 0$ would however allow x_i and x_j to belong to different cones with same effect.

6.3.1 Example CQO1

Consider the following conic quadratic problem which involves some linear constraints, a quadratic cone and a rotated quadratic cone.

$$\begin{aligned} \text{minimize} \quad & x_4 + x_5 + x_6 \\ \text{subject to} \quad & x_1 + x_2 + 2x_3 = 1, \\ & x_1, x_2, x_3 \geq 0, \\ & x_4 \geq \sqrt{x_1^2 + x_2^2}, \\ & 2x_5x_6 \geq x_3^2 \end{aligned} \tag{6.5}$$

Setting up the linear part

The linear parts (constraints, variables, objective) are set up using exactly the same methods as for linear problems, and we refer to [Sec. 6.1](#) for all the details. The same applies to technical aspects such as defining an optimization task, retrieving the solution and so on.

Setting up the conic constraints

The conic constraints are specified as columns in a list-typed matrix called *cones*, with rows for each associated detail. In example (6.5) we have two conic constraints:

```
NUMCONES <- 2
prob$cones <- matrix(list(), nrow=2, ncol=NUMCONES)
rownames(prob$cones) <- c("type", "sub")

prob$cones[,1] <- list("QUAD", c(4, 1, 2))
prob$cones[,2] <- list("RQUAD", c(5, 6, 3))
```

The first row selects the type of cone, such as quadratic *"MSK_CT_QUAD"* or rotated quadratic *"MSK_CT_RQUAD"*, noting that the prefix *MSK_CT_* is optional. The second row selects the vector of variables constrained to the cone, identified by index (counting from one in the R language).

Source code

Listing 6.3: Source code solving problem (6.5).

```
library("Rmosek")

cqo1 <- function()
{
  # Specify the non-conic part of the problem.
  prob <- list(sense="min")
  prob$c <- c(0, 0, 0, 1, 1, 1)
  prob$A <- Matrix(c(1, 1, 2, 0, 0, 0), nrow=1, sparse=TRUE)
```

(continues on next page)

(continued from previous page)

```
prob$bc <- rbind(blc=1,
                buc=1)
prob$bx <- rbind(blx=c(rep(0,3), rep(-Inf,3)),
                bux=rep(Inf,6))

# Specify the cones.
NUMCONES <- 2
prob$cones <- matrix(list(), nrow=2, ncol=NUMCONES)
rownames(prob$cones) <- c("type", "sub")

prob$cones[,1] <- list("QUAD", c(4, 1, 2))
prob$cones[,2] <- list("RQUAD", c(5, 6, 3))

#
# Use cbind to extend this chunk of cones if needed:
#
#   oldcones <- prob$cones
#   prob$cones <- cbind(oldcones, newcones)
#

# Solve the problem
r <- mosek(prob)

# Return the solution
stopifnot(identical(r$response$code, 0))
r$sol
}

cqo1()
```

6.4 Power Cone Optimization

Conic optimization is a generalization of linear optimization, allowing constraints of the type

$$x^t \in \mathcal{K}_t,$$

where x^t is a subset of the problem variables and \mathcal{K}_t is a convex cone. Since the set \mathbb{R}^n of real numbers is also a convex cone, we can simply write a compound conic constraint $x \in \mathcal{K}$ where $\mathcal{K} = \mathcal{K}_1 \times \cdots \times \mathcal{K}_l$ is a product of smaller cones and x is the full problem variable.

MOSEK can solve conic optimization problems of the form

$$\begin{aligned} & \text{minimize} && c^T x + c^f \\ & \text{subject to} && l^c \leq Ax \leq u^c, \\ & && l^x \leq x \leq u^x, \\ & && x \in \mathcal{K}, \end{aligned}$$

where the domain restriction, $x \in \mathcal{K}$, implies that all variables are partitioned into convex cones

$$x = (x^0, x^1, \dots, x^{p-1}), \quad \text{with } x^t \in \mathcal{K}_t \subseteq \mathbb{R}^{n_t}.$$

In this tutorial we describe how to use the power cone. The primal power cone of dimension n with parameter $0 < \alpha < 1$ is defined as:

$$\mathcal{P}_n^{\alpha, 1-\alpha} = \left\{ x \in \mathbb{R}^n : x_0^\alpha x_1^{1-\alpha} \geq \sqrt{\sum_{i=2}^{n-1} x_i^2}, \ x_0, x_1 \geq 0 \right\}.$$

In particular, the most important special case is the three-dimensional power cone family:

$$\mathcal{P}_3^{\alpha, 1-\alpha} = \{x \in \mathbb{R}^3 : x_0^\alpha x_1^{1-\alpha} \geq |x_2|, \ x_0, x_1 \geq 0\}.$$

For example, the conic constraint $(x, y, z) \in \mathcal{P}_3^{0.25, 0.75}$ is equivalent to $x^{0.25}y^{0.75} \geq |z|$, or simply $xy^3 \geq z^4$ with $x, y \geq 0$.

MOSEK also supports the dual power cone:

$$(\mathcal{P}_n^{\alpha, 1-\alpha})^* = \left\{ x \in \mathbb{R}^n : \left(\frac{x_0}{\alpha}\right)^\alpha \left(\frac{x_1}{1-\alpha}\right)^{1-\alpha} \geq \sqrt{\sum_{i=2}^{n-1} x_i^2}, x_0, x_1 \geq 0 \right\}.$$

For other types of cones supported by **MOSEK** see [Sec. 6.3](#), [Sec. 6.5](#), [Sec. 6.6](#). Different cone types can appear together in one optimization problem.

Furthermore, each variable may belong to one cone at most. The constraint $x_i - x_j = 0$ would however allow x_i and x_j to belong to different cones with same effect.

6.4.1 Example POW1

Consider the following optimization problem which involves powers of variables:

$$\begin{aligned} & \text{maximize} && x^{0.2}y^{0.8} + z^{0.4} - x \\ & \text{subject to} && x + y + \frac{1}{2}z = 2, \\ & && x, y, z \geq 0. \end{aligned} \tag{6.6}$$

With $(x, y, z) = (x_0, x_1, x_2)$ we convert it into conic form using auxiliary variables as bounds for the power expressions:

$$\begin{aligned} & \text{maximize} && x_3 + x_4 - x_0 \\ & \text{subject to} && x_0 + x_1 + \frac{1}{2}x_2 = 2, \\ & && (x_0, x_1, x_3) \in \mathcal{P}_3^{0.2, 0.8}, \\ & && (x_2, x_5, x_4) \in \mathcal{P}_3^{0.4, 0.6}, \\ & && x_5 = 1. \end{aligned} \tag{6.7}$$

Setting up the linear part

The linear parts (constraints, variables, objective) are set up using exactly the same methods as for linear problems, and we refer to [Sec. 6.1](#) for all the details. The same applies to technical aspects such as defining an optimization task, retrieving the solution and so on.

Setting up the conic constraints

The conic constraints are specified as columns in a list-typed matrix called *cones*, with rows for each associated detail. In example (6.7) we have two conic constraints:

```
NUMCONES <- 2
prob$cones <- matrix(list(), nrow=3, ncol=NUMCONES)
rownames(prob$cones) <- c("type", "sub", "conepar")

prob$cones[,1] <- list("PPOW", c(1, 2, 4), c(0.2, 0.8))
prob$cones[,2] <- list("PPOW", c(3, 6, 5), c(0.4, 0.6))
```

The first row selects the type of cone, i.e., the power cone *"MSK_CT_PPOW"*, noting that the prefix *MSK_CT_* is optional. The second row selects the vector of variables constrained to the cone, identified by index (counting from one in the R language). The third row selects the cone parameters. The existence of this third row is ignored by nonparametric cones (compare, e.g., to [Sec. 6.3.1](#)) and any parameter value can be assigned to non-parametric cones when mixed with parametric ones.

The code below produces the answer of (6.6) which is

```
[ 0.06389298  0.78308564  2.30604283 ]
```

Source code

Listing 6.4: Source code solving problem (6.6).

```
library("Rmosek")

pow1 <- function()
{
  # Specify the non-conic part of the problem.
  prob <- list(sense="max")
  prob$c <- c(-1, 0, 0, 1, 1, 0)
  prob$A <- Matrix(c(1, 1, 0.5, 0, 0, 0), nrow=1, sparse=TRUE)
  prob$bc <- rbind(blc=2,
                   buc=2)
  prob$bx <- rbind(blx=c(rep(-Inf,5), 1),
                   bux=c(rep( Inf,5), 1))

  # Specify the cones.
  NUMCONES <- 2
  prob$cones <- matrix(list(), nrow=3, ncol=NUMCONES)
  rownames(prob$cones) <- c("type", "sub", "conepar")

  prob$cones[,1] <- list("PPOW", c(1, 2, 4), c(0.2, 0.8))
  prob$cones[,2] <- list("PPOW", c(3, 6, 5), c(0.4, 0.6))

  #
  # Non-parametric cones (such as "QUAD") are ignorant to the existence of
  # the third row "conepar" and can be assigned any value (such as NaN).
  #

  # Solve the problem
  r <- mosek(prob)

  # Return the solution
  stopifnot(identical(r$response$code, 0))
  r$sol
}

pow1()
```

6.5 Conic Exponential Optimization

Conic optimization is a generalization of linear optimization, allowing constraints of the type

$$x^t \in \mathcal{K}_t,$$

where x^t is a subset of the problem variables and \mathcal{K}_t is a convex cone. Since the set \mathbb{R}^n of real numbers is also a convex cone, we can simply write a compound conic constraint $x \in \mathcal{K}$ where $\mathcal{K} = \mathcal{K}_1 \times \cdots \times \mathcal{K}_l$ is a product of smaller cones and x is the full problem variable.

MOSEK can solve conic optimization problems of the form

$$\begin{aligned} & \text{minimize} && c^T x + c^f \\ & \text{subject to} && l^c \leq Ax \leq u^c, \\ & && l^x \leq x \leq u^x, \\ & && x \in \mathcal{K}, \end{aligned}$$

where the domain restriction, $x \in \mathcal{K}$, implies that all variables are partitioned into convex cones

$$x = (x^0, x^1, \dots, x^{p-1}), \quad \text{with } x^t \in \mathcal{K}_t \subseteq \mathbb{R}^{n_t}.$$

In this tutorial we describe how to use the primal exponential cone defined as:

$$K_{\text{exp}} = \{x \in \mathbb{R}^3 : x_0 \geq x_1 \exp(x_2/x_1), x_0, x_1 \geq 0\}.$$

MOSEK also supports the dual exponential cone:

$$K_{\text{exp}}^* = \{s \in \mathbb{R}^3 : s_0 \geq -s_2 e^{-1} \exp(s_1/s_2), s_2 \leq 0, s_0 \geq 0\}.$$

For other types of cones supported by **MOSEK** see [Sec. 6.3](#), [Sec. 6.4](#), [Sec. 6.6](#). Different cone types can appear together in one optimization problem.

For example, the following constraint:

$$(x_4, x_0, x_2) \in K_{\text{exp}}$$

describes a convex cone in \mathbb{R}^3 given by the inequalities:

$$x_4 \geq x_0 \exp(x_2/x_0), x_0, x_4 \geq 0.$$

Furthermore, each variable may belong to one cone at most. The constraint $x_i - x_j = 0$ would however allow x_i and x_j to belong to different cones with same effect.

6.5.1 Example CEO1

Consider the following basic conic exponential problem which involves some linear constraints and an exponential inequality:

$$\begin{aligned} & \text{minimize} && x_0 + x_1 \\ & \text{subject to} && x_0 + x_1 + x_2 = 1, \\ & && x_0 \geq x_1 \exp(x_2/x_1), \\ & && x_0, x_1 \geq 0. \end{aligned} \tag{6.8}$$

The conic form of (6.8) is:

$$\begin{aligned} & \text{minimize} && x_0 + x_1 \\ & \text{subject to} && x_0 + x_1 + x_2 = 1, \\ & && (x_0, x_1, x_2) \in K_{\text{exp}}, \\ & && x \in \mathbb{R}^3. \end{aligned} \tag{6.9}$$

Setting up the linear part

The linear parts (constraints, variables, objective) are set up using exactly the same methods as for linear problems, and we refer to [Sec. 6.1](#) for all the details. The same applies to technical aspects such as defining an optimization task, retrieving the solution and so on.

Setting up the conic constraints

The conic constraints are specified as columns in a list-typed matrix called *cones*, with rows for each associated detail. In example (6.9) we have one conic constraint:

```
NUMCONES <- 1
prob$cones <- matrix(list(), nrow=2, ncol=NUMCONES)
rownames(prob$cones) <- c("type", "sub")

prob$cones[,1] <- list("PEXP", c(1, 2, 3))
```

The first row selects the type of cone, in this case the exponential cone *"MSK_CT_PEXP"*, noting that the prefix *MSK_CT_* is optional. The second row selects the vector of variables constrained to the cone, identified by index (counting from one in the R language).

Source code

Listing 6.5: Source code solving problem (6.8).

```
library("Rmosek")

ceo1 <- function()
{
  # Specify the non-conic part of the problem.
  prob <- list(sense="min")
  prob$c <- c(1, 1, 0)
  prob$A <- Matrix(c(1, 1, 1), nrow=1, sparse=TRUE)
  prob$bc <- rbind(blc=1,
                  buc=1)
  prob$bx <- rbind(blx=rep(-Inf,3),
                  bux=rep( Inf,3))

  # Specify the cones.
  NUMCONES <- 1
  prob$cones <- matrix(list(), nrow=2, ncol=NUMCONES)
  rownames(prob$cones) <- c("type", "sub")

  prob$cones[,1] <- list("PEXP", c(1, 2, 3))

  # Solve the problem
  r <- mosek(prob)

  # Return the solution
  stopifnot(identical(r$response$code, 0))
  r$sol
}

ceo1()
```

6.6 Semidefinite Optimization

Semidefinite optimization is a generalization of conic optimization, allowing the use of matrix variables belonging to the convex cone of positive semidefinite matrices

$$\mathcal{S}_+^r = \{X \in \mathcal{S}^r : z^T X z \geq 0, \quad \forall z \in \mathbb{R}^r\},$$

where \mathcal{S}^r is the set of $r \times r$ real-valued symmetric matrices.

MOSEK can solve semidefinite optimization problems of the form

$$\begin{aligned} & \text{minimize} && \sum_{j=0}^{n-1} c_j x_j + \sum_{j=0}^{p-1} \langle \overline{C}_j, \overline{X}_j \rangle + c^f \\ & \text{subject to} && \begin{aligned} l_i^c &\leq \sum_{j=0}^{n-1} a_{ij} x_j + \sum_{j=0}^{p-1} \langle \overline{A}_{ij}, \overline{X}_j \rangle &\leq u_i^c, & i = 0, \dots, m-1, \\ l_j^x &\leq x_j &\leq u_j^x, & j = 0, \dots, n-1, \\ & x \in \mathcal{K}, \overline{X}_j \in \mathcal{S}_+^{r_j}, & & j = 0, \dots, p-1 \end{aligned} \end{aligned}$$

where the problem has p symmetric positive semidefinite variables $\overline{X}_j \in \mathcal{S}_+^{r_j}$ of dimension r_j with symmetric coefficient matrices $\overline{C}_j \in \mathcal{S}^{r_j}$ and $\overline{A}_{i,j} \in \mathcal{S}^{r_j}$. We use standard notation for the matrix inner product, i.e., for $A, B \in \mathbb{R}^{m \times n}$ we have

$$\langle A, B \rangle := \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} A_{ij} B_{ij}.$$

6.6.1 Example SDO1

We consider the simple optimization problem with semidefinite and conic quadratic constraints:

$$\begin{aligned}
 & \text{minimize} && \left\langle \begin{bmatrix} 2 & 1 & 0 \\ 1 & 2 & 1 \\ 0 & 1 & 2 \end{bmatrix}, \bar{X} \right\rangle + x_0 \\
 & \text{subject to} && \left\langle \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \bar{X} \right\rangle + x_0 &= 1, \\
 & && \left\langle \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}, \bar{X} \right\rangle + x_1 + x_2 &= 1/2, \\
 & && x_0 \geq \sqrt{x_1^2 + x_2^2}, & \bar{X} \succeq 0,
 \end{aligned} \tag{6.10}$$

The problem description contains a 3-dimensional symmetric semidefinite variable which can be written explicitly as:

$$\bar{X} = \begin{bmatrix} \bar{X}_{00} & \bar{X}_{10} & \bar{X}_{20} \\ \bar{X}_{10} & \bar{X}_{11} & \bar{X}_{21} \\ \bar{X}_{20} & \bar{X}_{21} & \bar{X}_{22} \end{bmatrix} \in \mathcal{S}_+^3,$$

and a conic quadratic variable $(x_0, x_1, x_2) \in \mathcal{Q}^3$. The objective is to minimize

$$2(\bar{X}_{00} + \bar{X}_{10} + \bar{X}_{11} + \bar{X}_{21} + \bar{X}_{22}) + x_0,$$

subject to the two linear constraints

$$\begin{aligned}
 \bar{X}_{00} + \bar{X}_{11} + \bar{X}_{22} + x_0 &= 1, \\
 \bar{X}_{00} + \bar{X}_{11} + \bar{X}_{22} + 2(\bar{X}_{10} + \bar{X}_{20} + \bar{X}_{21}) + x_1 + x_2 &= 1/2.
 \end{aligned}$$

and can be modeled in R as shown in [Listing 6.6](#).

Listing 6.6: R implementation of model (6.10).

```
##
# Copyright : Copyright (c) MOSEK ApS, Denmark. All rights reserved.
#
# File :      sdo1.R
#
# Purpose :   To solve the mixed semidefinite and conic quadratic optimization problem
#
#              minimize      Tr [2, 1, 0; 1, 2, 1; 0, 1, 2]*X + x(1)
#
#              subject to    Tr [1, 0, 0; 0, 1, 0; 0, 0, 1]*X + x(1)                = 1
#                           Tr [1, 1, 1; 1, 1, 1; 1, 1, 1]*X                + x(2) + x(3) = 0.5
#                           X>=0,  x(1) >= sqrt(x(2)^2 + x(3)^2)
##
library("Rmosek")

getbarvarMatrix <- function(barvar, bardim)
{
  N <- as.integer(bardim)
  new("dspMatrix", x=barvar, uplo="L", Dim=c(N,N))
}

sdo1 <- function()
{
  # Specify the non-matrix variable part of the problem.
  prob <- list(sense="min")
  prob$c <- c(1, 0, 0)
  prob$A <- sparseMatrix(i=c(1, 2, 2),
```

(continues on next page)

```

        j=c(1, 2, 3),
        x=c(1, 1, 1), dims=c(2, 3))
prob$bc    <- rbind(blc=c(1, 0.5),
                   buc=c(1, 0.5))
prob$bx    <- rbind(blx=rep(-Inf,3),
                   bux=rep( Inf,3))
prob$cones <- cbind(list("QUAD", c(1, 2, 3)))

# Specify semidefinite matrix variables (one 3x3 block)
prob$bardim <- c(3)

# Block triplet format specifying the lower triangular part
# of the symmetric coefficient matrix 'barc':
prob$barc$j <- c(1, 1, 1, 1, 1)
prob$barc$k <- c(1, 2, 3, 2, 3)
prob$barc$l <- c(1, 2, 3, 1, 2)
prob$barc$v <- c(2, 2, 2, 1, 1)

# Block triplet format specifying the lower triangular part
# of the symmetric coefficient matrix 'barA':
prob$barA$i <- c(1, 1, 1, 2, 2, 2, 2, 2, 2)
prob$barA$j <- c(1, 1, 1, 1, 1, 1, 1, 1, 1)
prob$barA$k <- c(1, 2, 3, 1, 2, 3, 2, 3, 3)
prob$barA$l <- c(1, 2, 3, 1, 2, 3, 1, 1, 2)
prob$barA$v <- c(1, 1, 1, 1, 1, 1, 1, 1, 1)

# Solve the problem
r <- mosek(prob)

# Print matrix variable and return the solution
stopifnot(identical(r$response$code, 0))
print( list(barx=getbarvarMatrix(r$sol$itr$barx[[1]], prob$bardim[1])) )
r$sol
}

sdo1()

```

6.7 Affine conic constraints (new)

Rmosek package can solve conic optimization problems in another format:

$$\begin{aligned}
 & \text{minimize} && c^T x + c^f \\
 & \text{subject to} && l^c \leq Ax \leq u^c, \\
 & && l^x \leq x \leq u^x, \\
 & && Fx + g \in \mathcal{K},
 \end{aligned}$$

where $F \in \mathbb{R}^{k \times n}$ and $g \in \mathbb{R}^k$ specify an *affine conic constraint* of length (dimension) k . Usually \mathcal{K} will be a product of basic cones corresponding to individual constraints.

In this tutorial we demonstrate how to use the affine conic format. It supports all types of basic cones available in **MOSEK** and can be combined with semidefinite variables as in [Sec. 6.6](#).

6.7.1 Example AFFCO1

Consider the following simple optimization problem:

$$\begin{aligned}
 & \text{maximize} && x_1^{1/3} + (x_1 + x_2 + 0.1)^{1/4} \\
 & \text{subject to} && (x_1 - 0.5)^2 + (x_2 - 0.6)^2 \leq 1, \\
 & && x_1 - x_2 \leq 1.
 \end{aligned} \tag{6.11}$$

Adding auxiliary variables we convert this problem into an equivalent conic form:

$$\begin{aligned}
& \text{maximize} && t_1 + t_2 \\
& \text{subject to} && (1, x_1 - 0.5, x_2 - 0.6) \in \mathcal{Q}^3, \\
& && (x_1, 1, t_1) \in \mathcal{P}_3^{1/3, 2/3}, \\
& && (x_1 + x_2 + 0.1, 1, t_2) \in \mathcal{P}_3^{1/4, 3/4}, \\
& && x_1 - x_2 \leq 1.
\end{aligned} \tag{6.12}$$

Note that each of the vectors constrained to a cone is in a natural way an affine combination of the problem variables.

We first set up the linear part of the problem, including the number of variables, objective and all bounds precisely as in Sec. 6.1. Cones will be defined using the field K in the `problem` structure. We construct the matrices F, g for each of the three cones. For example, the constraint $(1, x_1 - 0.5, x_2 - 0.6) \in \mathcal{Q}^3$ is written in matrix form as

$$\begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ t_1 \\ t_2 \end{bmatrix} + \begin{bmatrix} 1 \\ -0.5 \\ -0.6 \end{bmatrix} \in \mathcal{Q}^3.$$

Below we set up the matrices and define the cone type as a quadratic cone of length 3. The last cone parameter is left empty, but must be included because we will use cone parameters for the power cones.

```
# The quadratic cone
FQ <- rbind(c(0,0,0,0), c(1,0,0,0), c(0,1,0,0))
gQ <- c(1, -0.5, -0.6)
cQ <- matrix(list("QUAD", 3, NULL), nrow=3, ncol=1)
```

Next we demonstrate how to do the same for the second of the power cone constraints. Its affine representation is:

$$\begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ t_1 \\ t_2 \end{bmatrix} + \begin{bmatrix} 0.1 \\ 1 \\ 0 \end{bmatrix} \in \mathcal{P}_3^{1/4, 3/4}.$$

The power cone is defined by its type, length, and the additional parameter which is a vector of exponents $\alpha, 1-\alpha$ appearing in the cone definition. In fact any pair of positive real numbers *proportional* to $(\alpha, 1-\alpha)$ may be used. They will be normalized to add up to 1:

```
# The power cone for (x_1+x_2+0.1, 1, t_2) \in POW3^(1/4, 3/4)
FP2 <- rbind(c(1,1,0,0), c(0,0,0,0), c(0,0,0,1))
gP2 <- c(0.1, 1, 0)
cP2 <- matrix(list("POW", 3, c(1.0, 3.0)), nrow=3, ncol=1)
```

Once affine conic descriptions of all cones are ready it remains to stack them vertically into the matrix F and vector g and concatenate the cone descriptions in one matrix, one column per cone. Below is the full code for problem (6.12).

Listing 6.7: Script implementing conic version of problem (6.11).

```
library("Rmosek")

affco1 <- function()
{
  prob <- list(sense="max")

  # Variables [x1; x2; t1; t2]
  prob$c <- c(0, 0, 1, 1)

  # Linear inequality x_1 - x_2 <= 1
```

(continues on next page)

```

prob$A <- Matrix(c(1, -1, 0, 0), nrow=1, sparse=TRUE)
prob$bc <- rbind(blc=-Inf, buc=1)
prob$bx <- rbind(blx=rep(-Inf,4), bux=rep(Inf,4))

# The quadratic cone
FQ <- rbind(c(0,0,0,0), c(1,0,0,0), c(0,1,0,0))
gQ <- c(1, -0.5, -0.6)
cQ <- matrix(list("QUAD", 3, NULL), nrow=3, ncol=1)

# The power cone for (x_1, 1, t_1) \in POW3^(1/3,2/3)
FP1 <- rbind(c(1,0,0,0), c(0,0,0,0), c(0,0,1,0))
gP1 <- c(0, 1, 0)
cP1 <- matrix(list("PPOW", 3, c(1/3, 2/3)), nrow=3, ncol=1)

# The power cone for (x_1+x_2+0.1, 1, t_2) \in POW3^(1/4,3/4)
FP2 <- rbind(c(1,1,0,0), c(0,0,0,0), c(0,0,0,1))
gP2 <- c(0.1, 1, 0)
cP2 <- matrix(list("PPOW", 3, c(1.0, 3.0)), nrow=3, ncol=1)

# All cones
prob$F <- rbind(FQ, FP1, FP2)
prob$g <- cbind(gQ, gP1, gP2)
prob$cones <- cbind(cQ, cP1, cP2)
rownames(prob$cones) <- c("type", "dim", "conepar")

r <- mosek(prob, list(soldetail=1))
stopifnot(identical(r$response$code, 0))

print(r$sol$itr$pobjval)
print(r$sol$itr$xx[1:2])
}

```

6.7.2 Example AFFCO2

Consider the following simple linear dynamical system. A point in \mathbb{R}^n moves along a trajectory given by $z(t) = z(0) \exp(At)$, where $z(0)$ is the starting position and $A = \mathbf{Diag}(a_1, \dots, a_n)$ is a diagonal matrix with $a_i < 0$. Find the time after which $z(t)$ is within euclidean distance d from the origin. Denoting the coordinates of the starting point by $z(0) = (z_1, \dots, z_n)$ we can write this as an optimization problem in one variable t :

$$\begin{aligned} & \text{minimize} && t \\ & \text{subject to} && \sqrt{\sum_i (z_i \exp(a_i t))^2} \leq d, \end{aligned}$$

which can be cast into conic form as:

$$\begin{aligned} & \text{minimize} && t \\ & \text{subject to} && (d, z_1 y_1, \dots, z_n y_n) \in \mathcal{Q}^{n+1}, \\ & && (y_i, 1, a_i t) \in K_{\exp}, \quad i = 1, \dots, n, \end{aligned} \tag{6.13}$$

with variable vector $x = [t, y_1, \dots, y_n]^T$.

We assemble all conic constraints in the form

$$Fx + g \in \mathcal{Q}^{n+1} \times (K_{\exp})^n.$$

For the conic quadratic constraint this representation is

$$\begin{bmatrix} 0 & 0_n^T \\ 0_n & \mathbf{Diag}(z_1, \dots, z_n) \end{bmatrix} \begin{bmatrix} t \\ y \end{bmatrix} + \begin{bmatrix} d \\ 0_n \end{bmatrix} \in \mathcal{Q}^{n+1}.$$

For the i -th exponential cone we have

$$\begin{bmatrix} 0 & e_i^T \\ 0 & 0_n \\ a_i & 0_n \end{bmatrix} \begin{bmatrix} t \\ y \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \in K_{\exp},$$

where e_i denotes a vector of length n with a single 1 in position i .

Listing 6.8: Script implementing problem (6.13).

```
firstHittingTime <- function(n, z, a, d)
{
  prob <- list(sense="min")
  # Variables [t, y1, ..., yn]
  prob$A <- Matrix(nrow=0, ncol=n+1)
  prob$bx <- rbind(rep(-Inf, n+1), rep(Inf, n+1))
  prob$c <- c(1, rep(0, n))

  # Quadratic cone
  FQ <- Diagonal(n+1, c(0, z))
  gQ <- c(d, rep(0, n))

  # All exponential cones
  FE <- sparseMatrix( c(seq(1, 3*n, by=3), seq(3, 3*n, by=3)),
                     c(2:(n+1), rep(1, n)),
                     x=c(rep(1, n), t(a)),
                     y=c(0, 1, 0), n)
  gE <- rep(c(0, 1, 0), n)

  # Assemble input data
  prob$F <- rbind(FQ, FE)
  prob$g <- c(gQ, gE)
  prob$cones <- cbind(matrix(list("QUAD", 1+n, NULL), nrow=3, ncol=1),
                     matrix(list("PEXP", 3, NULL), nrow=3, ncol=n))
  rownames(prob$cones) <- c("type", "dim", "conepar")

  # Solve
  r <- mosek(prob, list(soldetail=1))
  stopifnot(identical(r$response$code, 0))

  r$sol$itr$xx[1]
}
```

6.8 Geometric Programming

Geometric programs (GP) are a particular class of optimization problems which can be expressed in special polynomial form as positive sums of generalized monomials. More precisely, a geometric problem in canonical form is

$$\begin{aligned} & \text{minimize} && f_0(x) \\ & \text{subject to} && f_i(x) \leq 1, \quad i = 1, \dots, m, \\ & && x_j > 0, \quad j = 1, \dots, n, \end{aligned} \tag{6.14}$$

where each f_0, \dots, f_m is a *posynomial*, that is a function of the form

$$f(x) = \sum_k c_k x_1^{\alpha_{k1}} x_2^{\alpha_{k2}} \dots x_n^{\alpha_{kn}}$$

with arbitrary real α_{ki} and $c_k > 0$. The standard way to formulate GPs in convex form is to introduce a variable substitution

$$x_i = \exp(y_i).$$

Under this substitution all constraints in a GP can be reduced to the form

$$\log\left(\sum_k \exp(a_k^T y + b_k)\right) \leq 0 \tag{6.15}$$

involving a *log-sum-exp* bound. Moreover, constraints involving only a single monomial in x can be even more simply written as a linear inequality:

$$a_k^T y + b_k \leq 0$$

We refer to the **MOSEK Modeling Cookbook** and to [\[BKVH07\]](#) for more details on this reformulation. A geometric problem formulated in convex form can be entered into **MOSEK** with the help of exponential cones.

