



Tour de MOSEK 7: The short version

Erling D. Andersen
MOSEK ApS,
Fruebjergvej 3, Box 16,
2100 Copenhagen,
Denmark.

WWW: <http://www.mosek.com>

October 1, 2013

<http://www.mosek.com>

Introduction

Introduction

MOSEK

The tour

Stage 1: The
interior-point
optimizer

Stage 2:
Computational
results

Stage 3: Semidefinite
optimization

Stage 4: The new
mixed integer conic
optimizer

- MOSEK is software package for large scale optimization.
- Version 7 released May 2013.
- Linear and conic quadratic (+ mixed-integer).
- Conic quadratic optimization (+ mixed-integer).
- Convex(functional) optimization.
- C, JAVA, .NET and Python APIs.
- AMPL, AIMMS, GAMS, MATLAB, and R interfaces.
- Free for academic use. See <http://www.mosek.com>.

The tour

Introduction

The tour

New in MOSEK
version 7

Stage 1: The
interior-point
optimizer

Stage 2:
Computational
results

Stage 3: Semidefinite
optimization

Stage 4: The new
mixed integer conic
optimizer

- Interior-point optimizer
 - ◆ Improved linear algebra.
 - ◆ Handling of semi-definite optimization problems.
- Simplex optimizer.
 - ◆ Improved LU.
- New mixed integer optimizer for conic problems.
- Fusion, a modeling tool for conic problems.

Stage 1: The interior-point optimizer

Introduction

The tour

Stage 1: The interior-point optimizer

Linear algebra improvements

Improvements

Stage 2: Computational results

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

MOSEK solves:

$$(P) \quad \min \quad c^T x \\ \text{st} \quad Ax = b, \\ x \geq 0.$$

where A is **large** and **sparse** using a homogenous interior-point algorithm.

Each iterations requires solution of multiple:

$$\begin{bmatrix} -H^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} f^1 \\ f^2 \end{bmatrix} = \begin{bmatrix} r^1 \\ r^2 \end{bmatrix} \quad (1)$$

where H is a diagonal matrix.

- Known as the **augmented system**.

Introduction

The tour

Stage 1: The interior-point optimizer

Linear algebra improvements

Improvements

Stage 2: Computational results

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

Is reduced to:

$$A^T A f_2 = \quad (2)$$

- **The normal equation system!**
- Identical to the augmented system approach using a particular pivot order.
- System is reordered to preserve sparsity using approximative min degree (AMD) or graph partitioning (GP).
- Fixed pivot order.
- Well-known: Just one dense columns in A cause a lot of fill-in.

Introduction

The tour

Stage 1: The interior-point optimizer

Linear algebra improvements

Improvements

Stage 2: Computational results

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

- New graph partitioning based ordering inspired by METIS.
- Rewritten factorization
 - ◆ Employ vendor BLAS e.g. Intel MKL to exploit AVX.
 - ◆ Does not employ OpenMP.
- New dense column detection method based on graph partitioning.

Introduction

The tour

Stage 1: The interior-point optimizer

Linear algebra improvements

Improvements

Stage 2: Computational results

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

- Let \bar{S} be the index set of the sparse columns in A . And \bar{N} the index set of the dense ones.
- Solve a system with the matrix

$$\begin{bmatrix} A_{\bar{S}} H_{\bar{S}} A_{\bar{S}}^T & A_{\bar{N}} \\ A_{\bar{N}}^T & -H_{\bar{N}}^{-1} \end{bmatrix}$$

- Pivot order is fixed.
- Requires $A_{\bar{S}} H_{\bar{S}} A_{\bar{S}}^T$ to be of full rank.
 - ◆ May lead to numerical instability.
- Consequence: Minimize number of dense columns.

Introduction

The tour

Stage 1: The interior-point optimizer

Linear algebra improvements

Improvements

Stage 2: Computational results

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

Dens .	Num .
8	8
9	81
10	544
11	3782
12	17227
13	48321
14	62561
15	1
16	3
17	15
18	27
19	127
20	224
21	193

- Has dense columns!
- Which cutoff to use?