6.8.1 Example GP1

The following problem comes from [\[BKVH07\]](#). Consider maximizing the volume of a $h \times w \times d$ box subject to upper bounds on the area of the floor and of the walls and bounds on the ratios h/w and d/w :

$$\begin{aligned} & \text{maximize} && hwd \\ & \text{subject to} && 2(hw + hd) \leq A_{\text{wall}}, \\ & && wd \leq A_{\text{floor}}, \\ & && \alpha \leq h/w \leq \beta, \\ & && \gamma \leq d/w \leq \delta. \end{aligned} \tag{6.16}$$

The decision variables in the problem are h, w, d . We make a substitution

$$h = \exp(x), w = \exp(y), d = \exp(z)$$

after which (6.16) becomes

$$\begin{aligned} & \text{maximize} && x + y + z \\ & \text{subject to} && \log(\exp(x + y + \log(2/A_{\text{wall}})) + \exp(x + z + \log(2/A_{\text{wall}}))) \leq 0, \\ & && y + z \leq \log(A_{\text{floor}}), \\ & && \log(\alpha) \leq x - y \leq \log(\beta), \\ & && \log(\gamma) \leq z - y \leq \log(\delta). \end{aligned} \tag{6.17}$$

Next, we demonstrate how to implement a log-sum-exp constraint (6.15). It can be written as:

$$u_k \geq \exp(a_k^T y + b_k), \quad (\text{equiv. } (u_k, 1, a_k^T y + b_k) \in K_{\text{exp}}), \tag{6.18}$$

$$\sum_k u_k = 1.$$

This presentation requires one extra variable u_k for each monomial appearing in the original posynomial constraint. It is natural to express the cone membership using an affine conic constraint (see [Sec. 6.7](#)). In this case:

$$\begin{bmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ u_1 \\ u_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ \log(2/A_{\text{wall}}) \\ 0 \\ 1 \\ \log(2/A_{\text{wall}}) \end{bmatrix} \in K_{\text{exp}} \times K_{\text{exp}}.$$

We can now use this function to assemble all constraints in the model. The linear part of the problem is entered as in [Sec. 6.1](#).

Listing 6.9: Source code solving problem (6.17).

```
# Input data
Awall <- 200.0
Afloor <- 50.0
alpha <- 2.0
beta <- 10
gamma <- 2
delta <- 10

# Objective
prob <- list(sense="max")
prob$c <- c(1, 1, 1, 0, 0)

# Linear constraints:
```

(continues on next page)

```

# [ 0  0  0  1  1 ]      == 1
# [ 0  1  1  0  0 ]      <= log(Afloor)
# [ 1 -1  0  0  0 ]      in [log(alpha), log(beta)]
# [ 0 -1  1  0  0 ]      in [log(gamma), log(delta)]
#
prob$A <- rbind(c(0, 0, 0, 1, 1),
               c(0, 1, 1, 0, 0),
               c(1,-1, 0, 0, 0),
               c(0,-1, 1, 0, 0))

prob$bc <- rbind(c(1, -Inf,      log(alpha), log(gamma)),
               c(1, log(Afloor), log(beta), log(delta)))

prob$bx <- rbind(c(-Inf, -Inf, -Inf, -Inf, -Inf),
               c( Inf,  Inf,  Inf,  Inf,  Inf))

# The conic part FX+g \in Kexp x Kexp
#   x  y  z  u  v
# [ 0  0  0  1  0 ]    0
# [ 0  0  0  0  0 ]    1      in Kexp
# [ 1  1  0  0  0 ]    log(2/Awall)
#
# [ 0  0  0  0  1 ]    0
# [ 0  0  0  0  0 ]    1      in Kexp
# [ 1  0  1  0  0 ] + log(2/Awall)
#
#
prob$F <- rbind(c(0, 0, 0, 1, 0),
               c(0, 0, 0, 0, 0),
               c(1, 1, 0, 0, 0),
               c(0, 0, 0, 0, 1),
               c(0, 0, 0, 0, 0),
               c(1, 0, 1, 0, 0))

prob$g <- c(0, 1, log(2/Awall), 0, 1, log(2/Awall))

prob$cones <- matrix(rep(list("PEXP", 3, NULL), 2), ncol=2)
rownames(prob$cones) <- c("type", "dim", "conepar")

# Optimize and print results
r <- mosek(prob, list(soldetail=1))
print(exp(r$sol$itr$xx[1:3]))

```

6.9 Integer Optimization

An optimization problem where one or more of the variables are constrained to integer values is called a (mixed) integer optimization problem. **MOSEK** supports integer variables in combination with linear, quadratic and quadratically constrained and conic problems (except semidefinite). See the previous tutorials for an introduction to how to model these types of problems.

6.9.1 Example MILO1

We use the example

$$\begin{aligned}
 &\text{maximize} && x_0 + 0.64x_1 \\
 &\text{subject to} && 50x_0 + 31x_1 \leq 250, \\
 & && 3x_0 - 2x_1 \geq -4, \\
 & && x_0, x_1 \geq 0 \quad \text{and integer}
 \end{aligned} \tag{6.19}$$

to demonstrate how to set up and solve a problem with integer variables. It has the structure of a linear optimization problem (see Sec. 6.1) except for integrality constraints on the variables. Therefore, only

the specification of the integer constraints requires something new compared to the linear optimization problem discussed previously.

This is easily programmed in R using the piece code shown in Listing 6.10,

Listing 6.10: R implementation of problem (6.19).

```
##
# Copyright : Copyright (c) MOSEK ApS, Denmark. All rights reserved.
#
# File :      milo1.R
#
# Purpose :   To demonstrate how to solve a small mixed integer linear
#             optimization problem using the Rmosek package.
##
library("Rmosek")

milo1 <- function()
{
  # Specify the continuous part of the problem.
  prob <- list(sense="max")
  prob$c <- c(1, 0.64)
  prob$A <- Matrix(rbind(c(50, 31),
                        c( 3, -2)), sparse=TRUE)
  prob$bc <- rbind(blc=c(-Inf, -4),
                  buc=c( 250, Inf))
  prob$bx <- rbind(blx=c( 0, 0),
                  bux=c(Inf, Inf))

  # Specify the integer constraints
  prob$intsub <- c(1, 2)

  # Solve the problem
  r <- mosek(prob)

  # Return the solution
  stopifnot(identical(r$response$code, 0))
  r$sol
}

milo1()
```

where x_1 and x_2 are pointed out as integer variables.

The input arguments follow those of a linear or conic program with the additional identification of the integer variables. The column vector `intsub` should simply contain indexes to the subset of variables for which integrality is required. For instance if x should be a binary $\{0, 1\}$ -variable, its index in the problem formulation should be added to `intsub`, and its bounds $0 \leq x \leq 1$ should be specified explicitly.

If executed correctly you should be able to see the log of the interface and optimization process printed to the screen. The output structure will only include an integer solution `int`, since we are no longer in the continuous domain for which the interior-point algorithm operates. The structure also contains the problem status as well as the solution status based on certificates found by the **MOSEK** optimization library.

6.9.2 Specifying an initial solution

It is a common strategy to provide a starting feasible point (if one is known in advance) to the mixed-integer solver. This can in many cases reduce solution time.

It is not necessary to specify the whole solution. **MOSEK** will attempt to use it to speed up the computation. **MOSEK** will first try to construct a feasible solution by fixing integer variables to the values provided by the user (rounding if necessary) and optimizing over the continuous variables. The outcome of this process can be inspected via information items `"MSK_IINF_MIO_CONSTRUCT_SOLUTION"` and `"MSK_DINF_MIO_CONSTRUCT_SOLUTION_OBJ"`, and via the `Construct solution` objective entry

in the log. We concentrate on a simple example below.

$$\begin{aligned}
& \text{maximize} && 7x_0 + 10x_1 + x_2 + 5x_3 \\
& \text{subject to} && x_0 + x_1 + x_2 + x_3 \leq 2.5 \\
& && x_0, x_1, x_2 \in \mathbb{Z} \\
& && x_0, x_1, x_2, x_3 \geq 0
\end{aligned} \tag{6.20}$$

Solution values can be set using the appropriate fields in the problem structure.

Listing 6.11: Implementation of problem (6.20) specifying an initial solution.

```
# Specify start guess for the integer variables.
prob$sol$int$xx <- c(1, 1, 0, NaN)
```

The log output from the optimizer will in this case indicate that the inputted values were used to construct an initial feasible solution:

```
Construct solution objective      : 1.950000000000e+01
```

The same information can be obtained from the API:

Listing 6.12: Retrieving information about usage of initial solution

```
print(r$info$MIO_CONSTRUCT_SOLUTION)
print(r$dinfo$MIO_CONSTRUCT_SOLUTION_OBJ)
```

6.9.3 Example MICO1

Integer variables can also be used arbitrarily in conic problems (except semidefinite). We refer to the previous tutorials for how to set up a conic optimization problem. Here we present sample code that sets up a simple optimization problem:

$$\begin{aligned}
& \text{minimize} && x^2 + y^2 \\
& \text{subject to} && x \geq e^y + 3.8, \\
& && x, y \text{ integer.}
\end{aligned} \tag{6.21}$$

The canonical conic formulation of (6.21) suitable for Rmosek package is

$$\begin{aligned}
& \text{minimize} && t \\
& \text{subject to} && (t, x, y) \in \mathcal{Q}^3 && (t \geq \sqrt{x^2 + y^2}) \\
& && (x - 3.8, 1, y) \in K_{\text{exp}} && (x - 3.8 \geq e^y) \\
& && x, y \text{ integer,} \\
& && t \in \mathbb{R}.
\end{aligned} \tag{6.22}$$

Listing 6.13: Implementation of problem (6.22).

```
library("Rmosek")

micol <- function()
{
  # Specify the continuous part of the problem.
  # Variables are [t; x; y]
  prob <- list(sense="min")
  prob$c <- c(1, 0, 0)
  prob$A <- Matrix(nrow=0, ncol=3) # 0 constraints, 3 variables
  prob$bx <- rbind(blx=rep(-Inf,3), bux=rep(Inf,3))

  # Conic part of the problem
  prob$F <- rbind(diag(3), c(0,1,0), c(0,0,0), c(0,0,1))
  prob$g <- c(0, 0, 0, -3.8, 1, 0)
```

(continues on next page)

```

prob$cones <- cbind(matrix(list("QUAD", 3, NULL), nrow=3, ncol=1),
                     matrix(list("PEXP", 3, NULL), nrow=3, ncol=1))
rownames(prob$cones) <- c("type", "dim", "conepar")

# Specify the integer constraints
prob$intsub <- c(2, 3)

# Solve the problem
r <- mosek(prob)

# Return the solution
stopifnot(identical(r$response$code, 0))
print(r$sol$int$xx[2:3])
}

```

Note that the conic constraints are described using the format $Fx + g \in \mathcal{K}$, that is as *affine conic constraints*. See Sec. 6.7 for details.

Error and solution status handling were omitted for readability.

6.10 Problem Modification and Reoptimization

Often one might want to solve not just a single optimization problem, but a sequence of problems, each differing only slightly from the previous one. This section demonstrates how to modify and re-optimize an existing problem. The example we study is a simple production planning model.

Problem modifications regarding variables, cones, objective function and constraints can be grouped in categories:

- add/remove,
- coefficient modifications,
- bounds modifications.

Especially removing variables and constraints can be costly. Special care must be taken with respect to constraints and variable indexes that may be invalidated.

Depending on the type of modification, **MOSEK** may be able to optimize the modified problem more efficiently exploiting the information and internal state from the previous execution. After optimization, the solution is always stored internally, and is available before next optimization. The former optimal solution may be still feasible, but no longer optimal; or it may remain optimal if the modification of the objective function was small.

In general, **MOSEK** exploits dual information and availability of an optimal basis from the previous execution. The simplex optimizer is well suited for exploiting an existing primal or dual feasible solution. Restarting capabilities for interior-point methods are still not as reliable and effective as those for the simplex algorithm. More information can be found in Chapter 10 of the book [Chv83].

Parameter settings (see Sec. 7.4) can also be changed between optimizations.

6.10.1 Example: Production Planning

A company manufactures three types of products. Suppose the stages of manufacturing can be split into three parts: Assembly, Polishing and Packing. In the table below we show the time required for each stage as well as the profit associated with each product.

Product no.	Assembly (minutes)	Polishing (minutes)	Packing (minutes)	Profit (\$)
0	2	3	2	1.50
1	4	2	3	2.50
2	3	3	2	3.00

With the current resources available, the company has 100,000 minutes of assembly time, 50,000 minutes of polishing time and 60,000 minutes of packing time available per year. We want to know how many items of each product the company should produce each year in order to maximize profit?

Denoting the number of items of each type by x_0, x_1 and x_2 , this problem can be formulated as a linear optimization problem:

$$\begin{aligned} & \text{maximize} && 1.5x_0 &+& 2.5x_1 &+& 3.0x_2 \\ & \text{subject to} && 2x_0 &+& 4x_1 &+& 3x_2 &\leq 100000, \\ & && 3x_0 &+& 2x_1 &+& 3x_2 &\leq 50000, \\ & && 2x_0 &+& 3x_1 &+& 2x_2 &\leq 60000, \end{aligned} \tag{6.23}$$

and

$$x_0, x_1, x_2 \geq 0.$$

Code in Listing 6.14 loads and solves this problem.

Listing 6.14: Setting up and solving problem (6.23)

```
# Specify the c vector.
prob  <- list(sense="max")
prob$c <- c(1.5, 2.5, 3.0)

# Specify a in sparse format.
subi  <- c(1, 1, 1, 2, 2, 2, 3, 3, 3)
subj  <- c(1, 2, 3, 1, 2, 3, 1, 2, 3)
valij <- c(2, 4, 3, 3, 2, 3, 2, 3, 2)

prob$A <- sparseMatrix(subi,subj,x=valij);

# Specify bounds of the constraints.
prob$bc<- rbind(blc=rep(-Inf, 3),
                buc=c(100000, 50000, 60000))

# Specify bounds of the variables.
prob$bx<- rbind(blx=rep(0,3),
                bux=rep(Inf,3))

# Perform the optimization.
prob$iparam <- list(OPTIMIZER="OPTIMIZER_DUAL_SIMPLEX")
r <- mosek(prob)

# Show the optimal x solution.
print(r$sol$bas$xx)
```

6.10.2 Changing the Linear Constraint Matrix

Suppose we want to change the time required for assembly of product 0 to 3 minutes. This corresponds to setting $a_{0,0} = 3$, which is done by directly modifying the A matrix of the problem, as shown below.

```
prob$A[1,1] <- 3.0
```

The problem now has the form:

$$\begin{aligned} & \text{maximize} && 1.5x_0 &+& 2.5x_1 &+& 3.0x_2 \\ & \text{subject to} && 3x_0 &+& 4x_1 &+& 3x_2 &\leq 100000, \\ & && 3x_0 &+& 2x_1 &+& 3x_2 &\leq 50000, \\ & && 2x_0 &+& 3x_1 &+& 2x_2 &\leq 60000, \end{aligned} \tag{6.24}$$

and

$$x_0, x_1, x_2 \geq 0.$$

After this operation we can reoptimize the problem.

6.10.3 Appending Variables

We now want to add a new product with the following data:

Product no.	Assembly (minutes)	Polishing (minutes)	Packing (minutes)	Profit (\$)
3	4	0	1	1.00

This corresponds to creating a new variable x_3 , appending a new column to the A matrix and setting a new term in the objective. We do this in [Listing 6.15](#)

Listing 6.15: How to add a new variable (column)

```
prob$c      <- c(prob$c, 1.0)
prob$A      <- cbind(prob$A, c(4.0, 0.0, 1.0))
prob$bx     <- cbind(prob$bx, c(0.0, Inf))
```

After this operation the new problem is:

$$\begin{aligned}
 &\text{maximize} && 1.5x_0 + 2.5x_1 + 3.0x_2 + 1.0x_3 \\
 &\text{subject to} && 3x_0 + 4x_1 + 3x_2 + 4x_3 \leq 100000, \\
 & && 3x_0 + 2x_1 + 3x_2 \leq 50000, \\
 & && 2x_0 + 3x_1 + 2x_2 + 1x_3 \leq 60000,
 \end{aligned} \tag{6.25}$$

and

$$x_0, x_1, x_2, x_3 \geq 0.$$

6.10.4 Appending Constraints

Now suppose we want to add a new stage to the production process called *Quality control* for which 30000 minutes are available. The time requirement for this stage is shown below:

Product no.	Quality control (minutes)
0	1
1	2
2	1
3	1

This corresponds to adding the constraint

$$x_0 + 2x_1 + x_2 + x_3 \leq 30000$$

to the problem. This is done as follows.

Listing 6.16: Adding a new constraint.

```
prob$A      <- rbind(prob$A, c(1.0, 2.0, 1.0, 1.0))
prob$bc     <- cbind(prob$bc, c(-Inf, 30000.0))
```

Again, we can continue with re-optimizing the modified problem.

6.10.5 Changing bounds

One typical reoptimization scenario is to change bounds. Suppose for instance that we must operate with limited time resources, and we must change the upper bounds in the problem as follows:

Operation	Time available (before)	Time available (new)
Assembly	100000	80000
Polishing	50000	40000
Packing	60000	50000
Quality control	30000	22000

That means we would like to solve the problem:

$$\begin{aligned}
 &\text{maximize} && 1.5x_0 &+& 2.5x_1 &+& 3.0x_2 &+& 1.0x_3 \\
 &\text{subject to} && 3x_0 &+& 4x_1 &+& 3x_2 &+& 4x_3 &\leq 80000, \\
 &&& 3x_0 &+& 2x_1 &+& 3x_2 && &\leq 40000, \\
 &&& 2x_0 &+& 3x_1 &+& 2x_2 &+& 1x_3 &\leq 50000, \\
 &&& x_0 &+& 2x_1 &+& x_2 &+& x_3 &\leq 22000.
 \end{aligned} \tag{6.26}$$

In this case all we need to do is redefine the upper bound vector for the constraints, as shown in the next listing.

Listing 6.17: Change constraint bounds.

```

prob$bc<- rbind(blc=rep(-Inf, 4),
               buc=c(80000, 40000, 50000, 22000))
r <- mosek(prob)
print(r$sol$bas$xx)

```

Again, we can continue with re-optimizing the modified problem.

6.10.6 Advanced hot-start

In order to exploit the possibility of hot-starting the simplex algorithms it is necessary to pass the old basic solution when the modified problem is re-optimized. Without this operation the optimizer will simply start from scratch. Any subset of the basic solution may be provided, but to achieve the best results all fields of `res.sol.bas` should be present, that is `xx,xc,slx,sux,slc,suc,skx,skc`.

Listing 6.18: Passing the full basic solution.

```

# Reoptimize with changed coefficient
# Use previous solution to perform very simple hot-start.
# This part can be skipped, but then the optimizer will start
# from scratch on the new problem, i.e. without any hot-start.
prob$sol <- list(bas=r$sol$bas)
r <- mosek(prob)
print(r$sol$bas$xx)

```

If the dimensions of the problem (number of variables, constraints) have changed, the lengths of all fields have to be adjusted to be compatible with the reformulated problem. For example, here is an adjustment when adding a new variable:

Listing 6.19: Adjusting lengths in the solution fields related to variables.

```

# Reoptimize with a new variable and hot-start
# All parts of the solution must be extended to the new dimensions.
prob$sol <- list(bas=r$sol$bas)
prob$sol$bas$xx <- c(prob$sol$bas$xx, 0.0)
prob$sol$bas$slx <- c(prob$sol$bas$slx, 0.0)
prob$sol$bas$sux <- c(prob$sol$bas$sux, 0.0)
prob$sol$bas$skx <- c(prob$sol$bas$skx, "UN")
r <- mosek(prob)
print(r$sol$bas$xx)

```

If the optimizer used the data from the previous run to hot-start the optimizer for reoptimization, this will be indicated in the log:

```

Optimizer - hotstart          : yes

```

When performing re-optimizations, instead of removing a basic variable it may be more efficient to fix the variable at zero and then remove it when the problem is re-optimized and it has left the basis. This makes it easier for **MOSEK** to restart the simplex optimizer.

For a more in-depth treatment see the following sections:

- *Case studies* for more advanced and complicated optimization examples.

- *Problem Formulation and Solutions* for formal mathematical formulations of problems **MOSEK** can solve, dual problems and infeasibility certificates.

Chapter 7

Solver Interaction Tutorials

In this section we cover the interaction with the solver.

7.1 Accessing the solution

This section contains important information about the status of the solver and the status of the solution, which must be checked in order to properly interpret the results of the optimization.

7.1.1 Solver termination

The optimizer provides a **response code** of type *rescode*, relevant for error handling. It indicates if any errors occurred in any phase of optimization (including processing input data). It will also indicate system-related errors (such as an out of memory error, licensing error etc.). Finally, it will also indicate if the optimizer terminated correctly, but for a non-standard reason, for example because it reached a time limit or met another criterion set by the user. Such termination codes are not errors. The expected value for a typical successful optimization without any special settings is *"MSK_RES_OK"*.

If a runtime error causes the program to crash during optimization, the first debugging step is to enable logging and check the log output. See [Sec. 7.3](#).

If the optimization completes successfully, the next step is to check the solution status, as explained below.

7.1.2 Available solutions

MOSEK uses three kinds of optimizers and provides three types of solutions:

- **basic solution** from the simplex optimizer,
- **interior-point solution** from the interior-point optimizer,
- **integer solution** from the mixed-integer optimizer.

Under standard parameters settings the following solutions will be available for various problem types:

Table 7.1: Types of solutions available from **MOSEK**

	Simplex optimizer	Interior-point optimizer	Mixed-integer optimizer
Linear problem	<code>r\$sol\$bas</code>	<code>r\$sol\$itr</code>	
Nonlinear continuous problem		<code>r\$sol\$itr</code>	
Problem with integer variables			<code>r\$sol\$int</code>

For linear problems the user can force a specific optimizer choice making only one of the two solutions available. For example, if the user disables basis identification, then only the interior point solution will be available for a linear problem. Numerical issues may cause one of the solutions to be unknown even if another one is feasible.

Not all components of a solution are always available. For example, there is no dual solution for integer problems and no dual conic variables from the simplex optimizer.

The user will always need to specify which solution should be accessed.

7.1.3 Problem and solution status

Assuming that the optimization terminated without errors, the next important step is to check the problem and solution status. There is one for every type of solution, as explained above.

Problem status

Problem status (*prosta*) determines whether the problem is certified as feasible. Its values can roughly be divided into the following broad categories:

- **feasible** — the problem is feasible. For continuous problems and when the solver is run with default parameters, the feasibility status should ideally be *"MSK_PRO_STA_PRIM_AND_DUAL_FEAS"*.
- **primal/dual infeasible** — the problem is infeasible or unbounded or a combination of those. The exact problem status will indicate the type of infeasibility.
- **unknown** — the solver was unable to reach a conclusion, most likely due to numerical issues.

Solution status

Solution status (*solsta*) provides the information about what the solution values actually contain. The most important broad categories of values are:

- **optimal** (*"MSK_SOL_STA_OPTIMAL"*) — the solution values are feasible and optimal.
- **certificate** — the solution is in fact a certificate of infeasibility (primal or dual, depending on the solution).
- **unknown/undefined** — the solver could not solve the problem or this type of solution is not available for a given problem.

Problem and solution status can be found in the fields *prosta* and *solsta* of a solution structure *solution_info*, for instance *r\$sol\$itr\$prosta*, *r\$sol\$itr\$solsta* for the interior-point solution.

The solution status determines the action to be taken. For example, in some cases a suboptimal solution may still be valuable and deserve attention. It is the user's responsibility to check the status and quality of the solution.

Typical status reports

Here are the most typical optimization outcomes described in terms of the problem and solution statuses. Note that these do not cover all possible situations that can occur.

Table 7.2: Continuous problems (solution status for interior-point and basic solution)

Outcome	Problem status	Solution status
Optimal	<i>"MSK_PRO_STA_PRIM_AND_DUAL_FEAS"</i>	<i>"MSK_SOL_STA_OPTIMAL"</i>
Primal infeasible	<i>"MSK_PRO_STA_PRIM_INFEAS"</i>	<i>"MSK_SOL_STA_PRIM_INFEAS_CERT"</i>
Dual infeasible (unbounded)	<i>"MSK_PRO_STA_DUAL_INFEAS"</i>	<i>"MSK_SOL_STA_DUAL_INFEAS_CERT"</i>
Uncertain (stall, numerical issues, etc.)	<i>"MSK_PRO_STA_UNKNOWN"</i>	<i>"MSK_SOL_STA_UNKNOWN"</i>

Table 7.3: Integer problems (solution status for integer solution, others undefined)

Outcome	Problem status	Solution status
Integer optimal	<i>"MSK_PRO_STA_PRIM_FEAS"</i>	<i>"MSK_SOL_STA_INTEGER_OPTIMAL"</i>
Infeasible	<i>"MSK_PRO_STA_PRIM_INFEAS"</i>	<i>"MSK_SOL_STA_UNKNOWN"</i>
Integer feasible point	<i>"MSK_PRO_STA_PRIM_FEAS"</i>	<i>"MSK_SOL_STA_PRIM_FEAS"</i>
No conclusion	<i>"MSK_PRO_STA_UNKNOWN"</i>	<i>"MSK_SOL_STA_UNKNOWN"</i>

7.1.4 Retrieving solution values

After the meaning and quality of the solution (or certificate) have been established, we can query for the actual numerical values. They can be accessed using:

- `rsolitr$pobjval`, `r$solitrdobjval` — the primal and dual objective value.
- `rsolitr$xx` — solution values for the variables.
- `rsolitr$y`, `r$solitrslx` and so on — dual values for the linear constraints

and many other fields of the `solution_info` structure (replace `itr` with `bas` or `int` for other solution types). Note that if the optimization failed then the `r$sol` field may not exist and attempting to access it will cause an error.

7.1.5 Source code example

Below is a source code example with a simple framework for assessing and retrieving the solution to a conic optimization problem.

Listing 7.1: Sample framework for checking optimization result.

```
response <- function()
{
  # In this example we set up a simple problem
  prob <- setupProblem()

  # (Optionally) Uncomment the next line to get solution status Unknown
  # prob$iparam <- list(INTPNT_MAX_ITERATIONS=1)

  # Perform the optimization.
  r <- mosek(prob, list(verbose=0))
  # Use the line below instead to get more log output
  #r <- mosek(prob, list(verbose=10))

  # Expected result: The solution status of the interior-point solution is optimal.

  # Check if there was a fatal error
  if (r$response$code != 0 && startsWith(r$response$msg, "MSK_RES_ERR"))
  {
    print(sprintf("Optimization error: %s (%d)", r$response$msg, r$response$code))
  }
  else
  {
    if (r$sol$itr$solsta == 'OPTIMAL')
    {
      print('An optimal interior-point solution is located:')
      print(r$sol$itr$xx)
    }
    else if (r$sol$itr$solsta == 'DUAL_INFEAS_CER')
    {
      print('Dual infeasibility certificate found.')
    }
    else if (r$sol$itr$solsta == 'PRIM_INFEAS_CER')
    {
      print('Primal infeasibility certificate found.')
    }
    else if (r$sol$itr$solsta == 'UNKNOWN')
    {
      # The solutions status is unknown. The termination code
      # indicates why the optimizer terminated prematurely.
      print('The solution status is unknown.')
      print(sprintf('Termination code: %s (%d).', r$response$msg, r$response$code))
    }
  }
}
```

(continues on next page)

```

    }
    else
    {
        printf('An unexpected solution status is obtained.')
    }
}
}

```

7.2 Errors and exceptions

Response codes

The function `mosek` returns a **response code** (and its human-readable description), informing if optimization was performed correctly, and if not, what error occurred. The expected response, indicating successful execution, is always `"MSK_RES_OK"`. Typical errors include:

- referencing a nonexisting variable (for example with too large index),
- incompatible dimensions of input data matrices,
- NaN in the input data,
- duplicate conic variable,
- error in the optimizer.

A full list of response codes, error, warning and termination codes can be found in the [API reference](#). For example, if the objective vector contains a NaN then

```

r <- mosek(prob)
r$response

```

will produce as output:

```

$code
[1] 1470

$msg
[1] "MSK_RES_ERR_NAN_IN_C: c contains an invalid floating point value, i.e. a NaN."

```

Optimizer errors and warnings

The optimizer may also produce warning messages. They indicate non-critical but important events, that will not prevent solver execution, but may be an indication that something in the optimization problem might be improved. Warning messages are normally printed to a log stream (see [Sec. 7.3](#)). A typical warning is, for example:

```

MOSEK warning 53: A numerically large upper bound value 6.6e+09 is specified for constraint
↪ 'C69200' (46020).

```

Warnings can also be suppressed by setting the `MSK_IPAR_MAX_NUM_WARNINGS` parameter to zero, if they are well-understood.

Error and solution status handling example

Below is a source code example with a simple framework for handling major errors when assessing and retrieving the solution to a conic optimization problem.

Listing 7.2: Sample framework for checking optimization result.

```
response <- function()
{
  # In this example we set up a simple problem
  prob <- setupProblem()

  # (Optionally) Uncomment the next line to get solution status Unknown
  # prob$iparam <- list(INTPNT_MAX_ITERATIONS=1)

  # Perform the optimization.
  r <- mosek(prob, list(verbose=0))
  # Use the line below instead to get more log output
  #r <- mosek(prob, list(verbose=10))

  # Expected result: The solution status of the interior-point solution is optimal.

  # Check if there was a fatal error
  if (r$response$code != 0 && startsWith(r$response$msg, "MSK_RES_ERR"))
  {
    print(sprintf("Optimization error: %s (%d)", r$response$msg, r$response$code))
  }
  else
  {
    if (r$sol$itr$solsta == 'OPTIMAL')
    {
      print('An optimal interior-point solution is located:')
      print(r$sol$itr$xx)
    }
    else if (r$sol$itr$solsta == 'DUAL_INFEAS_CER')
    {
      print('Dual infeasibility certificate found.')
    }
    else if (r$sol$itr$solsta == 'PRIM_INFEAS_CER')
    {
      print('Primal infeasibility certificate found.')
    }
    else if (r$sol$itr$solsta == 'UNKNOWN')
    {
      # The solutions status is unknown. The termination code
      # indicates why the optimizer terminated prematurely.
      print('The solution status is unknown.')
      print(sprintf('Termination code: %s (%d).', r$response$msg, r$response$code))
    }
    else
    {
      printf('An unexpected solution status is obtained.')
    }
  }
}
}
```

7.3 Input/Output

7.3.1 Stream logging

By default the solver prints a log output analogous to the one produced by the command-line version of **MOSEK**. Logging may be turned off using the option `verbose`, for example:

```
r <- mosek(prob, list(verbose=0))
```

7.3.2 Log verbosity

The logging verbosity can be controlled by setting the relevant parameters, as for instance

- `MSK_IPAR_LOG`,
- `MSK_IPAR_LOG_INTPNT`,
- `MSK_IPAR_LOG_MIO`,
- `MSK_IPAR_LOG_CUT_SECOND_OPT`,
- `MSK_IPAR_LOG_SIM`, and
- `MSK_IPAR_LOG_SIM_MINOR`.

Each parameter controls the output level of a specific functionality or algorithm. The main switch is `MSK_IPAR_LOG` which affect the whole output. The actual log level for a specific functionality is determined as the minimum between `MSK_IPAR_LOG` and the relevant parameter. For instance, the log level for the output produce by the interior-point algorithm is tuned by the `MSK_IPAR_LOG_INTPNT`; the actual log level is defined by the minimum between `MSK_IPAR_LOG` and `MSK_IPAR_LOG_INTPNT`.

Tuning the solver verbosity may require adjusting several parameters. It must be noticed that verbose logging is supposed to be of interest during debugging and tuning. When output is no more of interest, the user can easily disable it globally with `MSK_IPAR_LOG`. Larger values of `MSK_IPAR_LOG` do not necessarily result in increased output.

By default **MOSEK** will reduce the amount of log information after the first optimization on a given problem. To get full log output on subsequent re-optimizations set `MSK_IPAR_LOG_CUT_SECOND_OPT` to zero.

7.3.3 Saving a problem to a file

An optimization problem can be dumped to a file using the function `mosek_write`. The file format will be determined from the filename's extension. Supported formats are listed in Sec. 14 together with a table of problem types supported by each.

For instance the problem can be written to an OPF file with

```
r <- mosek_write(prob, "dump.opf");
```

All formats can be compressed with `gzip` by appending the `.gz` extension, for example

```
r <- mosek_write(prob, "dump.task.gz");
```

Some remarks:

- The problem is written to the file as it is represented in the underlying *optimizer task*, that is including any auxiliary variables introduced by the R-to-C interface, if applicable.
- Unnamed variables are given generic names. It is therefore recommended to use meaningful variable names if the problem file is meant to be human-readable.
- The `task` format is **MOSEK**'s native file format which contains all the problem data as well as solver settings.

7.3.4 Reading a problem from a file

A problem saved in any of the supported file formats can be read directly into a *problem* structure using the function `mosek_read`. Afterwards the problem can be optimized, modified, etc.

```
r <- mosek_read("dump.opf");  
r$prob
```

7.4 Setting solver parameters

MOSEK comes with a large number of parameters that allows the user to tune the behavior of the optimizer. The typical settings which can be changed with solver parameters include:

- choice of the optimizer for linear problems,
- choice of primal/dual solver,
- turning presolve on/off,
- turning heuristics in the mixed-integer optimizer on/off,
- level of multi-threading,
- feasibility tolerances,
- solver termination criteria,
- behaviour of the license manager,

and more. All parameters have default settings which will be suitable for most typical users. The API reference contains:

- *Full list of parameters*
- *List of parameters grouped by topic*

Setting parameters

Each parameter is identified by a unique name and it can accept either integers, floating point values, symbolic strings or symbolic values. Parameters are set in the lists `iparam`, `dparam` and `sparam` of the structure `problem` and passed with the problem to `mosek`.

Some parameters can accept symbolic strings from a fixed set. The set of accepted values for every parameter is provided in the API reference.

For example, the following piece of code sets up parameters which choose and tune the interior point optimizer before solving a problem.

Listing 7.3: Parameter setting example.

```
# Set log level (integer parameter)
# and select interior-point optimizer... (integer parameter)
# ... without basis identification (integer parameter)
prob$iparam <- list(LOG=0, OPTIMIZER="OPTIMIZER_INTPNT", INTPNT_BASIS="BI_NEVER")

# Set relative gap tolerance (double parameter)
prob$dparam <- list(INTPNT_CO_TOL_REL_GAP=1.0e-7)

# Solve the problem
r <- mosek(prob)
```

7.5 Retrieving information items

After the optimization the user has access to the solution as well as to a report containing a large amount of additional *information items*. For example, one can obtain information about:

- **timing**: total optimization time, time spent in various optimizer subroutines, number of iterations, etc.
- **solution quality**: feasibility measures, solution norms, constraint and bound violations, etc.
- **problem structure**: counts of variables of different types, constraints, nonzeros, etc.
- **integer optimizer**: integrality gap, objective bound, number of cuts, etc.

and more. Information items are numerical values of integer, long integer or double type. The full list can be found in the API reference:

- *Double*
- *Integer*
- *Long*

Retrieving the values

Values of information items are only returned if the `getinfo` option is set to `TRUE`. They are available in the fields:

- `r$dinfo` for a double information item,
- `r$iinfo` for an integer or long integer information item.

Each information item is identified by a unique name. The example below reads two pieces of data from the solver: total optimization time and the number of interior-point iterations.

Listing 7.4: Information items example.

```
opts <- list(getinfo=TRUE)
r <- mosek(prob, opts)

print(r$dinfo$OPTIMIZER_TIME)
print(r$iinfo$INTPNT_ITER)
```

Chapter 8

Debugging Tutorials

This collection of tutorials contains basic techniques for debugging optimization problems using tools available in **MOSEK**: optimizer log, solution summary, infeasibility report, command-line tools. It is intended as a first line of technical help for issues such as: Why do I get solution status *unknown* and how can I fix it? Why is my model infeasible while it shouldn't be? Should I change some parameters? Can the model solve faster? etc.

The major steps when debugging a model are always:

- Consult the log output. It is enabled by default, but if necessary switch it on explicitly with:

```
r <- mosek(prob, list(verbose=10))
```

- Run the optimization and analyze the log output, see [Sec. 8.1](#). In particular:
 - check if the problem setup (number of constraints/variables etc.) matches your expectation.
 - check solution summary and solution status.
- Dump the problem to disk if necessary to continue analysis. See [Sec. 7.3.3](#).
 - use a human-readable text format, such as `*.opf` if you want to check the problem structure by hand. Assign names to variables and constraints to make them easier to identify.

```
r <- mosek_write(prob, "dump.opf");
```

- use the **MOSEK** native format `*.task.gz` when submitting a bug report or support question.

```
r <- mosek_write(prob, "dump.task.gz");
```

- Fix problem setup, improve the model, locate infeasibility or adjust parameters, depending on the diagnosis.

See the following sections for details.

8.1 Understanding optimizer log

The optimizer produces a log which splits roughly into four sections:

1. summary of the input data,
2. presolve and other pre-optimize problem setup stages,
3. actual optimizer iterations,
4. solution summary.

In this tutorial we show how to analyze the most important parts of the log when initially debugging a model: input data (1) and solution summary (4). For the iterations log (3) see [Sec. 12.3.4](#) or [Sec. 12.4.8](#).

8.1.1 Input data

If **MOSEK** behaves very far from expectations it may be due to errors in problem setup. The log file will begin with a summary of the structure of the problem, which looks for instance like:

```
Problem
  Name           :
  Objective sense : max
  Type           : CONIC (conic optimization problem)
  Constraints     : 20413
  Cones          : 2508
  Scalar variables : 20414
  Matrix variables : 0
  Integer variables : 0
```

This can be consulted to eliminate simple errors: wrong objective sense, wrong number of variables etc. Note that Fusion, and third-party modeling tools can introduce additional variables and constraints to the model. In the remaining **MOSEK** APIs the problem dimensions should match exactly what the user specified.

If this is not sufficient a bit more information can be obtained by dumping the problem to a file (see [Sec. 8](#)) and using the `anapro` option of any of the command line tools. This will produce a longer summary similar to:

```
** Variables
scalar: 20414      integer: 0      matrix: 0
low: 2082          up: 5014        ranged: 0      free: 12892    fixed: 426

** Constraints
all: 20413
low: 10028        up: 0           ranged: 0      free: 0        fixed: 10385

** Cones
QUAD: 1           dims: 2865: 1
RQUAD: 2507       dims: 3: 2507

** Problem data (numerics)
|c|               nnz: 10028      min=2.09e-05   max=1.00e+00
|A|               nnz: 597023     min=1.17e-10   max=1.00e+00
blx               fin: 2508       min=-3.60e+09   max=2.75e+05
bux               fin: 5440       min=0.00e+00   max=2.94e+08
blc               fin: 20413     min=-7.61e+05   max=7.61e+05
buc               fin: 10385     min=-5.00e-01   max=0.00e+00
```

Again, this can be used to detect simple errors, such as:

- Wrong type of cone was used or it has wrong dimension.
- The bounds for variables or constraints are incorrect or incomplete.
- The model is otherwise incomplete.
- Suspicious values of coefficients.
- For various data sizes the model does not scale as expected.

Finally saving the problem in a human-friendly text format such as LP or OPF (see [Sec. 8](#)) and analyzing it by hand can reveal if the model is correct.

Warnings and errors

At this stage the user can encounter warnings which should not be ignored, unless they are well-understood. They can also serve as hints as to numerical issues with the problem data. A typical warning of this kind is

```
MOSEK warning 53: A numerically large upper bound value 2.9e+08 is specified for variable
↪ 'absh[107]' (2613).
```

Warnings do not stop the problem setup. If, on the other hand, an error occurs then the model will become invalid. The user should make sure to test for errors/exceptions from all API calls that set up the problem and validate the data. See [Sec. 7.2](#) for more details.

8.1.2 Solution summary

The last item in the log is the solution summary.

Continuous problem

Optimal solution

A typical solution summary for a continuous (linear, conic, quadratic) problem looks like:

```
Problem status : PRIMAL_AND_DUAL_FEASIBLE
Solution status : OPTIMAL
Primal.  obj: 8.7560516107e+01    nrm: 1e+02    Viol.  con: 3e-12    var: 0e+00    cones: 3e-11
Dual.    obj: 8.7560521345e+01    nrm: 1e+00    Viol.  con: 5e-09    var: 9e-11    cones: 0e+00
```

It contains the following elements:

- Problem and solution status. For details see [Sec. 7.1.3](#).
- A summary of the primal solution: objective value, infinity norm of the solution vector \mathbf{xx} , maximal violations of constraints, variable bounds and cones. The violation of a linear constraint such as $a^T x \leq b$ is $\max(a^T x - b, 0)$. The violation of a conic constraint $x \in \mathcal{K}$ is the distance $\text{dist}(x, \mathcal{K})$.
- The same for the dual solution.

The features of the solution summary which characterize a very good and accurate solution and a well-posed model are:

- **Status:** The solution status is `OPTIMAL`.
- **Duality gap:** The primal and dual objective values are (almost) identical, which proves the solution is (almost) optimal.
- **Norms:** Ideally the norms of the solution and the objective values should not be too large. This of course depends on the input data, but a huge solution norm can be an indicator of issues with the scaling, conditioning and/or well-posedness of the model. It may also indicate that the problem is borderline between feasibility and infeasibility and sensitive to small perturbations in this respect.
- **Violations:** The violations are close to zero, which proves the solution is (almost) feasible. Observe that due to rounding errors it can be expected that the violations are proportional to the norm (`nrm:`) of the solution. It is rarely the case that violations are exactly zero.

Solution status UNKNOWN

A typical example with solution status `UNKNOWN` due to numerical problems will look like:

```
Problem status : UNKNOWN
Solution status : UNKNOWN
Primal.  obj: 1.3821656824e+01    nrm: 1e+01    Viol.  con: 2e-03    var: 0e+00    cones: 0e+00
Dual.    obj: 3.0119004098e-01    nrm: 5e+07    Viol.  con: 4e-16    var: 1e-01    cones: 0e+00
```

Note that:

- The primal and dual objective are very different.
- The dual solution has very large norm.
- There are considerable violations so the solution is likely far from feasible.

Follow the hints in [Sec. 8.2](#) to resolve the issue.

Solution status UNKNOWN with a potentially useful solution

Solution status UNKNOWN does not necessarily mean that the solution is completely useless. It only means that the solver was unable to make any more progress due to numerical difficulties, and it was not able to reach the accuracy required by the termination criteria (see [Sec. 12.3.2](#)). Consider for instance:

Problem status : UNKNOWN						
Solution status : UNKNOWN						
Primal.	obj:	3.4531019648e+04	nrm:	1e+05	Viol.	con: 7e-02 var: 0e+00 cones: 0e+00
Dual.	obj:	3.4529720645e+04	nrm:	8e+03	Viol.	con: 1e-04 var: 2e-04 cones: 0e+00

Such a solution may still be useful, and it is always up to the user to decide. It may be a good enough approximation of the optimal point. For example, the large constraint violation may be due to the fact that one constraint contained a huge coefficient.

Infeasibility certificate

A primal infeasibility certificate is stored in the dual variables:

Problem status : PRIMAL_INFEASIBLE						
Solution status : PRIMAL_INFEASIBLE_CER						
Dual.	obj:	2.9238975853e+02	nrm:	6e+02	Viol.	con: 0e+00 var: 1e-11 cones: 0e+00

It is a Farkas-type certificate as described in [Sec. 11.2.2](#). In particular, for a good certificate:

- The dual objective is positive for a minimization problem, negative for a maximization problem. Ideally it is well bounded away from zero.
- The norm is not too big and the violations are small (as for a solution).

If the model was not expected to be infeasible, the likely cause is an error in the problem formulation. Use the hints in [Sec. 8.1.1](#) and [Sec. 8.3](#) to locate the issue.

Just like a solution, the infeasibility certificate can be of better or worse quality. The infeasibility certificate above is very solid. However, there can be less clear-cut cases, such as for example:

Problem status : PRIMAL_INFEASIBLE						
Solution status : PRIMAL_INFEASIBLE_CER						
Dual.	obj:	1.6378689238e-06	nrm:	6e+05	Viol.	con: 7e-03 var: 2e-04 cones: 0e+00

This infeasibility certificate is more dubious because the dual objective is positive, but barely so in comparison with the large violations. It also has rather large norm. This is more likely an indication that the problem is borderline between feasibility and infeasibility or simply ill-posed and sensitive to tiny variations in input data. See [Sec. 8.3](#) and [Sec. 8.2](#).

The same remarks apply to dual infeasibility (i.e. unboundedness) certificates. Here the primal objective should be negative a minimization problem and positive for a maximization problem.

8.1.3 Mixed-integer problem

Optimal integer solution

For a mixed-integer problem there is no dual solution and a typical optimal solution report will look as follows:

Problem status : PRIMAL_FEASIBLE						
Solution status : INTEGER_OPTIMAL						
Primal.	obj:	6.0111122960e+06	nrm:	1e+03	Viol.	con: 2e-13 var: 2e-14 itg: 5e-15

The interpretation of all elements is as for a continuous problem. The additional field `itg` denotes the maximum violation of an integer variable from being an exact integer.

Feasible integer solution

If the solver found an integer solution but did not prove optimality, for instance because of a time limit, the solution status will be `PRIMAL_FEASIBLE`:

Problem status : PRIMAL_FEASIBLE							
Solution status : PRIMAL_FEASIBLE							
Primal.	obj:	6.0114607792e+06	nrm:	1e+03	Viol.	con:	2e-13
	var:	2e-13	itg:	4e-15			

In this case it is valuable to go back to the optimizer summary to see how good the best solution is:

31	35	1	0	6.0114607792e+06	6.0078960892e+06	0.06	4.1
Objective of best integer solution : 6.011460779193e+06							
Best objective bound : 6.007896089225e+06							

In this case the best integer solution found has objective value 6.011460779193e+06, the best proved lower bound is 6.007896089225e+06 and so the solution is guaranteed to be within 0.06% from optimum. The same data can be obtained as information items through an API. See also [Sec. 12.4](#) for more details.

Infeasible problem

If the problem is declared infeasible the summary is simply

Problem status : PRIMAL_INFEASIBLE							
Solution status : UNKNOWN							
Primal.	obj:	0.0000000000e+00	nrm:	0e+00	Viol.	con:	0e+00
	var:	0e+00	itg:	0e+00			

If infeasibility was not expected, consult [Sec. 8.3](#).

8.2 Addressing numerical issues

The suggestions in this section should help diagnose and solve issues with numerical instability, in particular UNKNOWN solution status or solutions with large violations. Since numerically stable models tend to solve faster, following these hints can also dramatically shorten solution times.

We always recommend that issues of this kind are addressed by reformulating or rescaling the model, since it is the modeler who has the best insight into the structure of the problem and can fix the cause of the issue.

8.2.1 Formulating problems

Scaling

Make sure that all the data in the problem are of comparable orders of magnitude. This applies especially to the linear constraint matrix. Use [Sec. 8.1.1](#) if necessary. For example a report such as

A	nnz: 597023	min=1.17e-6	max=2.21e+5
---	-------------	-------------	-------------

means that the ratio of largest to smallest elements in **A** is 10^{11} . In this case the user should rescale or reformulate the model to avoid such spread which makes it difficult for **MOSEK** to scale the problem internally. In many cases it may be possible to change the units, i.e. express the model in terms of rescaled variables (for instance work with millions of dollars instead of dollars, etc.).

Similarly, if the objective contains very different coefficients, say

$$\text{maximize } 10^{10}x + y$$

then it is likely to lead to inaccuracies. The objective will be dominated by the contribution from x and y will become insignificant.

Removing huge bounds

Never use a very large number as replacement for ∞ . Instead define the variable or constraint as unbounded from below/above. Similarly, avoid artificial huge bounds if you expect they will not become tight in the optimal solution.

Avoiding linear dependencies

As much as possible try to avoid linear dependencies and near-linear dependencies in the model. See [Example 8.3](#).

Avoiding ill-posedness

Avoid continuous models which are ill-posed: the solution space is degenerate, for example consists of a single point (technically, the Slater condition is not satisfied). In general, this refers to problems which are borderline between feasible and infeasible. See [Example 8.1](#).

Scaling the expected solution

Try to formulate the problem in such a way that the expected solution (both primal and dual) is not very large. Consult the solution summary [Sec. 8.1.2](#) to check the objective values or solution norms.

8.2.2 Further suggestions

Here are other simple suggestions that can help locate the cause of the issues. They can also be used as hints for how to tune the optimizer if fixing the root causes of the issue is not possible.

- Remove the objective and solve the feasibility problem. This can reveal issues with the objective.
- Change the objective or change the objective sense from minimization to maximization (if applicable). If the two objective values are almost identical, this may indicate that the feasible set is very small, possibly degenerate.
- Perturb the data, for instance bounds, very slightly, and compare the results.
- For linear problems: solve the problem using a different optimizer by setting the parameter `MSK_IPAR_OPTIMIZER` and compare the results.
- Force the optimizer to solve the primal/dual versions of the problem by setting the parameter `MSK_IPAR_INTPNT_SOLVE_FORM` or `MSK_IPAR_SIM_SOLVE_FORM`. **MOSEK** has a heuristic to decide whether to dualize, but for some problems the guess is wrong an explicit choice may give better results.
- Solve the problem without presolve or some of its parts by setting the parameter `MSK_IPAR_PRESOLVE_USE`, see [Sec. 12.1](#).
- Use different numbers of threads (`MSK_IPAR_NUM_THREADS`) and compare the results. Very different results indicate numerical issues resulting from round-off errors.

If the problem was dumped to a file, experimenting with various parameters is facilitated with the **MOSEK** Command Line Tool or **MOSEK** Python Console [Sec. 8.4](#).

8.2.3 Typical pitfalls

Example 8.1 (Ill-posedness). A toy example of this situation is the feasibility problem

$$(x - 1)^2 \leq 1, (x + 1)^2 \leq 1$$

whose only solution is $x = 0$ and moreover replacing any 1 on the right hand side by $1 - \varepsilon$ makes the problem infeasible and replacing it by $1 + \varepsilon$ yields a problem whose solution set is an interval (fully-dimensional). This is an example of ill-posedness.

Example 8.2 (Huge solution). If the norm of the expected solution is very large it may lead to numerical issues or infeasibility. For example the problem

$$(10^{-4}, x, 10^3) \in \mathcal{Q}_r^3$$

may be declared infeasible because the expected solution must satisfy $x \geq 5 \cdot 10^9$.

Example 8.3 (Near linear dependency). Consider the following problem:

$$\begin{array}{llllll} \text{minimize} & & & & & \\ \text{subject to} & x_1 & + & x_2 & & = 1, \\ & & & & x_3 & + & x_4 & = 1, \\ & - & x_1 & & - & x_3 & & = -1 + \varepsilon, \\ & & - & x_2 & & - & x_4 & = -1, \\ & x_1, & & x_2, & & x_3, & & x_4 & \geq 0. \end{array}$$

If we add the equalities together we obtain:

$$0 = \varepsilon$$

which is infeasible for any $\varepsilon \neq 0$. Here infeasibility is caused by a linear dependency in the constraint matrix coupled with a precision error represented by the ε . Indeed if a problem contains linear dependencies then the problem is either infeasible or contains redundant constraints. In the above case any of the equality constraints can be removed while not changing the set of feasible solutions. To summarize linear dependencies in the constraints can give rise to infeasible problems and therefore it is better to avoid them.

Example 8.4 (Presolving very tight bounds). Next consider the problem

$$\begin{array}{llll} \text{minimize} & & & \\ \text{subject to} & x_1 - 0.01x_2 & = & 0, \\ & x_2 - 0.01x_3 & = & 0, \\ & x_3 - 0.01x_4 & = & 0, \\ & x_1 & \geq & -10^{-9}, \\ & x_1 & \leq & 10^{-9}, \\ & x_4 & \geq & 10^{-4}. \end{array}$$

Now the **MOSEK** presolve will, for the sake of efficiency, fix variables (and constraints) that have tight bounds where tightness is controlled by the parameter `MSK_DPAR_PRESOLVE_TOL_X`. Since the bounds

$$-10^{-9} \leq x_1 \leq 10^{-9}$$

are tight, presolve will set $x_1 = 0$. It easy to see that this implies $x_4 = 0$, which leads to the incorrect conclusion that the problem is infeasible. However a tiny change of the value 10^{-9} makes the problem feasible. In general it is recommended to avoid ill-posed problems, but if that is not possible then one solution is to reduce parameters such as `MSK_DPAR_PRESOLVE_TOL_X` to say 10^{-10} . This will at least make sure that presolve does not make the undesired reduction.

8.3 Debugging infeasibility

This section contains hints for debugging problems that are unexpectedly infeasible. It is always a good idea to remove the objective, i.e. only solve a feasibility problem when debugging such issues.

8.3.1 Numerical issues

Infeasible problem status may be just an artifact of numerical issues appearing when the problem is badly-scaled, barely feasible or otherwise ill-conditioned so that it is unstable under small perturbations of the data or round-off errors. This may be visible in the solution summary if the infeasibility certificate has poor quality. See [Sec. 8.1.2](#) for how to diagnose that and [Sec. 8.2](#) for possible hints. [Sec. 8.2.3](#) contains examples of situations which may lead to infeasibility for numerical reasons.

We refer to [Sec. 8.2](#) for further information on dealing with those sort of issues. For the rest of this section we concentrate on the case when the solution summary leaves little doubt that the problem solved by the optimizer actually is infeasible.

8.3.2 Locating primal infeasibility

As an example of a primal infeasible problem consider minimizing the cost of transportation between a number of production plants and stores: Each plant produces a fixed number of goods, and each store has a fixed demand that must be met. Supply, demand and cost of transportation per unit are given in [Fig. 8.1](#).

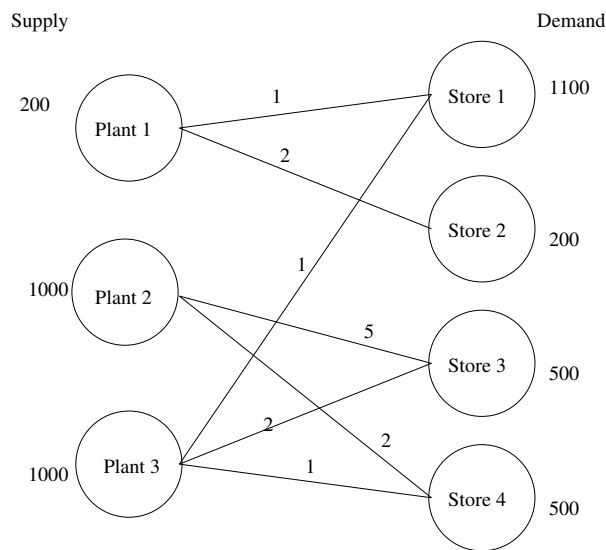


Fig. 8.1: Supply, demand and cost of transportation.

The problem represented in [Fig. 8.1](#) is infeasible, since the total demand

$$2300 = 1100 + 200 + 500 + 500$$

exceeds the total supply

$$2200 = 200 + 1000 + 1000$$

If we denote the number of transported goods from plant i to store j by x_{ij} , the problem can be formulated as the LP:

$$\begin{aligned}
 & \text{minimize} && x_{11} + 2x_{12} + 5x_{23} + 2x_{24} + x_{31} + 2x_{33} + x_{34} \\
 & \text{subject to} && s_0 : x_{11} + x_{12} \leq 200, \\
 & && s_1 : x_{23} + x_{24} \leq 1000, \\
 & && s_2 : x_{31} + x_{33} + x_{34} \leq 1000, \\
 & && d_1 : x_{11} + x_{31} = 1100, \\
 & && d_2 : x_{12} = 200, \\
 & && d_3 : x_{23} + x_{33} = 500, \\
 & && d_4 : x_{24} + x_{34} = 500, \\
 & && x_{ij} \geq 0.
 \end{aligned} \tag{8.1}$$

Solving problem (8.1) using **MOSEK** will result in an infeasibility status. The infeasibility certificate is contained in the dual variables and can be accessed from an API. The variables and constraints with nonzero solution values form an infeasible subproblem, which frequently is very small. See [Sec. 11.1.2](#) or [Sec. 11.2.2](#) for detailed specifications of infeasibility certificates.

A short infeasibility report can also be printed to the log stream. It can be turned on by setting the parameter `MSK_IPAR_INFEAS_REPORT_AUTO` to `"MSK_ON"`. This causes **MOSEK** to print a report on variables and constraints which are involved in infeasibility in the above sense, i.e. have nonzero values in the certificate. The parameter `MSK_IPAR_INFEAS_REPORT_LEVEL` controls the amount of information presented in the infeasibility report. The default value is 1. For the above example the report is

MOSEK PRIMAL INFEASIBILITY REPORT.					
Problem status: The problem is primal infeasible					
The following constraints are involved in the primal infeasibility.					
Index	Name	Lower bound	Upper bound	Dual lower	Dual upper
0	s0	NONE	2.000000e+002	0.000000e+000	1.000000e+000
2	s2	NONE	1.000000e+003	0.000000e+000	1.000000e+000
3	d1	1.100000e+003	1.100000e+003	1.000000e+000	0.000000e+000
4	d2	2.000000e+002	2.000000e+002	1.000000e+000	0.000000e+000
The following bound constraints are involved in the infeasibility.					
Index	Name	Lower bound	Upper bound	Dual lower	Dual upper
8	x33	0.000000e+000	NONE	1.000000e+000	0.000000e+000
10	x34	0.000000e+000	NONE	1.000000e+000	0.000000e+000

The infeasibility report is divided into two sections corresponding to constraints and variables. It is a selection of those lines from the problem solution which are important in understanding primal infeasibility. In this case the constraints s0, s2, d1, d2 and variables x33, x34 are of importance because of nonzero dual values. The columns `Dual lower` and `Dual upper` contain the values of dual variables s_l^c , s_u^c , s_l^x and s_u^x in the primal infeasibility certificate (see [Sec. 11.1.2](#)).

In our example the certificate means that an appropriate linear combination of constraints s0, s1 with coefficient $s_u^c = 1$, constraints d1 and d2 with coefficient $s_u^c - s_l^c = 0 - 1 = -1$ and lower bounds on x33 and x34 with coefficient $-s_l^x = -1$ gives a contradiction. Indeed, the combination of the four involved constraints is $x_{33} + x_{34} \leq -100$ (as indicated in the introduction, the difference between supply and demand).