Introduction

The tour

Stage 1: The interior-point optimizer

Linear algebra improvements

Improvements

Stage 2: Computational results

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

- Idea.
 - ◆ Try to emulate the optimal ordering for the augmented system.
 - ◆ Fixed pivot order.
 - ◆ Keep detection cost down.
- Solve a linear system of the form:

$$\begin{bmatrix} A_{\bar{S}} H_{\bar{S}} A_{\bar{S}}^T & A_{\bar{N}} \\ A_{\bar{N}}^T & -H_{\bar{N}}^{-1} \end{bmatrix}$$

Let (\bar{S}, \bar{N}) be an initial guess for the partition. And choose a reordering P so

$$P \begin{bmatrix} A_{\bar{S}} H_{\bar{S}} A_{\bar{S}}^T & A_{\bar{N}} \\ A_{\bar{N}}^T & -H_{\bar{N}}^{-1} \end{bmatrix} P^T = \begin{bmatrix} M_{11} & 0 & M_{13} \\ 0 & M_{22} & M_{23} \\ M_{31} & M_{32} & M_{33} \end{bmatrix}$$

- M_{11} and M_{22} should be of about identical size.
- M_{33} should be as small a possible.
- Ordering can be locate using graph partitioning i.e. use MeTIS or the like.
- Nodes that appears in both \bar{N} and M_{33} are the dense columns.
- A refined (\bar{S}, \bar{N}) is obtained.
- Many refinements possible!

Stage 2: Computational results

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Dense column handling

Factor speed improvements

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

- Comparison of MOSEK 6.0.0.155 and v7.0.0.87.
- Linux OS:

```
model name      : Intel(R) Xeon(R) CPU E3-1270 V2
cache size     : 8192 KB
cpu cores      : 4
```

- Using 2 thread unless otherwise stated.
- Problems
 - ◆ Private and public test problems

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Dense column handling

Factor speed improvements

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

Name	Time (s)	R. time	Iter.	R. iter.	Num. dense	
	v6	v7/v6	v6	v7/v6	v6	v7
difns8t4	5.11	0.91	27	1.04	74	89
net12	5.20	0.39	42	0.53	544	545
bienstock-310809-1	6.12	0.76	19	2.50	400	625
bas1	6.73	0.39	14	0.53	5	40
gonnew16	8.16	3.52	39	0.82	246	329
GON8IO	8.64	0.69	29	0.90	73	278
ind3	10.43	1.10	12	1.00	3	185
15dec2008	10.64	0.40	21	0.91	175	287
L1_nine12	11.14	0.20	15	1.31	29	0
pointlogic-210911-1	12.00	9.59	45	0.74	451	175
lt	13.17	0.56	47	0.48	292	506
_time_horizon_minimiser	15.66	0.31	14	1.07	23	0
dray17	18.07	0.38	90	1.24	55	448
ind2	29.48	0.76	12	1.00	318	1018
zhao4	29.79	0.29	31	1.06	680	0
neos3	35.44	1.61	9	2.90	1	2
friedlander-6	59.57	0.22	20	1.14	0	721
c3	69.65	1.75	9	1.60	57	0
avq1	69.84	3.95	12	1.15	541	1
ml2010-rmine14	78.02	0.92	25	1.35	28	28
TestA5	82.47	0.62	14	1.07	373	1487
dray5	92.63	0.39	52	0.68	0	1203
rusltplan	101.53	0.60	42	0.98	718	2094
stormG2_1000	107.62	0.48	108	0.50	119	119
tp-6	128.14	1.47	49	0.88	776	742
karted	138.35	4.11	20	0.95	193	590
scipmsk1	147.49	1.08	16	1.35	749	1
ts-palko	164.95	0.20	213	0.13	210	210
degme	169.48	2.02	62	0.56	883	890
160910-2	173.58	0.87	71	5.57	291	1893
gamshuge	729.91	0.78	98	1.08	270	44
G. avg		0.80		0.98		

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Dense column handling

Factor speed improvements

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

- Other changes contributes to difference.
 - ◆ New dualization heuristic.
 - ◆ Better programming, new compiler etc.
- Many dense columns in v7.
 - ◆ Does not affect stability much.
- New method seems to work well.
 - ◆ Can be relatively expensive for smallish problems.