It is also possible to extract the infeasible subproblem with the command-line tool. For an infeasible problem called `infeas.lp` the command:

```
mosek -d MSK_IPAR_INFEAS_REPORT_AUTO MSK_ON infeas.lp -info rinfeas.lp
```

will produce the file `rinfeas.bas.inf.lp` which contains the infeasible subproblem. Because of its size it may be easier to work with than the original problem file.

Returning to the transportation example, we discover that removing the fifth constraint $x_{12} = 200$ makes the problem feasible. Almost all undesired infeasibilities should be fixable at the modeling stage.

8.3.3 Locating dual infeasibility

A problem may also be *dual infeasible*. In this case the primal problem is usually unbounded, meaning that feasible solutions exist such that the objective tends towards infinity. For example, consider the problem

$$\begin{aligned}
& \text{maximize} && 200y_1 + 1000y_2 + 1000y_3 + 1100y_4 + 200y_5 + 500y_6 + 500y_7 \\
& \text{subject to} && y_1 + y_4 \leq 1, \quad y_1 + y_5 \leq 2, \quad y_2 + y_6 \leq 5, \quad y_2 + y_7 \leq 2, \\
& && y_3 + y_4 \leq 1, \quad y_3 + y_6 \leq 2, \quad y_3 + y_7 \leq 1 \\
& && y_1, y_2, y_3 \leq 0
\end{aligned}$$

which is dual to (8.1) (and therefore is dual infeasible). The dual infeasibility report may look as follows:

MOSEK DUAL INFEASIBILITY REPORT.					
Problem status: The problem is dual infeasible					
The following constraints are involved in the infeasibility.					
Index	Name	Activity	Objective	Lower bound	Upper bound
5	x33	-1.000000e+00		NONE	2.000000e+00
6	x34	-1.000000e+00		NONE	1.000000e+00
The following variables are involved in the infeasibility.					
Index	Name	Activity	Objective	Lower bound	Upper bound
0	y1	-1.000000e+00	2.000000e+02	NONE	0.000000e+00
2	y3	-1.000000e+00	1.000000e+03	NONE	0.000000e+00
3	y4	1.000000e+00	1.100000e+03	NONE	NONE
4	y5	1.000000e+00	2.000000e+02	NONE	NONE
Interior-point solution summary					
Problem status : DUAL_INFEASIBLE					
Solution status : DUAL_INFEASIBLE_CER					
Primal. obj: 1.0000000000e+02 nrm: 1e+00 Viol. con: 0e+00 var: 0e+00					

In the report we see that the variables y1, y3, y4, y5 and two constraints contribute to infeasibility with non-zero values in the Activity column. Therefore

$$(y_1, \dots, y_7) = (-1, 0, -1, 1, 1, 0, 0)$$

is the dual infeasibility certificate as in [Sec. 11.1.2](#). This just means, that along the ray

$$(0, 0, 0, 0, 0, 0, 0) + t(y_1, \dots, y_7) = (-t, 0, -t, t, t, 0, 0), \quad t > 0,$$

which belongs to the feasible set, the objective value $100t$ can be arbitrarily large, i.e. the problem is unbounded.

In the example problem we could

- Add a lower bound on y3. This will directly invalidate the certificate of dual infeasibility.
- Increase the objective coefficient of y3. Changing the coefficients sufficiently will invalidate the inequality $c^T y^* > 0$ and thus the certificate.

8.3.4 Suggestions

Primal infeasibility

When trying to understand what causes the unexpected primal infeasible status use the following hints:

- Remove the objective function. This does not change the infeasibility status but simplifies the problem, eliminating any possibility of issues related to the objective function.
- Remove cones, semidefinite variables and integer constraints. Solve only the linear part of the problem. Typical simple modeling errors will lead to infeasibility already at this stage.
- Consider whether your problem has some obvious necessary conditions for feasibility and examine if these are satisfied, e.g. total supply should be greater than or equal to total demand.
- Verify that coefficients and bounds are reasonably sized in your problem.
- See if there are any obvious contradictions, for instance a variable is bounded both in the variables and constraints section, and the bounds are contradictory.
- Consider replacing suspicious equality constraints by inequalities. For instance, instead of $x_{12} = 200$ see what happens for $x_{12} \geq 200$ or $x_{12} \leq 200$.

- Relax bounds of the suspicious constraints or variables.
- For integer problems, remove integrality constraints on some/all variables and see if the problem solves.
- Form an **elastic model**: allow to violate constraints at a cost. Introduce slack variables and add them to the objective as penalty. For instance, suppose we have a constraint

$$\begin{array}{ll}\text{minimize} & c^T x, \\ \text{subject to} & a^T x \leq b.\end{array}$$

which might be causing infeasibility. Then create a new variable y and form the problem which contains:

$$\begin{array}{ll}\text{minimize} & c^T x + y, \\ \text{subject to} & a^T x \leq b + y.\end{array}$$

Solving this problem will reveal by how much the constraint needs to be relaxed in order to become feasible. This is equivalent to inspecting the infeasibility certificate but may be more intuitive.

- If you think you have a feasible solution or its part, fix all or some of the variables to those values. Presolve will propagate them through the model and potentially reveal more localized sources of infeasibility.
- Dump the problem in OPF or LP format and verify that the problem that was passed to the optimizer corresponds to the problem expressed in the high-level modeling language, if any such was used.

Dual infeasibility

When trying to understand what causes the unexpected dual infeasible status use the following hints:

- Verify that the objective coefficients are reasonably sized.
- Check if no bounds and constraints are missing, for example if all variables that should be nonnegative have been declared as such etc.
- Strengthen bounds of the suspicious constraints or variables.
- Form an series of models with decreasing bounds on the objective, that is, instead of objective

$$\text{minimize } c^T x$$

solve the problem with an additional constraint such as

$$c^T x = -10^5$$

and inspect the solution to figure out the mechanism behind arbitrarily decreasing objective values. This is equivalent to inspecting the infeasibility certificate but may be more intuitive.

- Dump the problem in OPF or LP format and verify that the problem that was passed to the optimizer corresponds to the problem expressed in the high-level modeling language, if any such was used.

Please note that modifying the problem to invalidate the reported certificate does *not* imply that the problem becomes feasible — the reason for infeasibility may simply *move*, resulting a problem that is still infeasible, but for a different reason. More often, the reported certificate can be used to give a hint about errors or inconsistencies in the model that produced the problem.

8.4 Python Console

The **MOSEK** Python Console is an alternative to the **MOSEK** Command Line Tool. It can be used for interactive loading, solving and debugging optimization problems stored in files, for example **MOSEK** task files. It facilitates debugging techniques described in [Sec. 8](#).

8.4.1 Usage

The tool requires Python 2 or 3. The **MOSEK** interface for Python must be installed following the installation instructions for Python API or Python Fusion API. In the basic case it should be sufficient to execute the script

```
python setup.py install --user
```

in the directory containing the **MOSEK** Python module.

The Python Console is contained in the file `mosekconsole.py` in the folder with **MOSEK** binaries. It can be copied to an arbitrary location. The file is also available for [download here](#) (`mosekconsole.py`).

To run the console in interactive mode use

```
python mosekconsole.py
```

To run the console in batch mode provide a semicolon-separated list of commands as the second argument of the script, for example:

```
python mosekconsole.py "read data.task.gz; solve form=dual; writesol data"
```

The script is written using the **MOSEK** Python API and can be extended by the user if more specific functionality is required. We refer to the documentation of the Python API.

8.4.2 Examples

To read a problem from `data.task.gz`, solve it, and write solutions to `data.sol`, `data.bas` or `data.itg`:

```
read data.task.gz; solve; writesol data
```

To convert between file formats:

```
read data.task.gz; write data.mps
```

To set a parameter before solving:

```
read data.task.gz; param INTPNT_CO_TOL_DFEAS 1e-9; solve"
```

To list parameter values related to the mixed-integer optimizer in the task file:

```
read data.task.gz; param MIO
```

To print a summary of problem structure:

```
read data.task.gz; anapro
```

To solve a problem forcing the dual and switching off presolve:

```
read data.task.gz; solve form=dual presolve=no
```

To write an infeasible subproblem to a file for debugging purposes:

```
read data.task.gz; solve; infsub; write inf.opf
```

8.4.3 Full list of commands

Below is a brief description of all the available commands. Detailed information about a specific command `cmd` and its options can be obtained with

```
help cmd
```

Table 8.1: List of commands of the MOSEK Python Console.

Command	Description
help [command]	Print list of commands or info about a specific command
log filename	Save the session to a file
intro	Print MOSEK splashscreen
testlic	Test the license system
read filename	Load problem from file
reread	Reload last problem file
solve [options]	Solve current problem
write filename	Write current problem to file
param [name [value]]	Set a parameter or get parameter values
paramdef	Set all parameters to default values
info [name]	Get an information item
anapro	Analyze problem data
hist	Plot a histogram of problem data
histsol	Plot a histogram of the solutions
spy	Plot the sparsity pattern of the A matrix
truncate epsilon	Truncate small coefficients down to 0
anasol	Analyze solutions
removeitg	Remove integrality constraints
infsub	Replace current problem with its infeasible subproblem
writesol basename	Write solution(s) to file(s) with given basename
delsol	Remove all solutions from the task
exit	Leave

Chapter 9

Technical guidelines

This section contains some more in-depth technical guidelines for Rmosek package, not strictly necessary for basic use of **MOSEK**.

9.1 Multithreading

Parallelization

The interior-point and mixed-integer optimizers in **MOSEK** are parallelized. By default **MOSEK** will automatically select the number of threads. However, the maximum number of threads allowed can be changed by setting the parameter `MSK_IPAR_NUM_THREADS` and related parameters. This should never exceed the number of cores. See [Sec. 12](#) and [Sec. 12.4](#) for more details.

The speed-up obtained when using multiple threads is highly problem and hardware dependent. We recommend experimenting with various thread numbers to determine the optimal settings. For small problems using multiple threads may be counter-productive because of the associated overhead.

Determinism

By default the optimizer is run-to-run deterministic, which means that it will return the same answer each time it is run on the same machine with the same input, the same parameter settings (including number of threads) and no time limits.

Setting the number of threads

The number of threads the optimizer uses can be changed with the parameter `MSK_IPAR_NUM_THREADS`.

For conic problems (when the conic optimizer is used) the value set at the first optimization will remain fixed through the lifetime of the process. The thread pool will be reserved once for all and subsequent changes to `MSK_IPAR_NUM_THREADS` will have no effect. The only possibility at that point is to switch between multi-threaded and single-threaded interior-point optimization using the parameter `MSK_IPAR_INTPNT_MULTI_THREAD`.

9.2 Parallel optimization Using the The Multicore Package

The *multicore* package works by copying the full memory state of the R session to new processes. While this seems like a large overhead, in practice, the copy is delayed until modification assuring a smooth parallel execution. The downside is that this low-level memory state copy is not safe for all types of resources. As an example, parallel interactions with the GUI or on-screen devices can cause the R session to crash. It is thus recommended only to use the *multicore* package in console R.

In the Rmosek package a license is an externally acquired resource, and attempts to simply copy the memory state of this resource will provoke a session crash. Thus, licenses should always be released before the time of parallelization.

Always call `mosek_clean()` before a parallelizing operator. Failure to do so is likely to provoke session crashes.

A consequence of this is that each new process will be using a separate license. That is, your license system should allow as many licenses to be checked out simultaneously, as many parallel optimizations you want to run. Please note that unlimited academic and commercial licenses are available at **MOSEK**.

9.3 The license system

MOSEK is a commercial product that **always** needs a valid license to work. **MOSEK** uses a third party license manager to implement license checking. The number of license tokens provided determines the number of optimizations that can be run simultaneously.

By default a license token remains checked out from the first optimization until the end of the **MOSEK** session, i.e.

- a license token is checked out when `mosek` is called the first time and
- it is returned when R is terminated, or `mosek_clean` is called.

Starting the optimization when no license tokens are available will result in an error.

Default behaviour of the license system can be changed in several ways:

- Setting the parameter `MSK_IPAR_CACHE_LICENSE` to `"MSK_OFF"` will force **MOSEK** to return the license token immediately after the optimization completed.
- Setting the parameter `MSK_IPAR_LICENSE_WAIT` will force **MOSEK** to wait until a license token becomes available instead of returning with an error.
- The license can be manually released by calling `mosek_clean`.

Chapter 10

Case Studies

In this section we present some case studies in which the Rmosek package is used to solve real-life applications. These examples involve some more advanced modeling skills and possibly some input data. The user is strongly recommended to first read the basic tutorials of [Sec. 6](#) before going through these advanced case studies.

- *Portfolio Optimization*
 - **Keywords:** Markowitz model, variance, risk, efficient frontier, transaction cost, market impact cost, cardinality constraints
 - **Type:** Conic Quadratic, Power Cone, Mixed-Integer
- *Least squares and other norm minimization problems*
 - **Keywords:** Least squares, regression, 2-norm, 1-norm, p-norm, ridge, lasso
 - **Type:** Conic Quadratic, Power Cone

10.1 Portfolio Optimization

In this section the Markowitz portfolio optimization problem and variants are implemented using Rmosek package.

- *Basic Markowitz model*
- *Efficient frontier*
- *Factor model and efficiency*
- *Market impact costs*
- *Transaction costs*
- *Cardinality constraints*

10.1.1 The Basic Model

The classical Markowitz portfolio optimization problem considers investing in n stocks or assets held over a period of time. Let x_j denote the amount invested in asset j , and assume a stochastic model where the return of the assets is a random variable r with known mean

$$\mu = \mathbf{E}r$$

and covariance

$$\Sigma = \mathbf{E}(r - \mu)(r - \mu)^T.$$

The return of the investment is also a random variable $y = r^T x$ with mean (or expected return)

$$\mathbf{E}y = \mu^T x$$

and variance

$$\mathbf{E}(y - \mathbf{E}y)^2 = x^T \Sigma x.$$

The standard deviation

$$\sqrt{x^T \Sigma x}$$

is usually associated with risk.

The problem facing the investor is to rebalance the portfolio to achieve a good compromise between risk and expected return, e.g., maximize the expected return subject to a budget constraint and an upper bound (denoted γ) on the tolerable risk. This leads to the optimization problem

$$\begin{aligned} & \text{maximize} && \mu^T x \\ & \text{subject to} && \begin{aligned} e^T x &= w + e^T x^0, \\ x^T \Sigma x &\leq \gamma^2, \\ x &\geq 0. \end{aligned} \end{aligned} \tag{10.1}$$

The variables x denote the investment i.e. x_j is the amount invested in asset j and x_j^0 is the initial holding of asset j . Finally, w is the initial amount of cash available.

A popular choice is $x^0 = 0$ and $w = 1$ because then x_j may be interpreted as the relative amount of the total portfolio that is invested in asset j .

Since e is the vector of all ones then

$$e^T x = \sum_{j=1}^n x_j$$

is the total investment. Clearly, the total amount invested must be equal to the initial wealth, which is

$$w + e^T x^0.$$

This leads to the first constraint

$$e^T x = w + e^T x^0.$$

The second constraint

$$x^T \Sigma x \leq \gamma^2$$

ensures that the variance, is bounded by the parameter γ^2 . Therefore, γ specifies an upper bound of the standard deviation (risk) the investor is willing to undertake. Finally, the constraint

$$x_j \geq 0$$

excludes the possibility of short-selling. This constraint can of course be excluded if short-selling is allowed.

The covariance matrix Σ is positive semidefinite by definition and therefore there exist a matrix G such that

$$\Sigma = GG^T. \tag{10.2}$$

In general the choice of G is **not** unique and one possible choice of G is the Cholesky factorization of Σ . However, in many cases another choice is better for efficiency reasons as discussed in [Sec. 10.1.3](#). For a given G we have that

$$\begin{aligned} x^T \Sigma x &= x^T G G^T x \\ &= \|G^T x\|^2. \end{aligned}$$

Hence, we may write the risk constraint as

$$\gamma \geq \|G^T x\|$$

or equivalently

$$(\gamma, G^T x) \in \mathcal{Q}^{n+1},$$

where \mathcal{Q}^{n+1} is the $(n+1)$ -dimensional quadratic cone. Therefore, problem (10.1) can be written as

$$\begin{aligned} &\text{maximize} && \mu^T x \\ &\text{subject to} && e^T x = w + e^T x^0, \\ & && (\gamma, G^T x) \in \mathcal{Q}^{n+1}, \\ & && x \geq 0, \end{aligned} \tag{10.3}$$

which is a conic quadratic optimization problem that can easily be formulated and solved with Rmosek package. Subsequently we will use the example data

$$\mu = \begin{bmatrix} 0.1073 \\ 0.0737 \\ 0.0627 \end{bmatrix}$$

and

$$\Sigma = 0.1 \cdot \begin{bmatrix} 0.2778 & 0.0387 & 0.0021 \\ 0.0387 & 0.1112 & -0.0020 \\ 0.0021 & -0.0020 & 0.0115 \end{bmatrix}.$$

This implies

$$G^T = \sqrt{0.1} \begin{bmatrix} 0.5271 & 0.0734 & 0.0040 \\ 0 & 0.3253 & -0.0070 \\ 0 & 0 & 0.1069 \end{bmatrix}$$

Why a Conic Formulation?

Problem (10.1) is a convex quadratically constrained optimization problem that can be solved directly using **MOSEK**. Why then reformulate it as a conic quadratic optimization problem (10.3)? The main reason for choosing a conic model is that it is more robust and usually solves faster and more reliably. For instance it is not always easy to numerically validate that the matrix Σ in (10.1) is positive semidefinite due to the presence of rounding errors. It is also very easy to make a mistake so Σ becomes indefinite. These problems are completely eliminated in the conic formulation.

Moreover, observe the constraint

$$\|G^T x\| \leq \gamma$$

more numerically robust than

$$x^T \Sigma x \leq \gamma^2$$

for very small and very large values of γ . Indeed, if say $\gamma \approx 10^4$ then $\gamma^2 \approx 10^8$, which introduces a scaling issue in the model. Hence, using conic formulation we work with the standard deviation instead of variance, which usually gives rise to a better scaled model.

Example code

Listing 10.1 demonstrates how the basic Markowitz model (10.3) is implemented.

Listing 10.1: Code implementing problem (10.3).

```
BasicMarkowitz <- function(
  n,          # Number of assets
  mu,         # An n-dimensional vector of expected returns
  GT,         # A matrix with n columns so (GT')*GT = covariance matrix
  x0,         # Initial holdings
  w,          # Initial cash holding
  gamma)      # Maximum risk (=std. dev) accepted
{
  prob <- list(sense="max")
  prob$c <- mu
  prob$A <- Matrix(1.0, ncol=n)
  prob$bc <- rbind(blc=w+sum(x0),
                  buc=w+sum(x0))
  prob$bx <- rbind(blx=rep(0.0,n),
                  bux=rep(Inf,n))

  # Specify the affine conic constraints.
  NUMCONES <- 1
  prob$F <- rbind(
    Matrix(0.0,ncol=n),
    GT
  )
  prob$g <- c(gamma,rep(0,n))
  prob$cones <- matrix(list(), nrow=3, ncol=NUMCONES)
  rownames(prob$cones) <- c("type","dim","conepar")

  prob$cones[-3,1] <- list("QUAD", n+1)

  # Solve the problem
  r <- mosek(prob,list(verbose=1))
  stopifnot(identical(r$response$code, 0))

  # Return the solution
  x <- r$sol$itr$xx
  list(expret=drop(mu %*% x), stddev=gamma, x=x)
}
```

The source code should be self-explanatory except perhaps for

```
NUMCONES <- 1
prob$F <- rbind(
  Matrix(0.0,ncol=n),
  GT
)
prob$g <- c(gamma,rep(0,n))
prob$cones <- matrix(list(), nrow=3, ncol=NUMCONES)
rownames(prob$cones) <- c("type","dim","conepar")

prob$cones[-3,1] <- list("QUAD", n+1)
```

where the constraint

$$(\gamma, G^T x) \in \mathcal{Q}^{n+1}$$

is created as an *affine conic constraint format* of the form $Fx + g \in \mathcal{K}$, in this specific case:

$$\begin{bmatrix} 0 \\ G^T \end{bmatrix} x + \begin{bmatrix} \gamma \\ 0 \end{bmatrix} \in \mathcal{Q}^{n+1}.$$

10.1.2 The Efficient Frontier

The portfolio computed by the Markowitz model is efficient in the sense that there is no other portfolio giving a strictly higher return for the same amount of risk. An efficient portfolio is also sometimes called a Pareto optimal portfolio. Clearly, an investor should only invest in efficient portfolios and therefore it may be relevant to present the investor with all efficient portfolios so the investor can choose the portfolio that has the desired tradeoff between return and risk.

Given a nonnegative α the problem

$$\begin{aligned} & \text{maximize} && \mu^T x - \alpha x^T \Sigma x \\ & \text{subject to} && e^T x = w + e^T x^0, \\ & && x \geq 0. \end{aligned} \tag{10.4}$$

is one standard way to trade the expected return against penalizing variance. Note that, in contrast to the previous example, we explicitly use the variance ($\|G^T x\|_2^2$) rather than standard deviation ($\|G^T x\|_2$), therefore the conic model includes a rotated quadratic cone:

$$\begin{aligned} & \text{maximize} && \mu^T x - \alpha s \\ & \text{subject to} && e^T x = w + e^T x^0, \\ & && (s, 0.5, G^T x) \in Q_r^{n+2} \quad (\text{equiv. to } s \geq \|G^T x\|_2^2 = x^T \Sigma x), \\ & && x \geq 0. \end{aligned} \tag{10.5}$$

The parameter α specifies the tradeoff between expected return and variance. Ideally the problem (10.4) should be solved for all values $\alpha \geq 0$ but in practice it is impossible. Using the example data from Sec. 10.1.1, the optimal values of return and variance for several values of α are shown in the figure.

Example code

Listing 10.2 demonstrates how to compute the efficient portfolios for several values of α .

Listing 10.2: Code for the computation of the efficient frontier based on problem (10.4).

```
EfficientFrontier <- function(
  n,          # Number of assets
  mu,         # An n-dimensional vector of expected returns
  GT,         # A matrix with n columns so (GT')*GT = covariance matrix
  x0,         # Initial holdings
  w,          # Initial cash holding
  alphas)     # List of risk penalties (we maximize expected return - alpha * variance)
{
  prob <- list(sense="max")
  prob$A <- cbind(Matrix(1.0, ncol=n), 0.0)
  prob$bc <- rbind(blc=w+sum(x0),
                  buc=w+sum(x0))
  prob$bx <- rbind(blx=c(rep(0.0,n), -Inf),
                  bux=rep(Inf,n+1))

  # Specify the affine conic constraints.
  NUMCONES <- 1
  prob$F <- rbind(
    cbind(Matrix(0.0,ncol=n), 1.0),
    rep(0, n+1),
    cbind(GT, 0.0)
  )
  prob$g <- c(0, 0.5, rep(0, n))
  prob$cones <- matrix(list(), nrow=3, ncol=NUMCONES)
  rownames(prob$cones) <- c("type", "dim", "conevar")

  prob$cones[-3,1] <- list("RQUAD", n+2)

  frontier <- matrix(NA, ncol=3, nrow=length(alphas))
```

(continues on next page)

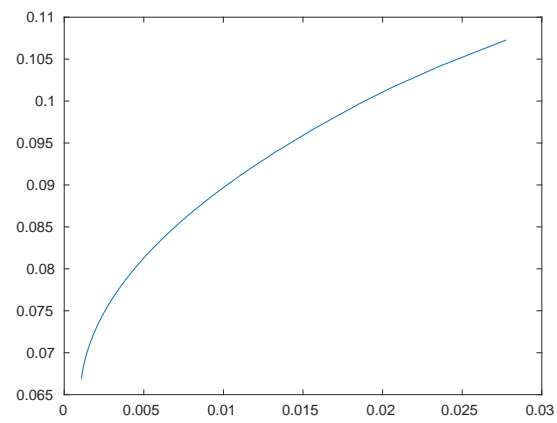


Fig. 10.1: The efficient frontier for the sample data.

```

colnames(frontier) <- c("alpha", "exp.ret.", "variance")

for (i in seq_along(alphas))
{
  prob[c] <- c(mu, -alphas[i])

  r <- mosek(prob, list(verbose=1))
  stopifnot(identical(r$response$code, 0))

  x      <- r$sol$itr$xx[1:n]
  gamma <- r$sol$itr$xx[n+1]

  frontier[i,] <- c(alphas[i], drop(mu*%x), gamma)
}

frontier
}

```

10.1.3 Factor model and efficiency

In practice it is often important to solve the portfolio problem very quickly. Therefore, in this section we discuss how to improve computational efficiency at the modeling stage.

The computational cost is of course to some extent dependent on the number of constraints and variables in the optimization problem. However, in practice a more important factor is the sparsity: the number of nonzeros used to represent the problem. Indeed it is often better to focus on the number of nonzeros in G see (10.2) and try to reduce that number by for instance changing the choice of G .

In other words if the computational efficiency should be improved then it is always good idea to start with focusing at the covariance matrix. As an example assume that

$$\Sigma = D + VV^T$$

where D is a positive definite diagonal matrix. Moreover, V is a matrix with n rows and p columns. Such a model for the covariance matrix is called a factor model and usually p is much smaller than n . In practice p tends to be a small number independent of n , say less than 100.

One possible choice for G is the Cholesky factorization of Σ which requires storage proportional to $n(n+1)/2$. However, another choice is

$$G^T = \begin{bmatrix} D^{1/2} \\ V^T \end{bmatrix}$$

because then

$$GG^T = D + VV^T.$$

This choice requires storage proportional to $n + pn$ which is much less than for the Cholesky choice of G . Indeed assuming p is a constant storage requirements are reduced by a factor of n .

The example above exploits the so-called factor structure and demonstrates that an alternative choice of G may lead to a significant reduction in the amount of storage used to represent the problem. This will in most cases also lead to a significant reduction in the solution time.

The lesson to be learned is that it is important to investigate how the covariance matrix is formed. Given this knowledge it might be possible to make a special choice for G that helps reducing the storage requirements and enhance the computational efficiency. More details about this process can be found in [And13].

10.1.4 Slippage Cost

The basic Markowitz model assumes that there are no costs associated with trading the assets and that the returns of the assets are independent of the amount traded. Neither of those assumptions is usually

valid in practice. Therefore, a more realistic model is

$$\begin{aligned} & \text{maximize} && \mu^T x \\ & \text{subject to} && e^T x + \sum_{j=1}^n T_j(\Delta x_j) = w + e^T x^0, \\ & && x^T \Sigma x \leq \gamma^2, \\ & && x \geq 0. \end{aligned} \tag{10.6}$$

Here Δx_j is the change in the holding of asset j i.e.

$$\Delta x_j = x_j - x_j^0$$

and $T_j(\Delta x_j)$ specifies the transaction costs when the holding of asset j is changed from its initial value. In the next two sections we show two different variants of this problem with two nonlinear cost functions T .

10.1.5 Market Impact Costs

If the initial wealth is fairly small and no short selling is allowed, then the holdings will be small and the traded amount of each asset must also be small. Therefore, it is reasonable to assume that the prices of the assets are independent of the amount traded. However, if a large volume of an asset is sold or purchased, the price, and hence return, can be expected to change. This effect is called market impact costs. It is common to assume that the market impact cost for asset j can be modeled by

$$T_j(\Delta x_j) = m_j |\Delta x_j|^{3/2}$$

where m_j is a constant that is estimated in some way by the trader. See [GK00] [p. 452] for details. From the [Modeling Cookbook](#) we know that $t \geq |z|^{3/2}$ can be modeled directly using the power cone $\mathcal{P}_3^{2/3, 1/3}$:

$$\{(t, z) : t \geq |z|^{3/2}\} = \{(t, z) : (t, 1, z) \in \mathcal{P}_3^{2/3, 1/3}\}$$

Hence, it follows that $\sum_{j=1}^n T_j(\Delta x_j) = \sum_{j=1}^n m_j |x_j - x_j^0|^{3/2}$ can be modeled by $\sum_{j=1}^n m_j t_j$ under the constraints

$$\begin{aligned} z_j &= |x_j - x_j^0|, \\ (t_j, 1, z_j) &\in \mathcal{P}_3^{2/3, 1/3}. \end{aligned}$$

Unfortunately this set of constraints is nonconvex due to the constraint

$$z_j = |x_j - x_j^0| \tag{10.7}$$

but in many cases the constraint may be replaced by the relaxed constraint

$$z_j \geq |x_j - x_j^0|, \tag{10.8}$$

For instance if the universe of assets contains a risk free asset then

$$z_j > |x_j - x_j^0| \tag{10.9}$$

cannot hold for an optimal solution.

If the optimal solution has the property (10.9) then the market impact cost within the model is larger than the true market impact cost and hence money are essentially considered garbage and removed by generating transaction costs. This may happen if a portfolio with very small risk is requested because the only way to obtain a small risk is to get rid of some of the assets by generating transaction costs. We generally assume that this is not the case and hence the models (10.7) and (10.8) are equivalent.

The above observations lead to

$$\begin{aligned} & \text{maximize} && \mu^T x \\ & \text{subject to} && e^T x + m^T t = w + e^T x^0, \\ & && (\gamma, G^T x) \in \mathcal{Q}^{n+1}, \\ & && (t_j, 1, x_j - x_j^0) \in \mathcal{P}_3^{2/3, 1/3}, \quad j = 1, \dots, n, \\ & && x \geq 0. \end{aligned} \tag{10.10}$$

The revised budget constraint

$$e^T x + m^T t = w + e^T x^0$$

specifies that the initial wealth covers the investment and the transaction costs. It should be mentioned that transaction costs of the form

$$t_j \geq |z_j|^p$$

where $p > 1$ is a real number can be modeled with the power cone as

$$(t_j, 1, z_j) \in \mathcal{P}_3^{1/p, 1-1/p}.$$

See the [Modeling Cookbook](#) for details.

Example code

Listing 10.3 demonstrates how to compute an optimal portfolio when market impact cost are included.

Listing 10.3: Implementation of model (10.10).

```
MarkowitzWithMarketImpact <- function(
  n,          # Number of assets
  mu,         # An n-dimensional vector of expected returns
  GT,         # A matrix with n columns so (GT')*GT = covariance matrix
  x0,         # Initial holdings
  w,          # Initial cash holding
  gamma,      # Maximum risk (=std. dev) accepted
  m)          # Market impacts (we use m_j/x_j-x0_j/^(3/2) for j'th asset)
{
  prob <- list(sense="max")
  prob$c <- c(mu, rep(0,n))
  prob$A <- cbind(Matrix(1.0,ncol=n), t(m))
  prob$bc <- rbind(blc=w+sum(x0),
                  buc=w+sum(x0))
  prob$bx <- rbind(blx=rep(0.0,2*n),
                  bux=rep(Inf,2*n))

  # Specify the affine conic constraints.
  # 1) Risk
  Fr <- rbind(
    Matrix(0.0,nrow=1,ncol=2*n),
    cbind(GT, Matrix(0.0,nrow=n,ncol=n))
  )
  gr <- c(gamma,rep(0,n))
  Kr <- matrix(list("QUAD", 1+n, NULL), nrow=3, ncol=1)

  # 2) Market impact (t_j >= |x_j-x0_j|^(3/2))
  # [ t_j ]
  # [ 1 ] \in PPOW(2,1)
  # [ x_j - x0_j ]
  Fm <- sparseMatrix(
    i=c(seq(from=1,by=3,len=n), seq(from=3,by=3,len=n)),
    j=c(seq(from=n+1,len=n), seq(from=1,len=n)),
    x=c(rep(1.0,n), rep(1.0,n)),
    dims=c(3*n, 2*n))
  gm <- rep(c(0,1,0), n)
  gm[seq(from=3,by=3,len=n)] <- -x0
  Km <- matrix(list("PPOW", 3, c(2,1)), nrow=3, ncol=n)

  prob$F <- rbind(Fr, Fm)
  prob$g <- c(gr, gm)
```

(continues on next page)

(continued from previous page)

```

prob$cones <- cbind(Kr, Km)
rownames(prob$cones) <- c("type", "dim", "conepar")

# Solve the problem
r <- mosek(prob, list(verbose=1))
stopifnot(identical(r$response$code, 0))

# Return the solution
x <- r$sol$itr$xx[1:n]
tx <- r$sol$itr$xx[(n+1):(2*n)]
list(expret=drop(mu %*% x), stddev=gamma, cost=drop(m %*% tx), x=x)
}

```

In the last part of the code we extend the affine conic constraint with triples of the form $(t_k, 1, x_k - x_k^0)$. Such a triple is constructed as an affine conic constraint with:

$$\begin{bmatrix} e_{n+k}^T \\ 0 \\ e_k^T \end{bmatrix} \cdot \begin{bmatrix} x \\ t \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ -x_k^0 \end{bmatrix}$$

where e_j denotes the vector of length $2n$ with 1 at position j and 0 otherwise. Membership of a sequence of triples in power cones $\mathcal{P}_3^{2/3, 1/3}$ is specified with the syntax:

```
Km <- matrix(list("PPOW", 3, c(2,1)), nrow=3, ncol=n)
```

Note that the construction `list("PPOW", d, c(a,b))` creates a power cone of dimension d with exponents

$$\frac{a}{a+b}, \frac{b}{a+b}.$$

10.1.6 Transaction Costs

Now assume there is a cost associated with trading asset j given by

$$T_j(\Delta x_j) = \begin{cases} 0, & \Delta x_j = 0, \\ f_j + g_j |\Delta x_j|, & \text{otherwise.} \end{cases}$$

Hence, whenever asset j is traded we pay a fixed setup cost f_j and a variable cost of g_j per unit traded. Given the assumptions about transaction costs in this section problem (10.6) may be formulated as

$$\begin{aligned} & \text{maximize} && \mu^T x \\ & \text{subject to} && e^T x + f^T y + g^T z = w + e^T x^0, \\ & && (\gamma, G^T x) \in \mathcal{Q}^{n+1}, \\ & && z_j \geq x_j - x_j^0, & j = 1, \dots, n, \\ & && z_j \geq x_j^0 - x_j, & j = 1, \dots, n, \\ & && z_j \leq U_j y_j, & j = 1, \dots, n, \\ & && y_j \in \{0, 1\}, & j = 1, \dots, n, \\ & && x \geq 0. \end{aligned} \tag{10.11}$$

First observe that

$$z_j \geq |x_j - x_j^0| = |\Delta x_j|.$$

We choose U_j as some a priori upper bound on the amount of trading in asset j and therefore if $z_j > 0$ then $y_j = 1$ has to be the case. This implies that the transaction cost for asset j is given by

$$f_j y_j + g_j z_j.$$

Example code

The following example code demonstrates how to compute an optimal portfolio when transaction costs are included.

Listing 10.4: Code solving problem (10.11).

```
MarkowitzWithTransactionCosts <- function(
  n,          # Number of assets
  mu,         # An n-dimensional vector of expected returns
  GT,         # A matrix with n columns so (GT')*GT = covariance matrix
  x0,         # Initial holdings
  w,          # Initial cash holding
  gamma,      # Maximum risk (=std. dev) accepted
  f,          # Fixed transaction cost
  g)          # Linear part of transaction cost
{
  # Upper bound on the traded amount
  u <- w+sum(x0)

  prob <- list(sense="max")
  prob$c <- c(mu, rep(0,2*n))

  # Specify linear constraints
  # [ e' g' f' ] [ x ] = w + e'*x0
  # [ I -I 0 ] * [ z ] <= x0
  # [ I I 0 ] [ y ] >= x0
  # [ 0 I -U ] <= 0
  prob$A <- rbind(cbind(Matrix(1.0,ncol=n), t(g), t(f)),
                  cbind(Diagonal(n, 1.0), -Diagonal(n, 1.0), Matrix(0,n,n)),
                  cbind(Diagonal(n, 1.0), Diagonal(n, 1.0), Matrix(0,n,n)),
                  cbind(Matrix(0,n,n), Diagonal(n, 1.0), Diagonal(n, -u)))
  prob$bc <- rbind(blc=c(w+sum(x0), rep(-Inf,n), x0, rep(-Inf,n)),
                  buc=c(w+sum(x0), x0, rep(Inf,n), rep(0.0,n)))
  # No shortselling and the linear bound 0 <= y <= 1
  prob$bx <- rbind(blx=c(rep(0.0,n), rep(-Inf,n), rep(0.0,n)),
                  bux=c(rep(Inf,n), rep(Inf, n), rep(1.0,n)))

  # Specify the affine conic constraints for risk
  prob$F <- rbind(
    Matrix(0.0,nrow=1,ncol=3*n),
    cbind(GT, Matrix(0.0,nrow=n,ncol=2*n))
  )
  prob$g <- c(gamma,rep(0,n))
  prob$cones <- matrix(list("QUAD", 1+n, NULL), nrow=3, ncol=1)
  rownames(prob$cones) <- c("type", "dim", "conepar")

  # Demand y to be integer (hence binary)
  prob$intsub <- (2*n+1):(3*n);

  # Solve the problem
  r <- mosek(prob,list(verbose=1))
  stopifnot(identical(r$response$code, 0))

  # Return the solution
  x <- r$sol$int$xx[1:n]
  z <- r$sol$int$xx[(n+1):(2*n)]
  y <- r$sol$int$xx[(2*n+1):(3*n)]
  list(expret=drop(mu %*% x), stddev=gamma, cost = drop(f %*% y)+drop(g %*% z), x=x)
}
```

10.1.7 Cardinality constraints

Another method to reduce costs involved with processing transactions is to only change positions in a small number of assets. In other words, at most k of the differences $|\Delta x_j| = |x_j - x_j^0|$ are allowed to be non-zero, where k is (much) smaller than the total number of assets n .

This type of constraint can be again modeled by introducing a binary variable y_j which indicates if $\Delta x_j \neq 0$ and bounding the sum of y_j . The basic Markowitz model then gets updated as follows:

$$\begin{aligned}
 & \text{maximize} && \mu^T x \\
 & \text{subject to} && e^T x = w + e^T x^0, \\
 & && (\gamma, G^T x) \in \mathcal{Q}^{n+1}, \\
 & && z_j \geq x_j - x_j^0, \quad j = 1, \dots, n, \\
 & && z_j \geq x_j^0 - x_j, \quad j = 1, \dots, n, \\
 & && z_j \leq U_j y_j, \quad j = 1, \dots, n, \\
 & && y_j \in \{0, 1\}, \quad j = 1, \dots, n, \\
 & && e^T y \leq k, \\
 & && x \geq 0,
 \end{aligned} \tag{10.12}$$

where U_j is some a priori chosen upper bound on the amount of trading in asset j .

Example code

The following example code demonstrates how to compute an optimal portfolio with cardinality bounds.

Listing 10.5: Code solving problem (10.12).

```

MarkowitzWithCardinality <- function(
  n,          # Number of assets
  mu,         # An n-dimensional vector of expected returns
  GT,         # A matrix with n columns so (GT')*GT = covariance matrix
  x0,         # Initial holdings
  w,          # Initial cash holding
  gamma,      # Maximum risk (=std. dev) accepted
  k)          # Cardinality bound
{
  # Upper bound on the traded amount
  u <- w+sum(x0)

  prob <- list(sense="max")
  prob$c <- c(mu, rep(0,2*n))

  # Specify linear constraints
  # [ e'  0  0 ]          = w + e'*x0
  # [ I  -I  0 ]  [ x ] <= x0
  # [ I  I  0 ] * [ z ] >= x0
  # [ 0  I  -U ]  [ y ] <= 0
  # [ 0  0  e' ]          <= k
  prob$A <- rbind(cbind(Matrix(1.0,ncol=n), Matrix(0.0,ncol=2*n)),
                  cbind(Diagonal(n, 1.0), -Diagonal(n, 1.0), Matrix(0,n,n)),
                  cbind(Diagonal(n, 1.0), Diagonal(n, 1.0), Matrix(0,n,n)),
                  cbind(Matrix(0,n,n), Diagonal(n, 1.0), Diagonal(n, -u)),
                  cbind(Matrix(0.0,ncol=2*n), Matrix(1.0,ncol=n)))
  prob$bc <- rbind(blc=c(w+sum(x0), rep(-Inf,n), x0, rep(-Inf,n), 0.0),
                  buc=c(w+sum(x0), x0, rep(Inf,n), rep(0.0,n), k))
  # No shortselling and the linear bound 0 <= y <= 1
  prob$bx <- rbind(blx=c(rep(0.0,n), rep(-Inf,n), rep(0.0,n)),
                  bux=c(rep(Inf,n), rep(Inf, n), rep(1.0,n)))

```

(continues on next page)

(continued from previous page)

```
# Specify the affine conic constraints for risk
prob$F <- rbind(
  Matrix(0.0,nrow=1,ncol=3*n),
  cbind(GT, Matrix(0.0,nrow=n,ncol=2*n))
)
prob$g <- c(gamma,rep(0,n))
prob$cones <- matrix(list("QUAD", 1+n, NULL), nrow=3, ncol=1)
rownames(prob$cones) <- c("type", "dim", "conepar")

# Demand y to be integer (hence binary)
prob$intsub <- (2*n+1):(3*n);

# Solve the problem
r <- mosek(prob,list(verbose=1))
stopifnot(identical(r$response$code, 0))

# Return the solution
x <- r$sol$int$xx[1:n]
list(card=k, expret=drop(mu %*% x), x=x)
}
```

If we solve our running example with $k = 1, 2, 3$ then we get the following solutions, with increasing expected returns:

Bound:	1	Expected return:	0,0627	Solution:	0,0000	0,0000	1,0000
Bound:	2	Expected return:	0,0669	Solution:	0,0939	0,0000	0,9061
Bound:	3	Expected return:	0,0685	Solution:	0,1010	0,1156	0,7834

10.2 Least Squares and Other Norm Minimization Problems

A frequently occurring problem in statistics and in many other areas of science is a norm minimization problem

$$\begin{aligned} & \text{minimize} && \|Fx - g\|, \\ & \text{subject to} && Ax = b, \end{aligned} \tag{10.13}$$

where $x \in \mathbb{R}^n$ and of course we can allow other types of constraints. The objective can involve various norms: infinity norm, 1-norm, 2-norm, p -norms and so on. For instance the most popular case of the 2-norm corresponds to the least squares linear regression, since it is equivalent to minimization of $\|Fx - g\|_2^2$.

10.2.1 Least squares, 2-norm

In the case of the 2-norm we specify the problem directly in conic quadratic form

$$\begin{aligned} & \text{minimize} && t, \\ & \text{subject to} && (t, Fx - g) \in \mathcal{Q}^{k+1}, \\ & && Ax = b. \end{aligned} \tag{10.14}$$

The first constraint of the problem can be represented as an affine conic constraint. This leads to the following model.

Listing 10.6: Script solving problem (10.14)

```
# Least squares regression
# minimize ||Fx-g||_2
norm_lse <- function(F,g,A,b)
{
  n <- dim(F)[2]
  k <- length(g)
```

(continues on next page)

(continued from previous page)

```
m <- dim(A)[1]

# Linear constraints in [x; t]
prob <- list(sense="min")
prob$A <- cbind(A, rep(0, m))
prob$bx <- rbind(rep(-Inf, n+1), rep(Inf, n+1))
prob$bc <- rbind(b, b)
prob$c <- c(rep(0, n), 1)

# Affine conic constraint
prob$F <- rbind( c(rep(0,n), 1),
                 cbind(F, rep(0, k)) )
prob$g <- c(0, -g)
prob$cones <- matrix(list("QUAD", k+1, NULL))
rownames(prob$cones) <- c("type", "dim", "conepar")

# Solve
r <- mosek(prob, list(verbose=1))
stopifnot(identical(r$response$code, 0))
r$sol$itr$xx[1:n]
}
```

10.2.2 Ridge regularisation

Regularisation is classically applied to reduce the impact of outliers and to control overfitting. In the conic version of *ridge* (*Tychonov*) *regression* we consider the problem

$$\begin{aligned} & \text{minimize} && \|Fx - g\|_2 + \gamma\|x\|_2, \\ & \text{subject to} && Ax = b, \end{aligned} \tag{10.15}$$

which can be written explicitly as

$$\begin{aligned} & \text{minimize} && t_1 + \gamma t_2, \\ & \text{subject to} && (t_1, Fx - g) \in \mathcal{Q}^{k+1}, \\ & && (t_2, x) \in \mathcal{Q}^{n+1}, \\ & && Ax = b. \end{aligned} \tag{10.16}$$

The implementation is a small extension of that from the previous section.

Listing 10.7: Script solving problem (10.16)

```
# Least squares regression with regularization
# minimize ||Fx-g||_2 + gamma*||x||_2
norm_lse_reg <- function(F,g,A,b,gamma)
{
  n <- dim(F)[2]
  k <- length(g)
  m <- dim(A)[1]

  # Linear constraints in [x; t1; t2]
  prob <- list(sense="min")
  prob$A <- cbind(A, rep(0, m), rep(0, m))
  prob$bx <- rbind(rep(-Inf, n+2), rep(Inf, n+2))
  prob$bc <- rbind(b, b)
  prob$c <- c(rep(0, n), 1, gamma)

  # Affine conic constraint
  prob$F <- rbind( c(rep(0,n), 1, 0),
                   cbind(F, rep(0, k), rep(0, k)),
                   c(rep(0,n), 0, 1),
                   cbind(diag(n), rep(0, n), rep(0, n)))
}
```

(continues on next page)

(continued from previous page)

```
prob$g <- c(0, -g, rep(0, n+1))
prob$cones <- cbind(matrix(list("QUAD", k+1, NULL)),
                    matrix(list("QUAD", n+1, NULL)))
rownames(prob$cones) <- c("type", "dim", "conepr")

# Solve
r <- mosek(prob, list(verbose=1))
stopifnot(identical(r$response$code, 0))
r$sol$itr$xx[1:n]
}
```

Note that classically least squares problems are formulated as quadratic problems and then the objective function would be written as

$$\|Fx - g\|_2^2 + \gamma\|x\|_2^2.$$

This version can easily be obtained by replacing the quadratic cone with an appropriate rotated quadratic cone in (10.16). Then the core of the implementation would change as follows:

Listing 10.8: Script solving classical quadratic ridge regression

```
# Affine conic constraint
prob$F <- rbind(      c(rep(0,n), 1,          0),
                     c(rep(0, n+2))),
             cbind(F,      rep(0, k), rep(0, k)),
             c(rep(0,n), 0,          1),
             c(rep(0, n+2))),
             cbind(diag(n), rep(0, n), rep(0, n)))
prob$g <- c(0, 0.5, -g, 0, 0.5, rep(0, n))
prob$cones <- cbind(matrix(list("RQUAD", k+2, NULL)),
                    matrix(list("RQUAD", n+2, NULL)))
rownames(prob$cones) <- c("type", "dim", "conepr")
```

Fig. 10.2 shows the solution to a polynomial fitting problem for a few variants of least squares regression with and without ridge regularization.

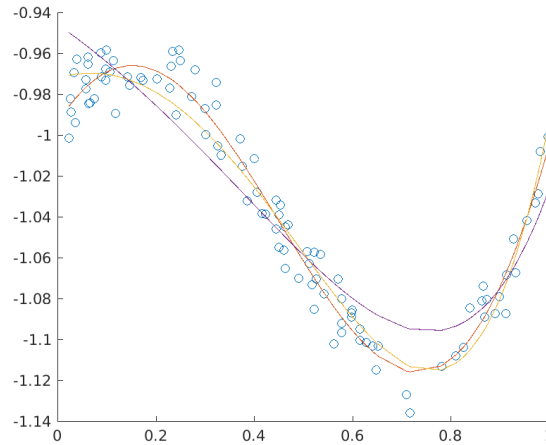


Fig. 10.2: Three fits to a dataset at various levels of regularization.

10.2.3 Lasso regularization

In *lasso* or *least absolute shrinkage and selection operator* the regularization term is the 1-norm of the solution

$$\begin{aligned} &\text{minimize} && \|Fx - g\|_2 + \gamma\|x\|_1, \\ &\text{subject to} && Ax = b. \end{aligned} \tag{10.17}$$

This variant typically tends to give preference to sparser solutions, i.e. solutions where only a few elements of x are nonzero, and therefore it is used as an efficient approximation to the cardinality constrained problem with an upper bound on the 0-norm of x . To see how it works we first implement (10.17) adding the constraint $t \geq \|x\|_1$ as a series of linear constraints

$$u_i \geq -x_i, \quad u_i \geq x_i, \quad t \geq \sum u_i,$$

so that eventually the problem becomes

$$\begin{aligned} & \text{minimize} && t_1 + \gamma t_2, \\ & \text{subject to} && u + x \geq 0, \\ & && u - x \geq 0, \\ & && t_2 - e^T u \geq 0, \\ & && Ax = b, \\ & && (t_1, Fx - g) \in \mathcal{Q}^{k+1}. \end{aligned}$$

Listing 10.9: Script solving problem (10.17)

```
# Least squares regression with lasso regularization
# minimize ||Fx-g||_2 + gamma*||x||_1
norm_lse_lasso <- function(F,g,A,b,gamma)
{
  n <- dim(F)[2]
  k <- length(g)
  m <- dim(A)[1]

  # Linear constraints in [x; u; t1; t2]
  prob <- list(sense="min")
  prob$A <- rbind(cbind(A, matrix(0, m, n+2)),
                 cbind(diag(n), diag(n), matrix(0, n, 2)),
                 cbind(-diag(n), diag(n), matrix(0, n, 2)),
                 c(rep(0, n), rep(-1, n), 0, 1))
  prob$bx <- rbind(rep(-Inf, 2*n+2), rep(Inf, 2*n+2))
  prob$bc <- rbind(c(b, rep(0, 2*n+1)), c(b, rep(Inf, 2*n+1)))
  prob$c <- c(rep(0, 2*n), 1, gamma)

  # Affine conic constraint
  prob$F <- rbind(c(rep(0, 2*n), 1, 0),
                 cbind(F, matrix(0, k, n+2)))
  prob$g <- c(0, -g)
  prob$cones <- matrix(list("QUAD", k+1, NULL))
  rownames(prob$cones) <- c("type", "dim", "conepar")

  # Solve
  r <- mosek(prob, list(verbose=1))
  stopifnot(identical(r$response$code, 0))
  r$sol$itr$xx[1:n]
}
```

The sparsity pattern of the solution of a large random regression problem can look for example as follows:

```
Lasso regularization
Gamma 0.0100 density 99% |Fx-g|_2: 54.3722
Gamma 0.1000 density 87% |Fx-g|_2: 54.3939
Gamma 0.3000 density 67% |Fx-g|_2: 54.5319
Gamma 0.6000 density 40% |Fx-g|_2: 54.8379
Gamma 0.9000 density 26% |Fx-g|_2: 55.0720
Gamma 1.3000 density 12% |Fx-g|_2: 55.1903
```

10.2.4 p-norm minimization

Now we consider the minimization of the p -norm defined for $p > 1$ as

$$\|y\|_p = \left(\sum_i |y_i|^p \right)^{1/p}. \quad (10.18)$$

We have the optimization problem:

$$\begin{aligned} & \text{minimize} && \|Fx - g\|_p, \\ & \text{subject to} && Ax = b. \end{aligned} \quad (10.19)$$

Increasing the value of p forces stronger penalization of outliers as ultimately, when $p \rightarrow \infty$, the p -norm $\|y\|_p$ converges to the infinity norm $\|y\|_\infty$ of y . According to the [Modeling Cookbook](#) the p -norm bound $t \geq \|Fx - g\|_p$ can be added to the model using a sequence of three-dimensional power cones and we obtain an equivalent problem

$$\begin{aligned} & \text{minimize} && t \\ & \text{subject to} && (r_i, t, (Fx - g)_i) \in \mathcal{P}_3^{1/p, 1-1/p}, \\ & && e^T r = t, \\ & && Ax = b. \end{aligned} \quad (10.20)$$

The power cones can be added one by one to the structure representing affine conic constraints. Each power cone will require one r_i , one copy of t and one row from F and g . An alternative solution is to create the vector

$$[r_1; \dots; r_k; t; \dots; t; Fx - g]$$

and then reshuffle its elements into

$$[r_1; t; F_1x - g_1; \dots; r_k; t; F_kx - g_k]$$

using an appropriate permutation matrix. This approach is demonstrated in the code below.

Listing 10.10: Script solving problem (10.20)

```
# P-norm minimization
# minimize \|Fx-g\|_p
norm_p_norm <- function(F,g,A,b,p)
{
  n <- dim(F)[2]
  k <- length(g)
  m <- dim(A)[1]

  # Linear constraints in [x; r; t]
  prob <- list(sense="min")
  prob$A <- rbind(cbind(A, matrix(0, m, k+1)),
                  c(rep(0, n), rep(1, k), -1))
  prob$bx <- rbind(rep(-Inf, n+k+1), rep(Inf, n+k+1))
  prob$bc <- rbind(c(b, 0), c(b, 0))
  prob$c <- c(rep(0, n+k), 1)

  # Permutation matrix which picks triples (r_i, t, F_ix-g_i)
  M <- cbind(sparseMatrix(seq(1,3*k,3), seq(1,k), x=rep(1,k), dims=c(3*k,k)),
             sparseMatrix(seq(2,3*k,3), seq(1,k), x=rep(1,k), dims=c(3*k,k)),
             sparseMatrix(seq(3,3*k,3), seq(1,k), x=rep(1,k), dims=c(3*k,k)))

  # Affine conic constraint
  prob$F <- M %*% rbind(cbind(matrix(0, k, n), diag(k), rep(0, k)),
                       cbind(matrix(0, k, n+k), rep(1, k)),
                       cbind(F, matrix(0, k, k+1)))
```

(continues on next page)

(continued from previous page)

```
prob$g <- as.numeric(M %*% c(rep(0, 2*k), -g))
prob$cones <- matrix(list("PPOW", 3, c(1, p-1)), nrow=3, ncol=k)
rownames(prob$cones) <- c("type", "dim", "conepr")

# Solve
r <- mosek(prob, list(verbose=1))
stopifnot(identical(r$sol$itr$prosta, "PRIMAL_AND_DUAL_FEASIBLE"))
r$sol$itr$xx[1:n]
}
```

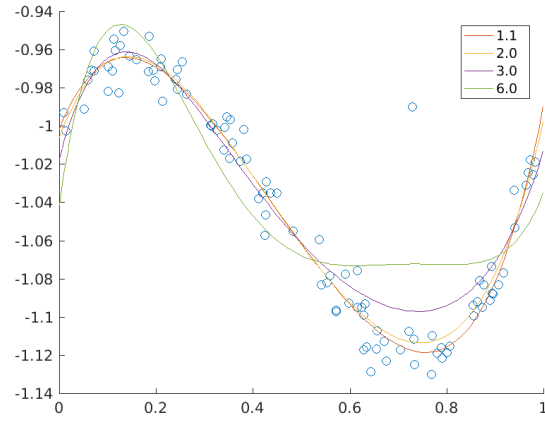


Fig. 10.3: p -norm minimizing fits of a polynomial of degree at most 5 to the data for various values of p .

Chapter 11

Problem Formulation and Solutions

In this chapter we will discuss the following issues:

- The formal, mathematical formulations of the problem types that **MOSEK** can solve and their duals.
- The solution information produced by **MOSEK**.
- The infeasibility certificate produced by **MOSEK** if the problem is infeasible.

For the underlying mathematical concepts, derivations and proofs see the [Modeling Cookbook](#) or any book on convex optimization. This chapter explains how the related data is organized specifically within the **MOSEK** API.

11.1 Linear Optimization

MOSEK accepts linear optimization problems of the form

$$\begin{array}{llllll} \text{minimize} & & c^T x + c^f & & & \\ \text{subject to} & l^c & \leq & Ax & \leq & u^c, \\ & l^x & \leq & x & \leq & u^x, \end{array} \quad (11.1)$$

where

- m is the number of constraints.
- n is the number of decision variables.
- $x \in \mathbb{R}^n$ is a vector of decision variables.
- $c \in \mathbb{R}^n$ is the linear part of the objective function.
- $c^f \in \mathbb{R}$ is a constant term in the objective
- $A \in \mathbb{R}^{m \times n}$ is the constraint matrix.
- $l^c \in \mathbb{R}^m$ is the lower limit on the activity for the constraints.
- $u^c \in \mathbb{R}^m$ is the upper limit on the activity for the constraints.
- $l^x \in \mathbb{R}^n$ is the lower limit on the activity for the variables.
- $u^x \in \mathbb{R}^n$ is the upper limit on the activity for the variables.

Lower and upper bounds can be infinite, or in other words the corresponding bound may be omitted.

A primal solution (x) is *(primal) feasible* if it satisfies all constraints in (11.1). If (11.1) has at least one primal feasible solution, then (11.1) is said to be (primal) feasible. In case (11.1) does not have a feasible solution, the problem is said to be *(primal) infeasible*.

11.1.1 Duality for Linear Optimization

Corresponding to the primal problem (11.1), there is a dual problem

$$\begin{aligned} & \text{maximize} && (l^c)^T s_l^c - (u^c)^T s_u^c + (l^x)^T s_l^x - (u^x)^T s_u^x + c^f \\ & && A^T y + s_l^c - s_u^c = c, \\ & \text{subject to} && -y + s_l^c - s_u^c = 0, \\ & && s_l^c, s_u^c, s_l^x, s_u^x \geq 0. \end{aligned} \quad (11.2)$$

If a bound in the primal problem is plus or minus infinity, the corresponding dual variable is fixed at 0, and we use the convention that the product of the bound value and the corresponding dual variable is 0. This is equivalent to removing variable $(s_l^x)_j$ from the dual problem. In other words:

$$l_j^x = -\infty \quad \Rightarrow \quad (s_l^x)_j = 0 \text{ and } l_j^x \cdot (s_l^x)_j = 0.$$

A solution

$$(y, s_l^c, s_u^c, s_l^x, s_u^x)$$

to the dual problem is feasible if it satisfies all the constraints in (11.2). If (11.2) has at least one feasible solution, then (11.2) is *(dual) feasible*, otherwise the problem is *(dual) infeasible*.

A solution

$$(x^*, y^*, (s_l^c)^*, (s_u^c)^*, (s_l^x)^*, (s_u^x)^*)$$

is denoted a *primal-dual feasible solution*, if (x^*) is a solution to the primal problem (11.1) and $(y^*, (s_l^c)^*, (s_u^c)^*, (s_l^x)^*, (s_u^x)^*)$ is a solution to the corresponding dual problem (11.2). We also define an auxiliary vector

$$(x^c)^* := Ax^*$$

containing the activities of linear constraints.

For a primal-dual feasible solution we define the *duality gap* as the difference between the primal and the dual objective value,

$$\begin{aligned} & c^T x^* + c^f - \{ (l^c)^T (s_l^c)^* - (u^c)^T (s_u^c)^* + (l^x)^T (s_l^x)^* - (u^x)^T (s_u^x)^* + c^f \} \\ &= \sum_{i=0}^{m-1} [(s_l^c)^*_i ((x_i^c)^* - l_i^c) + (s_u^c)^*_i (u_i^c - (x_i^c)^*)] \\ &+ \sum_{j=0}^{n-1} [(s_l^x)^*_j (x_j^x - l_j^x) + (s_u^x)^*_j (u_j^x - x_j^*)] \geq 0 \end{aligned} \quad (11.3)$$

where the first relation can be obtained by transposing and multiplying the dual constraints (11.2) by x^* and $(x^c)^*$ respectively, and the second relation comes from the fact that each term in each sum is nonnegative. It follows that the primal objective will always be greater than or equal to the dual objective.