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Dense column handling

Factor speed improvements

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

Name	Time (s)		Iter. v6	R. iter. v7/v6	Num. dense	
	v6	R. time v7/v6			v6	v7
model_2223	0.21	0.75	29	0.77	25	26
20130123_600	0.26	0.94	49	0.46	29	7
20130123_300	0.31	0.70	52	0.57	31	9
msprob3	0.32	1.48	32	1.15	73	31
ramsey3	0.45	1.90	12	1.31	199	212
20130123_900	0.46	0.66	59	0.50	39	9
050508-1	0.68	1.48	25	1.12	72	199
041208-1	1.08	0.50	25	1.00	5	93
230608-1	1.81	1.93	62	0.83	7	6
20130123_1000	2.79	0.55	82	0.43	26	12
280108-1	3.33	0.98	50	0.82	67	15
pcqo-250112-1	5.48	0.47	17	1.00	1176	1260
211107-1	6.34	0.86	50	0.86	0	1725
oxam5-230412	7.42	1.74	30	1.06	20	11
msci-p1to	8.49	0.08	24	0.76	42	331
autooc	8.65	1.80	28	0.90	41	230
bleyer-200312-1	16.88	0.24	11	1.50	0	1828
201107-3	17.12	0.39	50	1.02	0	0
oxam3-230412	65.72	0.63	94	0.49	20	11
G. avg		0.75		0.82		

Comment:

- Results is similar to the linear case.

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Dense column handling

Factor speed improvements

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

- Very limited and selective test.
- Only Mittlemans conic quadratic benchmark.

Name	Time (s) v6	R. time v7/v6	Iter. v6	R. iter. v7/v6
2013_firL2L1alph	38.57	0.67	13	1.00
2013_firL2Linfeps	143.07	0.65	21	1.23
2013_firL1Linfeps	180.41	0.52	91	1.03
2013_wbNRL	222.43	0.17	22	1.00
2013_firL1	287.86	0.36	21	1.00
2013_firL2L1eps	291.89	0.30	23	0.92
2013_firL1Linfalph	697.25	0.33	26	1.15
2013_firLinf	1295.83	0.42	25	1.23
2013_dsNRL	1334.43	0.28	42	0.91
2013_firL2Linfalph	1882.27	0.38	26	1.33
2013_firL2a	2398.54	0.29	9	0.70
G. avg		0.37		1.03

- Large speed-up.
- Dense problems!

Comparison to other optimizers

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Dense column handling

Factor speed improvements

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

```
27 Sep 2013  =====
              MISOCP and large SOCP Benchmark
              =====
              Hans D. Mittelmann (mittelmann@asu.edu)
```

Logfiles for these runs are at: plato.la.asu.edu/ftp/socp_logs/

```
MOSEK-7.0.0.85 www.mosek.com/
CPLEX-12.5.1 CPLEX
GUROBI-5.5.0 GUROBI
XPRESS-7.6 XPRESS
```

...

```
=====
```

problem	CPLEX	GUROBI	MOSEK	XPRESS
dsNRL	707	296	316	209
firL1	597	379	100	179
firL1Linfalph	1087	526	235	525
firL1Linfeps	331	48	92	105
firL2L1alph	85	99	28	17
firL2L1eps	573	371	77	581
firL2Linfalph	500	529	509	926
firL2Linfeps	290	152	94	134
firL2a	388	800	579	85
firLinf	672	420	502	1135
wbNRL	324	131	37	196

```
=====
```

Stage 3: Semidefinite optimization

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Stage 3: Semidefinite optimization

The primal problem

Computational results

Optimized problems

Result

Stage 4: The new mixed integer conic optimizer

$$\begin{aligned}
 \min \quad & \sum_{j=1}^n c_j x_j + \sum_{j=1}^{\bar{n}} \langle \bar{C}_j, \bar{X}_j \rangle \\
 \text{st} \quad & l_i^c \leq \sum_{j=1}^n a_{ij} x_j + \sum_{j=1}^{\bar{n}} \langle \bar{A}_{ij}, \bar{X}_j \rangle \leq u_i^c, \quad i = 1, \dots, m \\
 & l_j^x \leq x_j \leq u_j^x, \quad j = 1, \dots, n, \\
 & x \in \mathcal{K}, \\
 & \bar{X}_j \succeq 0, \quad j = 1, \dots, \bar{n}.
 \end{aligned}$$