It is well-known that a linear optimization problem has an optimal solution if and only if there exist feasible primal-dual solution so that the duality gap is zero, or, equivalently, that the *complementarity conditions*

$$\begin{aligned} (s_l^c)^*_i ((x_i^c)^* - l_i^c) &= 0, & i = 0, \dots, m-1, \\ (s_u^c)^*_i (u_i^c - (x_i^c)^*) &= 0, & i = 0, \dots, m-1, \\ (s_l^x)^*_j (x_j^x - l_j^x) &= 0, & j = 0, \dots, n-1, \\ (s_u^x)^*_j (u_j^x - x_j^*) &= 0, & j = 0, \dots, n-1, \end{aligned}$$

are satisfied.

If (11.1) has an optimal solution and **MOSEK** solves the problem successfully, both the primal and dual solution are reported, including a status indicating the exact state of the solution.

11.1.2 Infeasibility for Linear Optimization

Primal Infeasible Problems

If the problem (11.1) is infeasible (has no feasible solution), **MOSEK** will report a certificate of primal infeasibility: The dual solution reported is the certificate of infeasibility, and the primal solution is undefined.

A certificate of primal infeasibility is a feasible solution to the modified dual problem

$$\begin{aligned}
& \text{maximize} && (l^c)^T s_l^c - (u^c)^T s_u^c + (l^x)^T s_l^x - (u^x)^T s_u^x \\
& \text{subject to} && A^T y + s_l^x - s_u^x = 0, \\
& && -y + s_l^c - s_u^c = 0, \\
& && s_l^c, s_u^c, s_l^x, s_u^x \geq 0,
\end{aligned} \tag{11.4}$$

such that the objective value is strictly positive, i.e. a solution

$$(y^*, (s_l^c)^*, (s_u^c)^*, (s_l^x)^*, (s_u^x)^*)$$

to (11.4) so that

$$(l^c)^T (s_l^c)^* - (u^c)^T (s_u^c)^* + (l^x)^T (s_l^x)^* - (u^x)^T (s_u^x)^* > 0.$$

Such a solution implies that (11.4) is unbounded, and that (11.1) is infeasible.

Dual Infeasible Problems

If the problem (11.2) is infeasible (has no feasible solution), **MOSEK** will report a certificate of dual infeasibility: The primal solution reported is the certificate of infeasibility, and the dual solution is undefined.

A certificate of dual infeasibility is a feasible solution to the modified primal problem

$$\begin{aligned}
& \text{minimize} && c^T x \\
& \text{subject to} && \hat{l}^c \leq Ax \leq \hat{u}^c, \\
& && \hat{l}^x \leq x \leq \hat{u}^x,
\end{aligned} \tag{11.5}$$

where

$$\hat{l}_i^c = \begin{cases} 0 & \text{if } l_i^c > -\infty, \\ -\infty & \text{otherwise,} \end{cases} \quad \text{and} \quad \hat{u}_i^c := \begin{cases} 0 & \text{if } u_i^c < \infty, \\ \infty & \text{otherwise,} \end{cases}$$

and

$$\hat{l}_j^x = \begin{cases} 0 & \text{if } l_j^x > -\infty, \\ -\infty & \text{otherwise,} \end{cases} \quad \text{and} \quad \hat{u}_j^x := \begin{cases} 0 & \text{if } u_j^x < \infty, \\ \infty & \text{otherwise,} \end{cases}$$

such that

$$c^T x < 0.$$

Such a solution implies that (11.5) is unbounded, and that (11.2) is infeasible.

In case that both the primal problem (11.1) and the dual problem (11.2) are infeasible, **MOSEK** will report only one of the two possible certificates — which one is not defined (**MOSEK** returns the first certificate found).

11.1.3 Minimalization vs. Maximalization

When the objective sense of problem (11.1) is maximization, i.e.

$$\begin{aligned}
& \text{maximize} && c^T x + c^f \\
& \text{subject to} && l^c \leq Ax \leq u^c, \\
& && l^x \leq x \leq u^x,
\end{aligned}$$

the objective sense of the dual problem changes to minimization, and the domain of all dual variables changes sign in comparison to (11.2). The dual problem thus takes the form

$$\begin{aligned}
& \text{minimize} && (l^c)^T s_l^c - (u^c)^T s_u^c + (l^x)^T s_l^x - (u^x)^T s_u^x + c^f \\
& \text{subject to} && A^T y + s_l^x - s_u^x = c, \\
& && -y + s_l^c - s_u^c = 0, \\
& && s_l^c, s_u^c, s_l^x, s_u^x \leq 0.
\end{aligned}$$

This means that the duality gap, defined in (11.3) as the primal minus the dual objective value, becomes nonpositive. It follows that the dual objective will always be greater than or equal to the primal objective. The primal infeasibility certificate will be reported by **MOSEK** as a solution to the system

$$\begin{aligned} A^T y + s_l^x - s_u^x &= 0, \\ -y + s_l^c - s_u^c &= 0, \\ s_l^c, s_u^c, s_l^x, s_u^x &\leq 0, \end{aligned} \tag{11.6}$$

such that the objective value is strictly negative

$$(l^c)^T (s_l^c)^* - (u^c)^T (s_u^c)^* + (l^x)^T (s_l^x)^* - (u^x)^T (s_u^x)^* < 0.$$

Similarly, the certificate of dual infeasibility is an x satisfying the requirements of (11.5) such that $c^T x > 0$.

11.2 Conic Optimization

Conic optimization is an extension of linear optimization (see Sec. 11.1) allowing conic domains to be specified for subsets of the problem variables. A conic optimization problem to be solved by **MOSEK** can be written as

$$\begin{aligned} &\text{minimize} && c^T x + c^f \\ &\text{subject to} && l^c \leq Ax \leq u^c, \\ & && l^x \leq x \leq u^x, \\ & && x \in \mathcal{K}, \end{aligned} \tag{11.7}$$

where

- m is the number of constraints.
- n is the number of decision variables.
- $x \in \mathbb{R}^n$ is a vector of decision variables.
- $c \in \mathbb{R}^n$ is the linear part of the objective function.
- $c^f \in \mathbb{R}$ is a constant term in the objective
- $A \in \mathbb{R}^{m \times n}$ is the constraint matrix.
- $l^c \in \mathbb{R}^m$ is the lower limit on the activity for the constraints.
- $u^c \in \mathbb{R}^m$ is the upper limit on the activity for the constraints.
- $l^x \in \mathbb{R}^n$ is the lower limit on the activity for the variables.
- $u^x \in \mathbb{R}^n$ is the upper limit on the activity for the variables.

Lower and upper bounds can be infinite, or in other words the corresponding bound may be omitted.

The set \mathcal{K} is a Cartesian product of convex cones, namely $\mathcal{K} = \mathcal{K}_1 \times \cdots \times \mathcal{K}_p$. Having the domain restriction $x \in \mathcal{K}$, is thus equivalent to

$$x^t \in \mathcal{K}_t \subseteq \mathbb{R}^{n_t},$$

where $x = (x^1, \dots, x^p)$ is a partition of the problem variables. Please note that the n -dimensional Euclidean space \mathbb{R}^n is a cone itself, so simple linear variables are still allowed. The user only needs to specify subsets of variables which belong to non-trivial cones.

In this section we discuss the formulations which apply to the following cones supported by **MOSEK**:

- The set \mathbb{R}^n .
- The zero cone $\{(0, \dots, 0)\}$.
- Quadratic cone

$$\mathcal{Q}^n = \left\{ x \in \mathbb{R}^n : x_1 \geq \sqrt{\sum_{j=2}^n x_j^2} \right\}.$$

- Rotated quadratic cone

$$\mathcal{Q}_r^n = \left\{ x \in \mathbb{R}^n : 2x_1x_2 \geq \sum_{j=3}^n x_j^2, \quad x_1 \geq 0, \quad x_2 \geq 0 \right\}.$$

- Primal exponential cone

$$K_{\text{exp}} = \{x \in \mathbb{R}^3 : x_1 \geq x_2 \exp(x_3/x_2), \quad x_1, x_2 \geq 0\}$$

as well as its dual

$$K_{\text{exp}}^* = \{x \in \mathbb{R}^3 : x_1 \geq -x_3 e^{-1} \exp(x_2/x_3), \quad x_3 \leq 0, x_1 \geq 0\}.$$

- Primal power cone (with parameter $0 < \alpha < 1$)

$$\mathcal{P}_n^{\alpha, 1-\alpha} = \left\{ x \in \mathbb{R}^n : x_1^\alpha x_2^{1-\alpha} \geq \sqrt{\sum_{j=3}^n x_j^2}, \quad x_1, x_2 \geq 0 \right\}$$

as well as its dual

$$(\mathcal{P}_n^{\alpha, 1-\alpha})^* = \left\{ x \in \mathbb{R}^n : \left(\frac{x_1}{\alpha}\right)^\alpha \left(\frac{x_2}{1-\alpha}\right)^{1-\alpha} \geq \sqrt{\sum_{j=3}^n x_j^2}, \quad x_1, x_2 \geq 0 \right\}.$$

MOSEK supports also the cone of positive semidefinite matrices. Since that is handled through a separate interface, we discuss it in [Sec. 11.3](#).

11.2.1 Duality for Conic Optimization

Corresponding to the primal problem (11.7), there is a dual problem

$$\begin{aligned} & \text{maximize} && (l^c)^T s_l^c - (u^c)^T s_u^c + (l^x)^T s_l^x - (u^x)^T s_u^x + c^f \\ & \text{subject to} && A^T y + s_l^x - s_u^x + s_n^x = c \\ & && -y + s_l^c - s_u^c = 0, \\ & && s_l^c, s_u^c, s_l^x, s_u^x \geq 0, \\ & && s_n^x \in \mathcal{K}^*, \end{aligned} \tag{11.8}$$

where the dual cone \mathcal{K}^* is a Cartesian product of the cones dual to \mathcal{K}_t . In practice this means that s_n^x has one entry for each entry in x . Please note that the dual problem of the dual problem is identical to the original primal problem.

If a bound in the primal problem is plus or minus infinity, the corresponding dual variable is fixed at 0, and we use the convention that the product of the bound value and the corresponding dual variable is 0. This is equivalent to removing variable $(s_l^x)_j$ from the dual problem. In other words:

$$l_j^x = -\infty \quad \Rightarrow \quad (s_l^x)_j = 0 \text{ and } l_j^x \cdot (s_l^x)_j = 0.$$

A solution

$$(y, s_l^c, s_u^c, s_l^x, s_u^x, s_n^x)$$

to the dual problem is feasible if it satisfies all the constraints in (11.8). If (11.8) has at least one feasible solution, then (11.8) is *(dual) feasible*, otherwise the problem is *(dual) infeasible*.

A solution

$$(x^*, y^*, (s_l^c)^*, (s_u^c)^*, (s_l^x)^*, (s_u^x)^*, (s_n^x)^*)$$

is denoted a *primal-dual feasible solution*, if (x^*) is a solution to the primal problem (11.7) and $(y^*, (s_l^c)^*, (s_u^c)^*, (s_l^x)^*, (s_u^x)^*, (s_n^x)^*)$ is a solution to the corresponding dual problem (11.8). We also define an auxiliary vector

$$(x^c)^* := Ax^*$$

containing the activities of linear constraints.

For a primal-dual feasible solution we define the *duality gap* as the difference between the primal and the dual objective value,

$$\begin{aligned} & c^T x^* + c^f - \{ (l^c)^T (s_l^c)^* - (u^c)^T (s_u^c)^* + (l^x)^T (s_l^x)^* - (u^x)^T (s_u^x)^* + c^f \} \\ &= \sum_{i=0}^{m-1} [(s_l^c)^*_i ((x_i^c)^* - l_i^c) + (s_u^c)^*_i (u_i^c - (x_i^c)^*)] \\ &+ \sum_{j=0}^{n-1} [(s_l^x)^*_j (x_j - l_j^x) + (s_u^x)^*_j (u_j^x - x_j^*)] + \sum_{j=0}^{n-1} (s_n^x)^*_j x_j^* \geq 0 \end{aligned} \quad (11.9)$$

where the first relation can be obtained by transposing and multiplying the dual constraints (11.2) by x^* and $(x^c)^*$ respectively, and the second relation comes from the fact that each term in each sum is nonnegative. It follows that the primal objective will always be greater than or equal to the dual objective.

It is well-known that, under some non-degeneracy assumptions that exclude ill-posed cases, a conic optimization problem has an optimal solution if and only if there exist feasible primal-dual solution so that the duality gap is zero, or, equivalently, that the *complementarity conditions*

$$\begin{aligned} (s_l^c)^*_i ((x_i^c)^* - l_i^c) &= 0, & i = 0, \dots, m-1, \\ (s_u^c)^*_i (u_i^c - (x_i^c)^*) &= 0, & i = 0, \dots, m-1, \\ (s_l^x)^*_j (x_j^* - l_j^x) &= 0, & j = 0, \dots, n-1, \\ (s_u^x)^*_j (u_j^x - x_j^*) &= 0, & j = 0, \dots, n-1, \\ \sum_{j=0}^{n-1} (s_n^x)^*_j x_j^* &= 0. \end{aligned} \quad (11.10)$$

are satisfied.

If (11.7) has an optimal solution and **MOSEK** solves the problem successfully, both the primal and dual solution are reported, including a status indicating the exact state of the solution.

11.2.2 Infeasibility for Conic Optimization

Primal Infeasible Problems

If the problem (11.7) is infeasible (has no feasible solution), **MOSEK** will report a certificate of primal infeasibility: The dual solution reported is the certificate of infeasibility, and the primal solution is undefined.

A certificate of primal infeasibility is a feasible solution to the modified dual problem

$$\begin{aligned} & \text{maximize} && (l^c)^T s_l^c - (u^c)^T s_u^c + (l^x)^T s_l^x - (u^x)^T s_u^x \\ & \text{subject to} && \\ & && A^T y + s_l^x - s_u^x + s_n^x = 0, \\ & && -y + s_l^c - s_u^c = 0, \\ & && s_l^c, s_u^c, s_l^x, s_u^x \geq 0, \\ & && s_n^x \in \mathcal{K}^*, \end{aligned} \quad (11.11)$$

such that the objective value is strictly positive, i.e. a solution

$$(y^*, (s_l^c)^*, (s_u^c)^*, (s_l^x)^*, (s_u^x)^*, (s_n^x)^*)$$

to (11.11) so that

$$(l^c)^T (s_l^c)^* - (u^c)^T (s_u^c)^* + (l^x)^T (s_l^x)^* - (u^x)^T (s_u^x)^* > 0.$$

Such a solution implies that (11.11) is unbounded, and that (11.7) is infeasible.

Dual Infeasible Problems

If the problem (11.8) is infeasible (has no feasible solution), **MOSEK** will report a certificate of dual infeasibility: The primal solution reported is the certificate of infeasibility, and the dual solution is undefined.

A certificate of dual infeasibility is a feasible solution to the modified primal problem

$$\begin{aligned} & \text{minimize} && c^T x \\ & \text{subject to} && \hat{l}^c \leq Ax \leq \hat{u}^c, \\ & && \hat{l}^x \leq x \leq \hat{u}^x, \\ & && x \in K, \end{aligned} \quad (11.12)$$

where

$$\hat{l}_i^c = \begin{cases} 0 & \text{if } l_i^c > -\infty, \\ -\infty & \text{otherwise,} \end{cases} \quad \text{and} \quad \hat{u}_i^c := \begin{cases} 0 & \text{if } u_i^c < \infty, \\ \infty & \text{otherwise,} \end{cases} \quad (11.13)$$

and

$$\hat{l}_j^x = \begin{cases} 0 & \text{if } l_j^x > -\infty, \\ -\infty & \text{otherwise,} \end{cases} \quad \text{and} \quad \hat{u}_j^x := \begin{cases} 0 & \text{if } u_j^x < \infty, \\ \infty & \text{otherwise,} \end{cases} \quad (11.14)$$

such that

$$c^T x < 0.$$

Such a solution implies that (11.12) is unbounded, and that (11.8) is infeasible.

In case that both the primal problem (11.7) and the dual problem (11.8) are infeasible, **MOSEK** will report only one of the two possible certificates — which one is not defined (**MOSEK** returns the first certificate found).

11.2.3 Minimalization vs. Maximalization

When the objective sense of problem (11.7) is maximization, i.e.

$$\begin{aligned} & \text{maximize} && c^T x + c^f \\ & \text{subject to} && l^c \leq Ax \leq u^c, \\ & && l^x \leq x \leq u^x, \\ & && x \in \mathcal{K}, \end{aligned}$$

the objective sense of the dual problem changes to minimization, and the domain of all dual variables changes sign in comparison to (11.2). The dual problem thus takes the form

$$\begin{aligned} & \text{minimize} && (l^c)^T s_l^c - (u^c)^T s_u^c + (l^x)^T s_l^x - (u^x)^T s_u^x + c^f \\ & \text{subject to} && A^T y + s_l^x - s_u^x + s_n^x = c, \\ & && -y + s_l^c - s_u^c = 0, \\ & && s_l^c, s_u^c, s_l^x, s_u^x \leq 0, \\ & && -s_n^x \in \mathcal{K}^* \end{aligned}$$

This means that the duality gap, defined in (11.9) as the primal minus the dual objective value, becomes nonpositive. It follows that the dual objective will always be greater than or equal to the primal objective. The primal infeasibility certificate will be reported by **MOSEK** as a solution to the system

$$\begin{aligned} & A^T y + s_l^x - s_u^x + s_n^x = 0, \\ & -y + s_l^c - s_u^c = 0, \\ & s_l^c, s_u^c, s_l^x, s_u^x \leq 0, \\ & -s_n^x \in \mathcal{K}^* \end{aligned} \quad (11.15)$$

such that the objective value is strictly negative

$$(l^c)^T (s_l^c)^* - (u^c)^T (s_u^c)^* + (l^x)^T (s_l^x)^* - (u^x)^T (s_u^x)^* < 0.$$

Similarly, the certificate of dual infeasibility is an x satisfying the requirements of (11.12) such that $c^T x > 0$.

such that the objective value is strictly positive.

Similarly, a dual infeasibility certificate (11.12) is a feasible solution to

$$\begin{aligned}
& \text{minimize} && \sum_{j=0}^{n-1} c_j x_j + \sum_{j=0}^{p-1} \langle \bar{C}_j, \bar{X}_j \rangle \\
& \text{subject to} && \hat{l}_i^c \leq \sum_{j=0}^{n-1} a_{ij} x_j + \sum_{j=0}^{p-1} \langle \bar{A}_{ij}, \bar{X}_j \rangle \leq \hat{u}_i^c, \quad i = 0, \dots, m-1 \\
& && \hat{l}_j^x \leq x_j \leq \hat{u}_j^x, \quad j = 0, \dots, n-1 \\
& && x \in \mathcal{K}, \\
& && \bar{X}_j \in \mathcal{S}_+^{r_j}, \quad j = 0, \dots, p-1
\end{aligned} \tag{11.21}$$

where the modified bounds are as in (11.13) and (11.14) and the objective value is strictly negative.

11.4 Quadratic and Quadratically Constrained Optimization

A convex quadratic and quadratically constrained optimization problem has the form

$$\begin{aligned}
& \text{minimize} && \frac{1}{2} x^T Q^o x + c^T x + c^f \\
& \text{subject to} && l_k^c \leq \frac{1}{2} x^T Q^k x + \sum_{j=0}^{n-1} a_{kj} x_j \leq u_k^c, \quad k = 0, \dots, m-1, \\
& && l_j^x \leq x_j \leq u_j^x, \quad j = 0, \dots, n-1,
\end{aligned} \tag{11.22}$$

where all variables and bounds have the same meaning as for linear problems (see Sec. 11.1) and Q^o and all Q^k are symmetric matrices. Moreover, for convexity, Q^o must be a positive semidefinite matrix and Q^k must satisfy

$$\begin{aligned}
-\infty < l_k^c &\Rightarrow Q^k \text{ is negative semidefinite,} \\
u_k^c < \infty &\Rightarrow Q^k \text{ is positive semidefinite,} \\
-\infty < l_k^c \leq u_k^c < \infty &\Rightarrow Q^k = 0.
\end{aligned}$$

The convexity requirement is very important and **MOSEK** checks whether it is fulfilled.

11.4.1 A Recommendation

Any convex quadratic optimization problem can be reformulated as a conic quadratic optimization problem, see [Modeling Cookbook](#) and [And13]. In fact **MOSEK** does such conversion internally as a part of the solution process for the following reasons:

- the conic optimizer is numerically more robust than the one for quadratic problems.
- the conic optimizer is usually faster because quadratic cones are simpler than quadratic functions, even though the conic reformulation usually has more constraints and variables than the original quadratic formulation.
- it is easy to dualize the conic formulation if deemed worthwhile potentially leading to (huge) computational savings.

However, instead of relying on the automatic reformulation we recommend to formulate the problem as a conic problem from scratch because:

- it saves the computational overhead of the reformulation including the convexity check. A conic problem is convex by construction and hence no convexity check is needed for conic problems.
- usually the modeler can do a better reformulation than the automatic method because the modeler can exploit the knowledge of the problem at hand.

To summarize we recommend to formulate quadratic problems and in particular quadratically constrained problems directly in conic form.

11.4.2 Duality for Quadratic and Quadratically Constrained Optimization

The dual problem corresponding to the quadratic and quadratically constrained optimization problem (11.22) is given by

$$\begin{aligned}
& \text{maximize} && (l^c)^T s_l^c - (u^c)^T s_u^c + (l^x)^T s_l^x - (u^x)^T s_u^x + \frac{1}{2} x^T \left\{ \sum_{k=0}^{m-1} y_k Q^k - Q^o \right\} x + c^f \\
& \text{subject to} && A^T y + s_l^x - s_u^x + \left\{ \sum_{k=0}^{m-1} y_k Q^k - Q^o \right\} x = c, \\
& && -y + s_l^c - s_u^c = 0, \\
& && s_l^c, s_u^c, s_l^x, s_u^x \geq 0.
\end{aligned} \tag{11.23}$$

The dual problem is related to the dual problem for linear optimization (see Sec. 11.1.1), but depends on the variable x which in general can not be eliminated. In the solutions reported by **MOSEK**, the value of x is the same for the primal problem (11.22) and the dual problem (11.23).

11.4.3 Infeasibility for Quadratic Optimization

In case **MOSEK** finds a problem to be infeasible it reports a certificate of infeasibility. We write them out explicitly for quadratic problems, that is when $Q^k = 0$ for all k and quadratic terms appear only in the objective Q^o . In this case the constraints both in the primal and dual problem are linear, and **MOSEK** produces for them the same infeasibility certificate as for linear problems.

The certificate of primal infeasibility is a solution to the problem (11.4) such that the objective value is strictly positive.

The certificate of dual infeasibility is a solution to the problem (11.5) together with an additional constraint

$$Q^o x = 0$$

such that the objective value is strictly negative.

11.5 Affine Conic Constraints

Rmosek package allows conic problems to be specified in another format, namely with *affine conic constraints*. A conic problem can thus be specified as:

$$\begin{aligned}
& \text{minimize} && \sum_{j=1}^n c_j x_j + \sum_{j=1}^p \langle \bar{C}_j, \bar{X}_j \rangle + c^f \\
& \text{subject to} && \begin{aligned} l_i^c &\leq \sum_{j=1}^n a_{ij} x_j + \sum_{j=1}^p \langle \bar{A}_{ij}, \bar{X}_j \rangle &\leq u_i^c, & i = 1, \dots, m, \\ l_j^x &\leq x_j &\leq u_j^x, & j = 1, \dots, n, \end{aligned} \\
& && Fx + g \in \mathcal{K}, \\
& && \bar{X}_j \in \mathcal{S}_+^{r_j}, & j = 1, \dots, p
\end{aligned} \tag{11.24}$$

where all the data has the same meaning as in Sec. 11.2 and Sec. 11.3 and $F \in \mathbb{R}^{k \times n}$, $g \in \mathbb{R}^k$ specify the affine conic constraint.

Duality

The dual of problem (11.24) is

$$\begin{aligned}
& \text{maximize} && (l^c)^T s_l^c - (u^c)^T s_u^c + (l^x)^T s_l^x - (u^x)^T s_u^x - \dot{y}^T g + c^f \\
& \text{subject to} && A^T y + s_l^x - s_u^x + F^T \dot{y} = c, \\
& && -y + s_l^c - s_u^c = 0, \\
& && \bar{C}_j - \sum_{i=0}^{m-1} y_i \bar{A}_{ij} = \bar{S}_j, & j = 0, \dots, p-1 \\
& && s_l^c, s_u^c, s_l^x, s_u^x \geq 0, \\
& && \dot{y} \in \mathcal{K}^*, \\
& && \bar{S}_j \in \mathcal{S}_+^{r_j}, & j = 0, \dots, p-1.
\end{aligned} \tag{11.25}$$

Duality and infeasibility certificates behave analogously as in Sec. 11.2.

Remark

A problem of this form is internally converted into the problem:

$$\begin{aligned}
& \text{minimize} && \sum_{j=1}^n c_j x_j + \sum_{j=1}^p \langle \overline{C}_j, \overline{X}_j \rangle + c^f \\
& \text{subject to} && l_i^c \leq \sum_{j=1}^n a_{ij} x_j + \sum_{j=1}^p \langle \overline{A}_{ij}, \overline{X}_j \rangle \leq u_i^c, \quad i = 1, \dots, m, \\
& && g_i \leq z_i - \sum_{j=1}^n f_{ij} x_j \leq g_i, \quad i = 1, \dots, k, \\
& && l_j^x \leq x_j \leq u_j^x, \quad j = 1, \dots, n, \\
& && z \in \mathcal{K}, \\
& && \overline{X}_j \in \mathcal{S}_+^{r_j}, \quad j = 1, \dots, p
\end{aligned} \tag{11.26}$$

which conforms with the format in [Sec. 11.2](#) and [Sec. 11.3](#). The reformulated problem has $n+k$ variables, $m+k$ linear constraints and in total k variables in cones. The new columns and rows are appended at the end of the original ones. If a problem with affine conic constraints is saved to a file then this reformulation will be written.

Chapter 12

Optimizers

The most essential part of **MOSEK** are the optimizers:

- *primal simplex* (linear problems),
- *dual simplex* (linear problems),
- *interior-point* (linear, quadratic and conic problems),
- *mixed-integer* (problems with integer variables).

The structure of a successful optimization process is roughly:

- **Presolve**
 1. *Elimination*: Reduce the size of the problem.
 2. *Dualizer*: Choose whether to solve the primal or the dual form of the problem.
 3. *Scaling*: Scale the problem for better numerical stability.
- **Optimization**
 1. *Optimize*: Solve the problem using selected method.
 2. *Terminate*: Stop the optimization when specific termination criteria have been met.
 3. *Report*: Return the solution or an infeasibility certificate.

The preprocessing stage is transparent to the user, but useful to know about for tuning purposes. The purpose of the preprocessing steps is to make the actual optimization more efficient and robust. We discuss the details of the above steps in the following sections.

12.1 Presolve

Before an optimizer actually performs the optimization the problem is preprocessed using the so-called presolve. The purpose of the presolve is to

1. remove redundant constraints,
2. eliminate fixed variables,
3. remove linear dependencies,
4. substitute out (implied) free variables, and
5. reduce the size of the optimization problem in general.

After the presolved problem has been optimized the solution is automatically postsolved so that the returned solution is valid for the original problem. Hence, the presolve is completely transparent. For further details about the presolve phase, please see [\[AA95\]](#) and [\[AGMX96\]](#).

It is possible to fine-tune the behavior of the presolve or to turn it off entirely. If presolve consumes too much time or memory compared to the reduction in problem size gained it may be disabled. This is done by setting the parameter `MSK_IPAR_PRESOLVE_USE` to `"MSK_PRESOLVE_MODE_OFF"`. The two most time-consuming steps of the presolve are

- the eliminator, and
- the linear dependency check.

Therefore, in some cases it is worthwhile to disable one or both of these.

Numerical issues in the presolve

During the presolve the problem is reformulated so that it hopefully solves faster. However, in rare cases the presolved problem may be harder to solve than the original problem. The presolve may also be infeasible although the original problem is not. If it is suspected that presolved problem is much harder to solve than the original, we suggest to first turn the eliminator off by setting the parameter `MSK_IPAR_PRESOLVE_ELIMINATOR_MAX_NUM_TRIES` to 0. If that does not help, then trying to turn entire presolve off may help.

Since all computations are done in finite precision, the presolve employs some tolerances when concluding a variable is fixed or a constraint is redundant. If it happens that **MOSEK** incorrectly concludes a problem is primal or dual infeasible, then it is worthwhile to try to reduce the parameters `MSK_DPAR_PRESOLVE_TOL_X` and `MSK_DPAR_PRESOLVE_TOL_S`. However, if reducing the parameters actually helps then this should be taken as an indication that the problem is badly formulated.

Eliminator

The purpose of the eliminator is to eliminate free and implied free variables from the problem using substitution. For instance, given the constraints

$$\begin{aligned} y &= \sum_j x_j, \\ y, x &\geq 0, \end{aligned}$$

y is an implied free variable that can be substituted out of the problem, if deemed worthwhile. If the eliminator consumes too much time or memory compared to the reduction in problem size gained it may be disabled. This can be done by setting the parameter `MSK_IPAR_PRESOLVE_ELIMINATOR_MAX_NUM_TRIES` to 0. In rare cases the eliminator may cause that the problem becomes much hard to solve.

Linear dependency checker

The purpose of the linear dependency check is to remove linear dependencies among the linear equalities. For instance, the three linear equalities

$$\begin{aligned} x_1 + x_2 + x_3 &= 1, \\ x_1 + 0.5x_2 &= 0.5, \\ 0.5x_2 + x_3 &= 0.5. \end{aligned}$$

contain exactly one linear dependency. This implies that one of the constraints can be dropped without changing the set of feasible solutions. Removing linear dependencies is in general a good idea since it reduces the size of the problem. Moreover, the linear dependencies are likely to introduce numerical problems in the optimization phase. It is best practice to build models without linear dependencies, but that is not always easy for the user to control. If the linear dependencies are removed at the modeling stage, the linear dependency check can safely be disabled by setting the parameter `MSK_IPAR_PRESOLVE_LINDEP_USE` to `"MSK_OFF"`.

Dualizer

All linear, conic, and convex optimization problems have an equivalent dual problem associated with them. **MOSEK** has built-in heuristics to determine if it is more efficient to solve the primal or dual problem. The form (primal or dual) is displayed in the **MOSEK** log and available as an information item from the solver. Should the internal heuristics not choose the most efficient form of the problem it may be worthwhile to set the dualizer manually by setting the parameters:

- `MSK_IPAR_INTPTN_SOLVE_FORM`: In case of the interior-point optimizer.
- `MSK_IPAR_SIM_SOLVE_FORM`: In case of the simplex optimizer.

Note that currently only linear and conic (but not semidefinite) problems may be automatically dualized.

Scaling

Problems containing data with large and/or small coefficients, say $1.0e+9$ or $1.0e-7$, are often hard to solve. Significant digits may be truncated in calculations with finite precision, which can result in the optimizer relying on inaccurate data. Since computers work in finite precision, extreme coefficients should be avoided. In general, data around the same *order of magnitude* is preferred, and we will refer to a problem, satisfying this loose property, as being *well-scaled*. If the problem is not well scaled, **MOSEK** will try to scale (multiply) constraints and variables by suitable constants. **MOSEK** solves the scaled problem to improve the numerical properties.

The scaling process is transparent, i.e. the solution to the original problem is reported. It is important to be aware that the optimizer terminates when the termination criterion is met on the scaled problem, therefore significant primal or dual infeasibilities may occur after unscaling for badly scaled problems. The best solution of this issue is to reformulate the problem, making it better scaled.

By default **MOSEK** heuristically chooses a suitable scaling. The scaling for interior-point and simplex optimizers can be controlled with the parameters `MSK_IPAR_INTPNT_SCALING` and `MSK_IPAR_SIM_SCALING` respectively.

12.2 Linear Optimization

12.2.1 Optimizer Selection

Two different types of optimizers are available for linear problems: The default is an interior-point method, and the alternative is the simplex method (primal or dual). The optimizer can be selected using the parameter `MSK_IPAR_OPTIMIZER`.

The Interior-point or the Simplex Optimizer?

Given a linear optimization problem, which optimizer is the best: the simplex or the interior-point optimizer? It is impossible to provide a general answer to this question. However, the interior-point optimizer behaves more predictably: it tends to use between 20 and 100 iterations, almost independently of problem size, but cannot perform warm-start. On the other hand the simplex method can take advantage of an initial solution, but is less predictable from cold-start. The interior-point optimizer is used by default.

The Primal or the Dual Simplex Variant?

MOSEK provides both a primal and a dual simplex optimizer. Predicting which simplex optimizer is faster is impossible, however, in recent years the dual optimizer has seen several algorithmic and computational improvements, which, in our experience, make it faster on average than the primal version. Still, it depends much on the problem structure and size. Setting the `MSK_IPAR_OPTIMIZER` parameter to `"MSK_OPTIMIZER_FREE_SIMPLEX"` instructs **MOSEK** to choose one of the simplex variants automatically.

To summarize, if you want to know which optimizer is faster for a given problem type, it is best to try all the options.

12.2.2 The Interior-point Optimizer

The purpose of this section is to provide information about the algorithm employed in the **MOSEK** interior-point optimizer for linear problems and about its termination criteria.

The homogeneous primal-dual problem

In order to keep the discussion simple it is assumed that **MOSEK** solves linear optimization problems of standard form

$$\begin{aligned} & \text{minimize} && c^T x \\ & \text{subject to} && Ax = b, \\ & && x \geq 0. \end{aligned} \tag{12.1}$$

This is in fact what happens inside **MOSEK**; for efficiency reasons **MOSEK** converts the problem to standard form before solving, then converts it back to the input form when reporting the solution.

Since it is not known beforehand whether problem (12.1) has an optimal solution, is primal infeasible or is dual infeasible, the optimization algorithm must deal with all three situations. This is the reason why **MOSEK** solves the so-called homogeneous model

$$\begin{aligned} Ax - b\tau &= 0, \\ A^T y + s - c\tau &= 0, \\ -c^T x + b^T y - \kappa &= 0, \\ x, s, \tau, \kappa &\geq 0, \end{aligned} \tag{12.2}$$

where y and s correspond to the dual variables in (12.1), and τ and κ are two additional scalar variables. Note that the homogeneous model (12.2) always has solution since

$$(x, y, s, \tau, \kappa) = (0, 0, 0, 0, 0)$$

is a solution, although not a very interesting one. Any solution

$$(x^*, y^*, s^*, \tau^*, \kappa^*)$$

to the homogeneous model (12.2) satisfies

$$x_j^* s_j^* = 0 \text{ and } \tau^* \kappa^* = 0.$$

Moreover, there is always a solution that has the property $\tau^* + \kappa^* > 0$.

First, assume that $\tau^* > 0$. It follows that

$$\begin{aligned} A \frac{x^*}{\tau^*} &= b, \\ A^T \frac{y^*}{\tau^*} + \frac{s^*}{\tau^*} &= c, \\ -c^T \frac{x^*}{\tau^*} + b^T \frac{y^*}{\tau^*} &= 0, \\ x^*, s^*, \tau^*, \kappa^* &\geq 0. \end{aligned}$$

This shows that $\frac{x^*}{\tau^*}$ is a primal optimal solution and $(\frac{y^*}{\tau^*}, \frac{s^*}{\tau^*})$ is a dual optimal solution; this is reported as the optimal interior-point solution since

$$(x, y, s) = \left\{ \frac{x^*}{\tau^*}, \frac{y^*}{\tau^*}, \frac{s^*}{\tau^*} \right\}$$

is a primal-dual optimal solution (see Sec. 11.1 for the mathematical background on duality and optimality).

On other hand, if $\kappa^* > 0$ then

$$\begin{aligned} Ax^* &= 0, \\ A^T y^* + s^* &= 0, \\ -c^T x^* + b^T y^* &= \kappa^*, \\ x^*, s^*, \tau^*, \kappa^* &\geq 0. \end{aligned}$$

This implies that at least one of

$$c^T x^* < 0 \tag{12.3}$$

or

$$b^T y^* > 0 \tag{12.4}$$

is satisfied. If (12.3) is satisfied then x^* is a certificate of dual infeasibility, whereas if (12.4) is satisfied then y^* is a certificate of primal infeasibility.

In summary, by computing an appropriate solution to the homogeneous model, all information required for a solution to the original problem is obtained. A solution to the homogeneous model can be computed using a primal-dual interior-point algorithm [And09].

Interior-point Termination Criterion

For efficiency reasons it is not practical to solve the homogeneous model exactly. Hence, an exact optimal solution or an exact infeasibility certificate cannot be computed and a reasonable termination criterion has to be employed.

In the k -th iteration of the interior-point algorithm a trial solution

$$(x^k, y^k, s^k, \tau^k, \kappa^k)$$

to homogeneous model is generated, where

$$x^k, s^k, \tau^k, \kappa^k > 0.$$

Optimal case

Whenever the trial solution satisfies the criterion

$$\begin{aligned} \left\| A \frac{x^k}{\tau^k} - b \right\|_{\infty} &\leq \epsilon_p (1 + \|b\|_{\infty}), \\ \left\| A^T \frac{y^k}{\tau^k} + \frac{s^k}{\tau^k} - c \right\|_{\infty} &\leq \epsilon_d (1 + \|c\|_{\infty}), \text{ and} \\ \min \left(\frac{(x^k)^T s^k}{(\tau^k)^2}, \left| \frac{c^T x^k}{\tau^k} - \frac{b^T y^k}{\tau^k} \right| \right) &\leq \epsilon_g \max \left(1, \frac{\min(|c^T x^k|, |b^T y^k|)}{\tau^k} \right), \end{aligned} \quad (12.5)$$

the interior-point optimizer is terminated and

$$\frac{(x^k, y^k, s^k)}{\tau^k}$$

is reported as the primal-dual optimal solution. The interpretation of (12.5) is that the optimizer is terminated if

- $\frac{x^k}{\tau^k}$ is approximately primal feasible,
- $\left\{ \frac{y^k}{\tau^k}, \frac{s^k}{\tau^k} \right\}$ is approximately dual feasible, and
- the duality gap is almost zero.

Dual infeasibility certificate

On the other hand, if the trial solution satisfies

$$-\epsilon_i c^T x^k > \frac{\|c\|_{\infty}}{\max(1, \|b\|_{\infty})} \|Ax^k\|_{\infty}$$

then the problem is declared dual infeasible and x^k is reported as a certificate of dual infeasibility. The motivation for this stopping criterion is as follows: First assume that $\|Ax^k\|_{\infty} = 0$; then x^k is an exact certificate of dual infeasibility. Next assume that this is not the case, i.e.

$$\|Ax^k\|_{\infty} > 0,$$

and define

$$\bar{x} := \epsilon_i \frac{\max(1, \|b\|_{\infty})}{\|Ax^k\|_{\infty} \|c\|_{\infty}} x^k.$$

It is easy to verify that

$$\|A\bar{x}\|_{\infty} = \epsilon_i \frac{\max(1, \|b\|_{\infty})}{\|c\|_{\infty}} \text{ and } -c^T \bar{x} > 1,$$

which shows \bar{x} is an approximate certificate of dual infeasibility, where ϵ_i controls the quality of the approximation. A smaller value means a better approximation.

Primal infeasibility certificate

Finally, if

$$\epsilon_i b^T y^k > \frac{\|b\|_\infty}{\max(1, \|c\|_\infty)} \|A^T y^k + s^k\|_\infty$$

then y^k is reported as a certificate of primal infeasibility.

Adjusting optimality criteria

It is possible to adjust the tolerances ε_p , ε_d , ε_g and ε_i using parameters; see table for details.

Table 12.1: Parameters employed in termination criterion

ToleranceParameter	name
ε_p	<i>MSK_DPAR_INTPNT_TOL_PFEAS</i>
ε_d	<i>MSK_DPAR_INTPNT_TOL_DFEAS</i>
ε_g	<i>MSK_DPAR_INTPNT_TOL_REL_GAP</i>
ε_i	<i>MSK_DPAR_INTPNT_TOL_INFEAS</i>

The default values of the termination tolerances are chosen such that for a majority of problems appearing in practice it is not possible to achieve much better accuracy. Therefore, tightening the tolerances usually is not worthwhile. However, an inspection of (12.5) reveals that the quality of the solution depends on $\|b\|_\infty$ and $\|c\|_\infty$; the smaller the norms are, the better the solution accuracy.

The interior-point method as implemented by **MOSEK** will converge toward optimality and primal and dual feasibility at the same rate [And09]. This means that if the optimizer is stopped prematurely then it is very unlikely that either the primal or dual solution is feasible. Another consequence is that in most cases all the tolerances, ε_p , ε_d , ε_g and ε_i , have to be relaxed together to achieve an effect.

If the optimizer terminates without locating a solution that satisfies the termination criteria, for example because of a stall or other numerical issues, then it will check if the solution found up to that point satisfies the same criteria with all tolerances multiplied by the value of *MSK_DPAR_INTPNT_CO_TOL_NEAR_REL*. If this is the case, the solution is still declared as optimal.

The basis identification discussed in Sec. 12.2.2 requires an optimal solution to work well; hence basis identification should be turned off if the termination criterion is relaxed.

To conclude the discussion in this section, relaxing the termination criterion is usually not worthwhile.

Basis Identification

An interior-point optimizer does not return an optimal basic solution unless the problem has a unique primal and dual optimal solution. Therefore, the interior-point optimizer has an optional post-processing step that computes an optimal basic solution starting from the optimal interior-point solution. More information about the basis identification procedure may be found in [AY96]. In the following we provide an overall idea of the procedure.

There are some cases in which a basic solution could be more valuable:

- a basic solution is often more accurate than an interior-point solution,
- a basic solution can be used to warm-start the simplex algorithm in case of reoptimization,
- a basic solution is in general more sparse, i.e. more variables are fixed to zero. This is particularly appealing when solving continuous relaxations of mixed integer problems, as well as in all applications in which sparser solutions are preferred.

To illustrate how the basis identification routine works, we use the following trivial example:

$$\begin{aligned} &\text{minimize} && x + y \\ &\text{subject to} && x + y = 1, \\ &&& x, y \geq 0. \end{aligned}$$

It is easy to see that all feasible solutions are also optimal. In particular, there are two basic solutions, namely

$$\begin{aligned} (x_1^*, y_1^*) &= (1, 0), \\ (x_2^*, y_2^*) &= (0, 1). \end{aligned}$$

(continued from previous page)

Number of branches	: 4425
Number of relaxations solved	: 4410
Number of interior point iterations	: 25
Number of simplex iterations	: 221131

The first lines contain a summary of the problem as seen by the optimizer. This is followed by the iteration log. The columns have the following meaning:

- BRANCHES: Number of branches generated.
- RELAXS: Number of relaxations solved.
- ACT_NDS: Number of active branch bound nodes.
- DEPTH: Depth of the recently solved node.
- BEST_INT_OBJ: The best integer objective value, \bar{z} .
- BEST_RELAX_OBJ: The best objective bound, \underline{z} .
- REL_GAP(%): Relative optimality gap, $100\% \cdot \epsilon_{\text{rel}}$
- TIME: Time (in seconds) from the start of optimization.

Following that a summary of the optimization process is printed.

Chapter 13

Rmosek API Reference

- *Command reference:*
 - *Functions*
 - *Data structures*
- **Optimizer parameters:**
 - *Double, Integer, String*
 - *Full list*
 - *Browse by topic*
- **Optimizer information items:**
 - *Double, Integer, Long*
- *Optimizer response codes*
- *Constants*
- *Nonlinear API (scopt)*

13.1 Command Reference

The Rmosek interface is composed of a small set of functions and structures.

- *Function list*
- *Structures and data types list*

13.1.1 Functions

- *mosek*
- *mosek_version*
- *mosek_clean*
- *mosek_read*
- *mosek_write*

`r = mosek(prob, opts)`

Solve an optimization problem using the **MOSEK** optimization library. This is the main interface to **MOSEK**.

Parameters

- `prob` (*problem*) – A structure containing the optimization problem.
- `opts` (*options*) – The solver options.

Return `r` (*result*) – A structure containing solutions and other results. See [Sec. 7.1](#).

13.2 Parameters grouped by topic

Analysis

- *MSK_DPAR_ANA_SOL_INFEAS_TOL*
- *MSK_IPAR_ANA_SOL_BASIS*
- *MSK_IPAR_ANA_SOL_PRINT_VIOLATED*
- *MSK_IPAR_LOG_ANA_PRO*

Basis identification

- *MSK_DPAR_SIM_LU_TOL_REL_PIV*
- *MSK_IPAR_BI_CLEAN_OPTIMIZER*
- *MSK_IPAR_BI_IGNORE_MAX_ITER*
- *MSK_IPAR_BI_IGNORE_NUM_ERROR*
- *MSK_IPAR_BI_MAX_ITERATIONS*
- *MSK_IPAR_INTPNT_BASIS*
- *MSK_IPAR_LOG_BI*
- *MSK_IPAR_LOG_BI_FREQ*

Conic interior-point method

- *MSK_DPAR_INTPNT_CO_TOL_DFEAS*
- *MSK_DPAR_INTPNT_CO_TOL_INFEAS*
- *MSK_DPAR_INTPNT_CO_TOL_MU_RED*
- *MSK_DPAR_INTPNT_CO_TOL_NEAR_REL*
- *MSK_DPAR_INTPNT_CO_TOL_PFEAS*
- *MSK_DPAR_INTPNT_CO_TOL_REL_GAP*

Data check

- *MSK_DPAR_DATA_SYM_MAT_TOL*
- *MSK_DPAR_DATA_SYM_MAT_TOL_HUGE*
- *MSK_DPAR_DATA_SYM_MAT_TOL_LARGE*
- *MSK_DPAR_DATA_TOL_AIJ_HUGE*
- *MSK_DPAR_DATA_TOL_AIJ_LARGE*
- *MSK_DPAR_DATA_TOL_BOUND_INF*
- *MSK_DPAR_DATA_TOL_BOUND_WRN*
- *MSK_DPAR_DATA_TOL_C_HUGE*
- *MSK_DPAR_DATA_TOL_CJ_LARGE*

- *MSK_DPAR_DATA_TOL_QIJ*
- *MSK_DPAR_DATA_TOL_X*
- *MSK_DPAR_SEMIDEFINITE_TOL_APPROX*
- *MSK_IPAR_CHECK_CONVEXITY*
- *MSK_IPAR_LOG_CHECK_CONVEXITY*

Data input/output

- *MSK_IPAR_INFEAS_REPORT_AUTO*
- *MSK_IPAR_LOG_FILE*
- *MSK_IPAR_OPF_WRITE_HEADER*
- *MSK_IPAR_OPF_WRITE_HINTS*
- *MSK_IPAR_OPF_WRITE_LINE_LENGTH*
- *MSK_IPAR_OPF_WRITE_PARAMETERS*
- *MSK_IPAR_OPF_WRITE_PROBLEM*
- *MSK_IPAR_OPF_WRITE_SOL_BAS*
- *MSK_IPAR_OPF_WRITE_SOL_ITG*
- *MSK_IPAR_OPF_WRITE_SOL_ITR*
- *MSK_IPAR_OPF_WRITE_SOLUTIONS*
- *MSK_IPAR_PARAM_READ_CASE_NAME*
- *MSK_IPAR_PARAM_READ_IGN_ERROR*
- *MSK_IPAR_PTF_WRITE_TRANSFORM*
- *MSK_IPAR_READ_DEBUG*
- *MSK_IPAR_READ_KEEP_FREE_CON*
- *MSK_IPAR_READ_LP_DROP_NEW_VARS_IN_BOU*
- *MSK_IPAR_READ_LP_QUOTED_NAMES*
- *MSK_IPAR_READ_MPS_FORMAT*
- *MSK_IPAR_READ_MPS_WIDTH*
- *MSK_IPAR_READ_TASK_IGNORE_PARAM*
- *MSK_IPAR_SOL_READ_NAME_WIDTH*
- *MSK_IPAR_SOL_READ_WIDTH*
- *MSK_IPAR_WRITE_BAS_CONSTRAINTS*
- *MSK_IPAR_WRITE_BAS_HEAD*
- *MSK_IPAR_WRITE_BAS_VARIABLES*
- *MSK_IPAR_WRITE_COMPRESSION*
- *MSK_IPAR_WRITE_DATA_PARAM*
- *MSK_IPAR_WRITE_FREE_CON*

"MSK_BK_LO"
The constraint or variable has a finite lower bound and an infinite upper bound.

"MSK_BK_UP"
The constraint or variable has an infinite lower bound and an finite upper bound.

"MSK_BK_FX"
The constraint or variable is fixed.

"MSK_BK_FR"
The constraint or variable is free.

"MSK_BK_RA"
The constraint or variable is ranged.

mark
Mark

"MSK_MARK_LO"
The lower bound is selected for sensitivity analysis.

"MSK_MARK_UP"
The upper bound is selected for sensitivity analysis.

simdegen
Degeneracy strategies

"MSK_SIM_DEGEN_NONE"
The simplex optimizer should use no degeneration strategy.

"MSK_SIM_DEGEN_FREE"
The simplex optimizer chooses the degeneration strategy.

"MSK_SIM_DEGEN_AGGRESSIVE"
The simplex optimizer should use an aggressive degeneration strategy.

"MSK_SIM_DEGEN_MODERATE"
The simplex optimizer should use a moderate degeneration strategy.

"MSK_SIM_DEGEN_MINIMUM"
The simplex optimizer should use a minimum degeneration strategy.

transpose
Transposed matrix.

"MSK_TRANSPOSE_NO"
No transpose is applied.

"MSK_TRANSPOSE_YES"
A transpose is applied.

uplo
Triangular part of a symmetric matrix.

"MSK_UPLO_LO"
Lower part.

"MSK_UPLO_UP"
Upper part.

simreform
Problem reformulation.

"MSK_SIM_REFORMULATION_ON"
Allow the simplex optimizer to reformulate the problem.

"MSK_SIM_REFORMULATION_OFF"
Disallow the simplex optimizer to reformulate the problem.

"MSK_SIM_REFORMULATION_FREE"
The simplex optimizer can choose freely.

"MSK_SIM_REFORMULATION_AGGRESSIVE"
The simplex optimizer should use an aggressive reformulation strategy.

simdupvec
Exploit duplicate columns.

"MSK_SIM_EXPLOIT_DUPVEC_ON"
 Allow the simplex optimizer to exploit duplicated columns.

"MSK_SIM_EXPLOIT_DUPVEC_OFF"
 Disallow the simplex optimizer to exploit duplicated columns.

"MSK_SIM_EXPLOIT_DUPVEC_FREE"
 The simplex optimizer can choose freely.

simhotstart
 Hot-start type employed by the simplex optimizer

"MSK_SIM_HOTSTART_NONE"
 The simplex optimizer performs a coldstart.

"MSK_SIM_HOTSTART_FREE"
 The simplex optimizer chooses the hot-start type.

"MSK_SIM_HOTSTART_STATUS_KEYS"
 Only the status keys of the constraints and variables are used to choose the type of hot-start.

intpntthotstart
 Hot-start type employed by the interior-point optimizers.

"MSK_INTPNT_HOTSTART_NONE"
 The interior-point optimizer performs a coldstart.

"MSK_INTPNT_HOTSTART_PRIMAL"
 The interior-point optimizer exploits the primal solution only.

"MSK_INTPNT_HOTSTART_DUAL"
 The interior-point optimizer exploits the dual solution only.

"MSK_INTPNT_HOTSTART_PRIMAL_DUAL"
 The interior-point optimizer exploits both the primal and dual solution.

purify
 Solution purification employed optimizer.

"MSK_PURIFY_NONE"
 The optimizer performs no solution purification.

"MSK_PURIFY_PRIMAL"
 The optimizer purifies the primal solution.

"MSK_PURIFY_DUAL"
 The optimizer purifies the dual solution.

"MSK_PURIFY_PRIMAL_DUAL"
 The optimizer purifies both the primal and dual solution.

"MSK_PURIFY_AUTO"
 TBD

callbackcode
 Progress callback codes

"MSK_CALLBACK_BEGIN_BI"
 The basis identification procedure has been started.

"MSK_CALLBACK_BEGIN_CONIC"
 The callback function is called when the conic optimizer is started.

"MSK_CALLBACK_BEGIN_DUAL_BI"
 The callback function is called from within the basis identification procedure when the dual phase is started.

"MSK_CALLBACK_BEGIN_DUAL_SENSITIVITY"
 Dual sensitivity analysis is started.

"MSK_CALLBACK_BEGIN_DUAL_SETUP_BI"
 The callback function is called when the dual BI phase is started.

"MSK_CALLBACK_BEGIN_DUAL_SIMPLEX"
 The callback function is called when the dual simplex optimizer started.

"MSK_CALLBACK_BEGIN_DUAL_SIMPLEX_BI"
The callback function is called from within the basis identification procedure when the dual simplex clean-up phase is started.

"MSK_CALLBACK_BEGIN_FULL_CONVEXITY_CHECK"
Begin full convexity check.

"MSK_CALLBACK_BEGIN_INFEAS_ANA"
The callback function is called when the infeasibility analyzer is started.

"MSK_CALLBACK_BEGIN_INTPNT"
The callback function is called when the interior-point optimizer is started.

"MSK_CALLBACK_BEGIN_LICENSE_WAIT"
Begin waiting for license.

"MSK_CALLBACK_BEGIN_MIO"
The callback function is called when the mixed-integer optimizer is started.

"MSK_CALLBACK_BEGIN_OPTIMIZER"
The callback function is called when the optimizer is started.

"MSK_CALLBACK_BEGIN_PRESOLVE"
The callback function is called when the presolve is started.

"MSK_CALLBACK_BEGIN_PRIMAL_BI"
The callback function is called from within the basis identification procedure when the primal phase is started.

"MSK_CALLBACK_BEGIN_PRIMAL_REPAIR"
Begin primal feasibility repair.

"MSK_CALLBACK_BEGIN_PRIMAL_SENSITIVITY"
Primal sensitivity analysis is started.

"MSK_CALLBACK_BEGIN_PRIMAL_SETUP_BI"
The callback function is called when the primal BI setup is started.

"MSK_CALLBACK_BEGIN_PRIMAL_SIMPLEX"
The callback function is called when the primal simplex optimizer is started.

"MSK_CALLBACK_BEGIN_PRIMAL_SIMPLEX_BI"
The callback function is called from within the basis identification procedure when the primal simplex clean-up phase is started.

"MSK_CALLBACK_BEGIN_QCQO_REFORMULATE"
Begin QCQO reformulation.

"MSK_CALLBACK_BEGIN_READ"
MOSEK has started reading a problem file.

"MSK_CALLBACK_BEGIN_ROOT_CUTGEN"
The callback function is called when root cut generation is started.

"MSK_CALLBACK_BEGIN_SIMPLEX"
The callback function is called when the simplex optimizer is started.

"MSK_CALLBACK_BEGIN_SIMPLEX_BI"
The callback function is called from within the basis identification procedure when the simplex clean-up phase is started.

"MSK_CALLBACK_BEGIN_TO_CONIC"
Begin conic reformulation.

"MSK_CALLBACK_BEGIN_WRITE"
MOSEK has started writing a problem file.

"MSK_CALLBACK_CONIC"
The callback function is called from within the conic optimizer after the information database has been updated.

"MSK_CALLBACK_DUAL_SIMPLEX"
The callback function is called from within the dual simplex optimizer.

"MSK_CALLBACK_END_BI"
The callback function is called when the basis identification procedure is terminated.

"MSK_CALLBACK_END_CONIC"
The callback function is called when the conic optimizer is terminated.

"MSK_CALLBACK_END_DUAL_BI"
The callback function is called from within the basis identification procedure when the dual phase is terminated.

"MSK_CALLBACK_END_DUAL_SENSITIVITY"
Dual sensitivity analysis is terminated.

"MSK_CALLBACK_END_DUAL_SETUP_BI"
The callback function is called when the dual BI phase is terminated.

"MSK_CALLBACK_END_DUAL_SIMPLEX"
The callback function is called when the dual simplex optimizer is terminated.

"MSK_CALLBACK_END_DUAL_SIMPLEX_BI"
The callback function is called from within the basis identification procedure when the dual clean-up phase is terminated.

"MSK_CALLBACK_END_FULL_CONVEXITY_CHECK"
End full convexity check.

"MSK_CALLBACK_END_INFEAS_ANA"
The callback function is called when the infeasibility analyzer is terminated.

"MSK_CALLBACK_END_INTPNT"
The callback function is called when the interior-point optimizer is terminated.

"MSK_CALLBACK_END_LICENSE_WAIT"
End waiting for license.

"MSK_CALLBACK_END_MIO"
The callback function is called when the mixed-integer optimizer is terminated.

"MSK_CALLBACK_END_OPTIMIZER"
The callback function is called when the optimizer is terminated.

"MSK_CALLBACK_END_PRESOLVE"
The callback function is called when the presolve is completed.

"MSK_CALLBACK_END_PRIMAL_BI"
The callback function is called from within the basis identification procedure when the primal phase is terminated.

"MSK_CALLBACK_END_PRIMAL_REPAIR"
End primal feasibility repair.

"MSK_CALLBACK_END_PRIMAL_SENSITIVITY"
Primal sensitivity analysis is terminated.

"MSK_CALLBACK_END_PRIMAL_SETUP_BI"
The callback function is called when the primal BI setup is terminated.

"MSK_CALLBACK_END_PRIMAL_SIMPLEX"
The callback function is called when the primal simplex optimizer is terminated.

"MSK_CALLBACK_END_PRIMAL_SIMPLEX_BI"
The callback function is called from within the basis identification procedure when the primal clean-up phase is terminated.

"MSK_CALLBACK_END_QCQO_REFORMULATE"
End QCQO reformulation.

"MSK_CALLBACK_END_READ"
MOSEK has finished reading a problem file.

"MSK_CALLBACK_END_ROOT_CUTGEN"
The callback function is called when root cut generation is terminated.

"MSK_CALLBACK_END_SIMPLEX"
The callback function is called when the simplex optimizer is terminated.

"MSK_CALLBACK_END_SIMPLEX_BI"
The callback function is called from within the basis identification procedure when the simplex clean-up phase is terminated.

"MSK_CALLBACK_END_TO_CONIC"
End conic reformulation.

"MSK_CALLBACK_END_WRITE"
MOSEK has finished writing a problem file.

"MSK_CALLBACK_IM_BI"
The callback function is called from within the basis identification procedure at an intermediate point.

"MSK_CALLBACK_IM_CONIC"
The callback function is called at an intermediate stage within the conic optimizer where the information database has not been updated.

"MSK_CALLBACK_IM_DUAL_BI"
The callback function is called from within the basis identification procedure at an intermediate point in the dual phase.

"MSK_CALLBACK_IM_DUAL_SENSIVITY"
The callback function is called at an intermediate stage of the dual sensitivity analysis.

"MSK_CALLBACK_IM_DUAL_SIMPLEX"
The callback function is called at an intermediate point in the dual simplex optimizer.

"MSK_CALLBACK_IM_FULL_CONVEXITY_CHECK"
The callback function is called at an intermediate stage of the full convexity check.

"MSK_CALLBACK_IM_INTPNT"
The callback function is called at an intermediate stage within the interior-point optimizer where the information database has not been updated.

"MSK_CALLBACK_IM_LICENSE_WAIT"
MOSEK is waiting for a license.

"MSK_CALLBACK_IM_LU"
The callback function is called from within the LU factorization procedure at an intermediate point.

"MSK_CALLBACK_IM_MIO"
The callback function is called at an intermediate point in the mixed-integer optimizer.

"MSK_CALLBACK_IM_MIO_DUAL_SIMPLEX"
The callback function is called at an intermediate point in the mixed-integer optimizer while running the dual simplex optimizer.

"MSK_CALLBACK_IM_MIO_INTPNT"
The callback function is called at an intermediate point in the mixed-integer optimizer while running the interior-point optimizer.

"MSK_CALLBACK_IM_MIO_PRIMAL_SIMPLEX"
The callback function is called at an intermediate point in the mixed-integer optimizer while running the primal simplex optimizer.

"MSK_CALLBACK_IM_ORDER"
The callback function is called from within the matrix ordering procedure at an intermediate point.

"MSK_CALLBACK_IM_PRESOLVE"
The callback function is called from within the presolve procedure at an intermediate stage.

"MSK_CALLBACK_IM_PRIMAL_BI"
The callback function is called from within the basis identification procedure at an intermediate point in the primal phase.

checkconvexitytype
Types of convexity checks.

"MSK_CHECK_CONVEXITY_NONE"
No convexity check.

"MSK_CHECK_CONVEXITY_SIMPLE"
Perform simple and fast convexity check.

"MSK_CHECK_CONVEXITY_FULL"
Perform a full convexity check.

compresstype
Compression types

"MSK_COMPRESS_NONE"
No compression is used.

"MSK_COMPRESS_FREE"
The type of compression used is chosen automatically.

"MSK_COMPRESS_GZIP"
The type of compression used is gzip compatible.

"MSK_COMPRESS_ZSTD"
The type of compression used is zstd compatible.

conetype
Cone types

"MSK_CT_QUAD"
The cone is a quadratic cone.

"MSK_CT_RQUAD"
The cone is a rotated quadratic cone.

"MSK_CT_PEXP"
A primal exponential cone.

"MSK_CT_DEXP"
A dual exponential cone.

"MSK_CT_PPOW"
A primal power cone.

"MSK_CT_DPOW"
A dual power cone.

"MSK_CT_ZERO"
The zero cone.

nametype
Name types

"MSK_NAME_TYPE_GEN"
General names. However, no duplicate and blank names are allowed.

"MSK_NAME_TYPE_MPS"
MPS type names.

"MSK_NAME_TYPE_LP"
LP type names.

scopr
SCopt operator types

"MSK_OPR_ENT"
Entropy

"MSK_OPR_EXP"
Exponential

"MSK_OPR_LOG"
Logarithm

"MSK_OPR_POW"
Power

"MSK_OPR_SQRT"
Square root

symmattype
Cone types

"MSK_SYMMAT_TYPE_SPARSE"
Sparse symmetric matrix.

dataformat
Data format types

"MSK_DATA_FORMAT_EXTENSION"
The file extension is used to determine the data file format.

"MSK_DATA_FORMAT_MPS"
The data file is MPS formatted.

"MSK_DATA_FORMAT_LP"
The data file is LP formatted.

"MSK_DATA_FORMAT_OP"
The data file is an optimization problem formatted file.

"MSK_DATA_FORMAT_FREE_MPS"
The data a free MPS formatted file.

"MSK_DATA_FORMAT_TASK"
Generic task dump file.

"MSK_DATA_FORMAT_PTF"
(P)retty (T)ext (F)format.

"MSK_DATA_FORMAT_CB"
Conic benchmark format,

"MSK_DATA_FORMAT_JSON_TASK"
JSON based task format.

dinfitem
Double information items

"MSK_DINF_BI_CLEAN_DUAL_TIME"
Time spent within the dual clean-up optimizer of the basis identification procedure since its invocation.

"MSK_DINF_BI_CLEAN_PRIMAL_TIME"
Time spent within the primal clean-up optimizer of the basis identification procedure since its invocation.

"MSK_DINF_BI_CLEAN_TIME"
Time spent within the clean-up phase of the basis identification procedure since its invocation.

"MSK_DINF_BI_DUAL_TIME"
Time spent within the dual phase basis identification procedure since its invocation.

"MSK_DINF_BI_PRIMAL_TIME"
Time spent within the primal phase of the basis identification procedure since its invocation.

"MSK_DINF_BI_TIME"
Time spent within the basis identification procedure since its invocation.

"MSK_DINF_INTPNT_DUAL_FEAS"
Dual feasibility measure reported by the interior-point optimizer. (For the interior-point optimizer this measure is not directly related to the original problem because a homogeneous model is employed.)

"MSK_DINF_INTPNT_DUAL_OBJ"
Dual objective value reported by the interior-point optimizer.

"MSK_DINF_INTPNT_FACTOR_NUM_FLOPS"
An estimate of the number of flops used in the factorization.

"MSK_FEATURE_PTS"
Base system.

"MSK_FEATURE_PTON"
Conic extension.

liinfitem
Long integer information items.

"MSK_LIINF_BI_CLEAN_DUAL_DEG_ITER"
Number of dual degenerate clean iterations performed in the basis identification.

"MSK_LIINF_BI_CLEAN_DUAL_ITER"
Number of dual clean iterations performed in the basis identification.

"MSK_LIINF_BI_CLEAN_PRIMAL_DEG_ITER"
Number of primal degenerate clean iterations performed in the basis identification.

"MSK_LIINF_BI_CLEAN_PRIMAL_ITER"
Number of primal clean iterations performed in the basis identification.

"MSK_LIINF_BI_DUAL_ITER"
Number of dual pivots performed in the basis identification.

"MSK_LIINF_BI_PRIMAL_ITER"
Number of primal pivots performed in the basis identification.

"MSK_LIINF_INTPNT_FACTOR_NUM_NZ"
Number of non-zeros in factorization.

"MSK_LIINF_MIO_ANZ"
Number of non-zero entries in the constraint matrix of the problem to be solved by the mixed-integer optimizer.

"MSK_LIINF_MIO_INTPNT_ITER"
Number of interior-point iterations performed by the mixed-integer optimizer.

"MSK_LIINF_MIO_PRE SOLVED_ANZ"
Number of non-zero entries in the constraint matrix of the problem after the mixed-integer optimizer's presolve.

"MSK_LIINF_MIO_SIMPLEX_ITER"
Number of simplex iterations performed by the mixed-integer optimizer.

"MSK_LIINF_RD_NUMANZ"
Number of non-zeros in A that is read.

"MSK_LIINF_RD_NUMQNZ"
Number of Q non-zeros.

iinfitem
Integer information items.

"MSK_IINF_ANA_PRO_NUM_CON"
Number of constraints in the problem.

"MSK_IINF_ANA_PRO_NUM_CON_EQ"
Number of equality constraints.

"MSK_IINF_ANA_PRO_NUM_CON_FR"
Number of unbounded constraints.

"MSK_IINF_ANA_PRO_NUM_CON_LO"
Number of constraints with a lower bound and an infinite upper bound.

"MSK_IINF_ANA_PRO_NUM_CON_RA"
Number of constraints with finite lower and upper bounds.

"MSK_IINF_ANA_PRO_NUM_CON_UP"
Number of constraints with an upper bound and an infinite lower bound.

"MSK_IINF_ANA_PRO_NUM_VAR"
Number of variables in the problem.

"MSK_IINF_ANA_PRO_NUM_VAR_BIN"
Number of binary (0-1) variables.

"MSK_IINF_ANA_PRO_NUM_VAR_CONT"
Number of continuous variables.

"MSK_IINF_ANA_PRO_NUM_VAR_EQ"
Number of fixed variables.

"MSK_IINF_ANA_PRO_NUM_VAR_FR"
Number of free variables.

"MSK_IINF_ANA_PRO_NUM_VAR_INT"
Number of general integer variables.

"MSK_IINF_ANA_PRO_NUM_VAR_LO"
Number of variables with a lower bound and an infinite upper bound.

"MSK_IINF_ANA_PRO_NUM_VAR_RA"
Number of variables with finite lower and upper bounds.

"MSK_IINF_ANA_PRO_NUM_VAR_UP"
Number of variables with an upper bound and an infinite lower bound.

"MSK_IINF_INTPNT_FACTOR_DIM_DENSE"
Dimension of the dense sub system in factorization.

"MSK_IINF_INTPNT_ITER"
Number of interior-point iterations since invoking the interior-point optimizer.

"MSK_IINF_INTPNT_NUM_THREADS"
Number of threads that the interior-point optimizer is using.

"MSK_IINF_INTPNT_SOLVE_DUAL"
Non-zero if the interior-point optimizer is solving the dual problem.

"MSK_IINF_MIO_ABSGAP_SATISFIED"
Non-zero if absolute gap is within tolerances.

"MSK_IINF_MIO_CLIQUETABLE_SIZE"
Size of the clique table.

"MSK_IINF_MIO_CONSTRUCT_SOLUTION"
This item informs if **MOSEK** constructed an initial integer feasible solution.

- -1: tried, but failed,
- 0: no partial solution supplied by the user,
- 1: constructed feasible solution.

"MSK_IINF_MIO_NODE_DEPTH"
Depth of the last node solved.

"MSK_IINF_MIO_NUM_ACTIVE_NODES"
Number of active branch and bound nodes.

"MSK_IINF_MIO_NUM_BRANCH"
Number of branches performed during the optimization.

"MSK_IINF_MIO_NUM_CLIQUETCUTS"
Number of clique cuts.

"MSK_IINF_MIO_NUM_CMIR_CUTS"
Number of Complemented Mixed Integer Rounding (CMIR) cuts.

"MSK_IINF_MIO_NUM_GOMORY_CUTS"
Number of Gomory cuts.

"MSK_IINF_MIO_NUM_IMPLIED_BOUND_CUTS"
Number of implied bound cuts.

"MSK_IINF_MIO_NUM_INT_SOLUTIONS"
Number of integer feasible solutions that have been found.

"MSK_IINF_MIO_NUM_KNAPSACK_COVER_CUTS"
Number of clique cuts.

"MSK_IINF_RD_NUMVAR"
Number of variables read.

"MSK_IINF_RD_PROTOTYPE"
Problem type.

"MSK_IINF_SIM_DUAL_DEG_ITER"
The number of dual degenerate iterations.

"MSK_IINF_SIM_DUAL_HOTSTART"
If 1 then the dual simplex algorithm is solving from an advanced basis.

"MSK_IINF_SIM_DUAL_HOTSTART_LU"
If 1 then a valid basis factorization of full rank was located and used by the dual simplex algorithm.

"MSK_IINF_SIM_DUAL_INF_ITER"
The number of iterations taken with dual infeasibility.

"MSK_IINF_SIM_DUAL_ITER"
Number of dual simplex iterations during the last optimization.

"MSK_IINF_SIM_NUMCON"
Number of constraints in the problem solved by the simplex optimizer.

"MSK_IINF_SIM_NUMVAR"
Number of variables in the problem solved by the simplex optimizer.

"MSK_IINF_SIM_PRIMAL_DEG_ITER"
The number of primal degenerate iterations.

"MSK_IINF_SIM_PRIMAL_HOTSTART"
If 1 then the primal simplex algorithm is solving from an advanced basis.

"MSK_IINF_SIM_PRIMAL_HOTSTART_LU"
If 1 then a valid basis factorization of full rank was located and used by the primal simplex algorithm.

"MSK_IINF_SIM_PRIMAL_INF_ITER"
The number of iterations taken with primal infeasibility.

"MSK_IINF_SIM_PRIMAL_ITER"
Number of primal simplex iterations during the last optimization.

"MSK_IINF_SIM_SOLVE_DUAL"
Is non-zero if dual problem is solved.

"MSK_IINF_SOL_BAS_PROSTA"
Problem status of the basic solution. Updated after each optimization.

"MSK_IINF_SOL_BAS_SOLSTA"
Solution status of the basic solution. Updated after each optimization.

"MSK_IINF_SOL_ITG_PROSTA"
Problem status of the integer solution. Updated after each optimization.

"MSK_IINF_SOL_ITG_SOLSTA"
Solution status of the integer solution. Updated after each optimization.

"MSK_IINF_SOL_ITR_PROSTA"
Problem status of the interior-point solution. Updated after each optimization.

"MSK_IINF_SOL_ITR_SOLSTA"
Solution status of the interior-point solution. Updated after each optimization.

"MSK_IINF_STO_NUM_A_REALLOC"
Number of times the storage for storing A has been changed. A large value may indicate that memory fragmentation may occur.

infotype
Information item types

"MSK_INF_DOU_TYPE"
Is a double information type.

"MSK_INF_INT_TYPE"
Is an integer.

"MSK_INF_LINT_TYPE"
Is a long integer.

iomode
Input/output modes

"MSK_IOMODE_READ"
The file is read-only.

"MSK_IOMODE_WRITE"
The file is write-only. If the file exists then it is truncated when it is opened. Otherwise it is created when it is opened.

"MSK_IOMODE_READWRITE"
The file is to read and write.

branchdir
Specifies the branching direction.

"MSK_BRANCH_DIR_FREE"
The mixed-integer optimizer decides which branch to choose.

"MSK_BRANCH_DIR_UP"
The mixed-integer optimizer always chooses the up branch first.

"MSK_BRANCH_DIR_DOWN"
The mixed-integer optimizer always chooses the down branch first.

"MSK_BRANCH_DIR_NEAR"
Branch in direction nearest to selected fractional variable.

"MSK_BRANCH_DIR_FAR"
Branch in direction farthest from selected fractional variable.

"MSK_BRANCH_DIR_ROOT_LP"
Chose direction based on root lp value of selected variable.

"MSK_BRANCH_DIR_GUIDED"
Branch in direction of current incumbent.

"MSK_BRANCH_DIR_PSEUDOCOST"
Branch based on the pseudocost of the variable.

miocontsoltype
Continuous mixed-integer solution type

"MSK_MIO_CONT_SOL_NONE"
No interior-point or basic solution are reported when the mixed-integer optimizer is used.

"MSK_MIO_CONT_SOL_ROOT"
The reported interior-point and basic solutions are a solution to the root node problem when mixed-integer optimizer is used.

"MSK_MIO_CONT_SOL_ITG"
The reported interior-point and basic solutions are a solution to the problem with all integer variables fixed at the value they have in the integer solution. A solution is only reported in case the problem has a primal feasible solution.

"MSK_MIO_CONT_SOL_ITG_REL"
In case the problem is primal feasible then the reported interior-point and basic solutions are a solution to the problem with all integer variables fixed at the value they have in the integer solution. If the problem is primal infeasible, then the solution to the root node problem is reported.

miomode
Integer restrictions

"MSK_MIO_MODE_IGNORED"
The integer constraints are ignored and the problem is solved as a continuous problem.

"MSK_MIO_MODE_SATISFIED"
Integer restrictions should be satisfied.

mionodeseltype
Mixed-integer node selection types

"MSK_MIO_NODE_SELECTION_FREE"
The optimizer decides the node selection strategy.

"MSK_MIO_NODE_SELECTION_FIRST"
The optimizer employs a depth first node selection strategy.

"MSK_MIO_NODE_SELECTION_BEST"
The optimizer employs a best bound node selection strategy.

"MSK_MIO_NODE_SELECTION_PSEUDO"
The optimizer employs selects the node based on a pseudo cost estimate.

mpsformat
MPS file format type

"MSK_MPS_FORMAT_STRICT"
It is assumed that the input file satisfies the MPS format strictly.

"MSK_MPS_FORMAT_RELAXED"
It is assumed that the input file satisfies a slightly relaxed version of the MPS format.

"MSK_MPS_FORMAT_FREE"
It is assumed that the input file satisfies the free MPS format. This implies that spaces are not allowed in names. Otherwise the format is free.

"MSK_MPS_FORMAT_CPLEX"
The CPLEX compatible version of the MPS format is employed.

objsense
Objective sense types

"MSK_OBJECTIVE_SENSE_MINIMIZE"
The problem should be minimized.

"MSK_OBJECTIVE_SENSE_MAXIMIZE"
The problem should be maximized.

onoffkey
On/off

"MSK_ON"
Switch the option on.

"MSK_OFF"
Switch the option off.

optimizertype
Optimizer types

"MSK_OPTIMIZER_CONIC"
The optimizer for problems having conic constraints.

"MSK_OPTIMIZER_DUAL_SIMPLEX"
The dual simplex optimizer is used.

"MSK_OPTIMIZER_FREE"
The optimizer is chosen automatically.

"MSK_OPTIMIZER_FREE_SIMPLEX"
One of the simplex optimizers is used.

"MSK_OPTIMIZER_INTPNT"
The interior-point optimizer is used.

"MSK_OPTIMIZER_MIXED_INT"
The mixed-integer optimizer.

"MSK_OPTIMIZER_PRIMAL_SIMPLEX"
The primal simplex optimizer is used.

orderingtype
Ordering strategies

"MSK_ORDER_METHOD_FREE"
The ordering method is chosen automatically.

"MSK_ORDER_METHOD_APPMINLOC"
Approximate minimum local fill-in ordering is employed.

"MSK_ORDER_METHOD_EXPERIMENTAL"
This option should not be used.

"MSK_ORDER_METHOD_TRY_GRAPHPAR"
Always try the graph partitioning based ordering.

"MSK_ORDER_METHOD_FORCE_GRAPHPAR"
Always use the graph partitioning based ordering even if it is worse than the approximate minimum local fill ordering.

"MSK_ORDER_METHOD_NONE"
No ordering is used.

presolvemode
Presolve method.

"MSK_PRESOLVE_MODE_OFF"
The problem is not presolved before it is optimized.

"MSK_PRESOLVE_MODE_ON"
The problem is presolved before it is optimized.

"MSK_PRESOLVE_MODE_FREE"
It is decided automatically whether to presolve before the problem is optimized.

parametertype
Parameter type

"MSK_PAR_INVALID_TYPE"
Not a valid parameter.

"MSK_PAR_DOU_TYPE"
Is a double parameter.

"MSK_PAR_INT_TYPE"
Is an integer parameter.

"MSK_PAR_STR_TYPE"
Is a string parameter.

problemitem
Problem data items

"MSK_PI_VAR"
Item is a variable.

"MSK_PI_CON"
Item is a constraint.

"MSK_PI_CONE"
Item is a cone.

problemtypes
Problem types

"MSK_PROBTYPE_LO"
The problem is a linear optimization problem.

"MSK_PROBTYPE_QO"
The problem is a quadratic optimization problem.

"MSK_PROBTYPE_QCQO"
The problem is a quadratically constrained optimization problem.

"MSK_PROBTYPE_CONIC"
A conic optimization.

"MSK_PROBTYPE_MIXED"
General nonlinear constraints and conic constraints. This combination can not be solved by **MOSEK**.

prosta
 Problem status keys

"MSK_PRO_STA_UNKNOWN"
 Unknown problem status.

"MSK_PRO_STA_PRIM_AND_DUAL_FEAS"
 The problem is primal and dual feasible.

"MSK_PRO_STA_PRIM_FEAS"
 The problem is primal feasible.

"MSK_PRO_STA_DUAL_FEAS"
 The problem is dual feasible.

"MSK_PRO_STA_PRIM_INFEAS"
 The problem is primal infeasible.

"MSK_PRO_STA_DUAL_INFEAS"
 The problem is dual infeasible.

"MSK_PRO_STA_PRIM_AND_DUAL_INFEAS"
 The problem is primal and dual infeasible.

"MSK_PRO_STA_ILL_POSED"
 The problem is ill-posed. For example, it may be primal and dual feasible but have a positive duality gap.

"MSK_PRO_STA_PRIM_INFEAS_OR_UNBOUNDED"
 The problem is either primal infeasible or unbounded. This may occur for mixed-integer problems.

xmlwriteroutputtype
 XML writer output mode

"MSK_WRITE_XML_MODE_ROW"
 Write in row order.

"MSK_WRITE_XML_MODE_COL"
 Write in column order.

rescodetype
 Response code type

"MSK_RESPONSE_OK"
 The response code is OK.

"MSK_RESPONSE_WRN"
 The response code is a warning.

"MSK_RESPONSE_TRM"
 The response code is an optimizer termination status.

"MSK_RESPONSE_ERR"
 The response code is an error.

"MSK_RESPONSE_UNK"
 The response code does not belong to any class.

scalingtype
 Scaling type

"MSK_SCALING_FREE"
 The optimizer chooses the scaling heuristic.

"MSK_SCALING_NONE"
 No scaling is performed.

"MSK_SCALING_MODERATE"
 A conservative scaling is performed.

"MSK_SCALING_AGGRESSIVE"
 A very aggressive scaling is performed.

scalingmethod
 Scaling method

"MSK_SCALING_METHOD_POW2"
Scales only with power of 2 leaving the mantissa untouched.

"MSK_SCALING_METHOD_FREE"
The optimizer chooses the scaling heuristic.

sensitivitytype
Sensitivity types

"MSK_SENSITIVITY_TYPE_BASIS"
Basis sensitivity analysis is performed.

simseltype
Simplex selection strategy

"MSK_SIM_SELECTION_FREE"
The optimizer chooses the pricing strategy.

"MSK_SIM_SELECTION_FULL"
The optimizer uses full pricing.

"MSK_SIM_SELECTION_ASE"
The optimizer uses approximate steepest-edge pricing.

"MSK_SIM_SELECTION_DEVEX"
The optimizer uses devex steepest-edge pricing (or if it is not available an approximate steep-edge selection).

"MSK_SIM_SELECTION_SE"
The optimizer uses steepest-edge selection (or if it is not available an approximate steep-edge selection).

"MSK_SIM_SELECTION_PARTIAL"
The optimizer uses a partial selection approach. The approach is usually beneficial if the number of variables is much larger than the number of constraints.

solitem
Solution items

"MSK_SOL_ITEM_XC"
Solution for the constraints.

"MSK_SOL_ITEM_XX"
Variable solution.

"MSK_SOL_ITEM_Y"
Lagrange multipliers for equations.

"MSK_SOL_ITEM_SLC"
Lagrange multipliers for lower bounds on the constraints.

"MSK_SOL_ITEM_SUC"
Lagrange multipliers for upper bounds on the constraints.

"MSK_SOL_ITEM_SLX"
Lagrange multipliers for lower bounds on the variables.

"MSK_SOL_ITEM_SUX"
Lagrange multipliers for upper bounds on the variables.

"MSK_SOL_ITEM_SNX"
Lagrange multipliers corresponding to the conic constraints on the variables.

solsta
Solution status keys

"MSK_SOL_STA_UNKNOWN"
Status of the solution is unknown.

"MSK_SOL_STA_OPTIMAL"
The solution is optimal.

"MSK_SOL_STA_PRIM_FEAS"
The solution is primal feasible.

"MSK_SOL_STA_DUAL_FEAS"
The solution is dual feasible.

"MSK_SOL_STA_PRIM_AND_DUAL_FEAS"
The solution is both primal and dual feasible.

"MSK_SOL_STA_PRIM_INFEAS_CER"
The solution is a certificate of primal infeasibility.

"MSK_SOL_STA_DUAL_INFEAS_CER"
The solution is a certificate of dual infeasibility.

"MSK_SOL_STA_PRIM_ILLPOSED_CER"
The solution is a certificate that the primal problem is illposed.

"MSK_SOL_STA_DUAL_ILLPOSED_CER"
The solution is a certificate that the dual problem is illposed.

"MSK_SOL_STA_INTEGER_OPTIMAL"
The primal solution is integer optimal.

solttype
Solution types

"MSK_SOL_BAS"
The basic solution.

"MSK_SOL_ITR"
The interior solution.

"MSK_SOL_ITG"
The integer solution.

solveform
Solve primal or dual form

"MSK_SOLVE_FREE"
The optimizer is free to solve either the primal or the dual problem.

"MSK_SOLVE_PRIMAL"
The optimizer should solve the primal problem.

"MSK_SOLVE_DUAL"
The optimizer should solve the dual problem.

stakey
Status keys

"MSK_SK_UNK"
The status for the constraint or variable is unknown.

"MSK_SK_BAS"
The constraint or variable is in the basis.

"MSK_SK_SUPBAS"
The constraint or variable is super basic.

"MSK_SK_LOW"
The constraint or variable is at its lower bound.

"MSK_SK_UPR"
The constraint or variable is at its upper bound.

"MSK_SK_FIX"
The constraint or variable is fixed.

"MSK_SK_INF"
The constraint or variable is infeasible in the bounds.

startpointtype
Starting point types

"MSK_STARTING_POINT_FREE"
The starting point is chosen automatically.

"MSK_STARTING_POINT_GUESS"
The optimizer guesses a starting point.

By default **MOSEK** writes ranged constraints this way.

If the files must adhere to the LP standard, ranged constraints must either be split into upper bounded and lower bounded constraints or be written as an equality with a slack variable. For example the expression (14.1) may be written as

$$x_1 + x_2 - sl_1 = 0, \quad -5 \leq sl_1 \leq 5.$$

Bounds

Bounds on the variables can be specified in the bound section beginning with one of the keywords

```
bound
bounds
```

The bounds section is optional but should, if present, follow the **subject to** section. All variables listed in the bounds section must occur in either the objective or a constraint.

The default lower and upper bounds are 0 and $+\infty$. A variable may be declared free with the keyword **free**, which means that the lower bound is $-\infty$ and the upper bound is $+\infty$. Furthermore it may be assigned a finite lower and upper bound. The bound definitions for a given variable may be written in one or two lines, and bounds can be any number or $\pm\infty$ (written as **+inf/-inf/+infinity/-infinity**) as in the example

```
bounds
x1 free
x2 <= 5
0.1 <= x2
x3 = 42
2 <= x4 < +inf
```

Variable Types

The final two sections are optional and must begin with one of the keywords

```
bin
binaries
binary
```

and

```
gen
general
```

Under **general** all integer variables are listed, and under **binary** all binary (integer variables with bounds 0 and 1) are listed:

```
general
x1 x2
binary
x3 x4
```

Again, all variables listed in the binary or general sections must occur in either the objective or a constraint.

Terminating Section

Finally, an LP formatted file must be terminated with the keyword

```
end
```

14.1.2 LP File Examples

Linear example `lo1.lp`

```
\ File: lo1.lp
maximize
obj: 3 x1 + x2 + 5 x3 + x4
subject to
c1: 3 x1 + x2 + 2 x3 = 30
c2: 2 x1 + x2 + 3 x3 + x4 >= 15
c3: 2 x2 + 3 x4 <= 25
bounds
0 <= x1 <= +infinity
0 <= x2 <= 10
0 <= x3 <= +infinity
0 <= x4 <= +infinity
end
```

Mixed integer example `mil01.lp`

```
maximize
obj: x1 + 6.4e-01 x2
subject to
c1: 5e+01 x1 + 3.1e+01 x2 <= 2.5e+02
c2: 3e+00 x1 - 2e+00 x2 >= -4e+00
bounds
0 <= x1 <= +infinity
0 <= x2 <= +infinity
general
x1 x2
end
```

14.1.3 LP Format peculiarities

Comments

Anything on a line after a `\` is ignored and is treated as a comment.

Names

A name for an objective, a constraint or a variable may contain the letters `a-z`, `A-Z`, the digits `0-9` and the characters

```
!"#$%&()/,.;?@_`'|~
```

The first character in a name must not be a number, a period or the letter `e` or `E`. Keywords must not be used as names.

MOSEK accepts any character as valid for names, except `\0`. A name that is not allowed in LP file will be changed and a warning will be issued.

The algorithm for making names LP valid works as follows: The name is interpreted as an `utf-8` string. For a Unicode character `c`:

- If `c==_` (underscore), the output is `__` (two underscores).
- If `c` is a valid LP name character, the output is just `c`.
- If `c` is another character in the ASCII range, the output is `_XX`, where `XX` is the hexadecimal code for the character.
- If `c` is a character in the range `127-65535`, the output is `_uXXXX`, where `XXXX` is the hexadecimal code for the character.

(continued from previous page)

```
QUADOBJ
  [vname1] [vname2] [value1]
QCMATRIX  [cname1]
  [vname1] [vname2] [value1]
BOUNDS
  ?? [name] [vname1] [value1]
CSECTION  [kname1] [value1] [ktype]
  [vname1]
ENDATA
```

Here the names in capitals are keywords of the MPS format and names in brackets are custom defined names or values. A couple of notes on the structure:

- Fields: All items surrounded by brackets appear in *fields*. The fields named “valueN” are numerical values. Hence, they must have the format

```
[+|-]XXXXXXX.XXXXXX[e|E][+|-]XXX]
```

where

```
X = [0|1|2|3|4|5|6|7|8|9].
```

- Sections: The MPS file consists of several sections where the names in capitals indicate the beginning of a new section. For example, COLUMNS denotes the beginning of the columns section.
- Comments: Lines starting with an * are comment lines and are ignored by **MOSEK**.
- Keys: The question marks represent keys to be specified later.
- Extensions: The sections QSECTION and CSECTION are specific **MOSEK** extensions of the MPS format. The sections QMATRIX, QUADOBJ and QCMATRIX are included for sake of compatibility with other vendors extensions to the MPS format.
- The standard MPS format is a fixed format, i.e. everything in the MPS file must be within certain fixed positions. **MOSEK** also supports a *free format*. See [Sec. 14.2.5](#) for details.

Linear example lo1.mps

A concrete example of a MPS file is presented below:

```
* File: lo1.mps
NAME          lo1
OBJSENSE
  MAX
ROWS
  N  obj
  E  c1
  G  c2
  L  c3
COLUMNS
  x1      obj      3
  x1      c1       3
  x1      c2       2
  x2      obj      1
  x2      c1       1
  x2      c2       1
  x2      c3       2
  x3      obj      5
  x3      c1       2
  x3      c2       3
  x4      obj      1
  x4      c2       1
```

(continues on next page)

bound, 10, 76, 79
dual, 77, 80
integer, 27
linear optimization, 10
semidefinite, 20