Explanation:

- $\langle A, B \rangle := \text{tr}(A^T B)$.
- x_j is a scalar variable.
- \bar{X}_j is a square matrix variable.
- \mathcal{K} represents conic quadratic constraints.
- $\bar{X}_j \succeq 0$ represents $\bar{X}_j = \bar{X}_j^T$ and \bar{X}_j is PSD.
- \bar{C}_j and \bar{A}_j are assumed to be symmetric.

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Stage 3: Semidefinite optimization

The primal problem

Computational results

Optimized problems

Result

Stage 4: The new mixed integer conic optimizer

- Windows server 2008.
- CPU: Intel XEON 3-1270. 4 Cores. 3.5 GHz.
- MOSEK:
 - ◆ Only 2 threads allowed.
- Test problems:
 - ◆ Taken from: <http://plato.asu.edu/ftp/sdp/>
 - ◆ Exclude the smallest and largest with respect to running time.

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Stage 3: Semidefinite optimization

The primal problem Computational results

Optimized problems

Result

Stage 4: The new mixed integer conic optimizer

Name	# con.	# cone	# var.	# mat. var.
G40_mb	2001	0	0	1
G40mc	2000	0	0	1
G48mc	3000	0	0	1
buck4	1200	0	1200	2
butcher	6434	0	3697	1
cancer_100	10469	0	0	1
cphil12	12376	0	0	1
foot	2209	0	0	1
hand	1297	0	0	1
inc_1200	5175	0	5174	1
inc_600	2515	0	2514	1
mater-6	20463	0	2	4966
neu3g	8007	0	0	1
nonc_500	4990	0	0	501
rabmo	5004	0	1654	1
reimer5	6187	0	6502	1
rendl1_600_0	601	0	0	1
ros_500	4988	0	0	499
rose15	3860	0	2	1
sensor_1000	5549	0	5546	1
sensor_500	3540	0	3537	1
shmup3	420	0	840	2
shmup4	800	0	1600	2
swissroll	3380	0	0	1
taha1a	3002	0	0	14
taha1b	8007	0	3	21
trto5	3280	0	3280	1
vibra5	3280	0	3280	2
yalsdp	5051	0	0	3

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Stage 3: Semidefinite optimization

The primal problem

Computational results

Optimized problems

Result

Stage 4: The new mixed integer conic optimizer

Name	P. obj.	# sig. fig.	# iter	time(s)
G40_mb	2.8643213490e+03	9	28	696.7
G40mc	5.7295764053e+03	9	22	403.7
G48mc	1.1999999909e+04	9	8	390.2
buck4	4.8624158709e+02	7	43	66.7
butcher	1.3999995866e+01	8	24	111.1
cancer_100	2.7623368607e+04	9	27	309.5
cphil12	0.0000000000e+00	10	6	134.3
foot	5.8530052646e+05	9	52	3444.3
hand	2.4747760883e+04	8	29	376.3
inc_1200	1.1527230666e+00	4	49	325.5
inc_600	6.6542132801e-01	4	52	46.3
mater-6	-1.3353870357e+02	11	43	50.6
neu3g	-2.0432660497e-08	8	18	185.1
nonc_500	-6.2571885516e-02	6	32	1.8
rabmo	3.7272482995e+00	9	20	88.8
reimer5	1.5183500067e+01	8	17	229.9
rend11_600_0	5.5796870423e+04	8	29	26.1
ros_500	-2.4949949663e+00	7	29	1.0
rose15	-2.9095377285e-05	7	21	26.2
sensor_1000	-2.3359891483e+00	4	39	180.9
sensor_500	-1.7389724937e+01	5	29	28.8
shmup3	2.0962607260e+03	5	53	175.2
shmup4	7.9882423009e+03	5	85	1718.8
swissroll	5.6238493722e+05	5	31	69.5
taha1a	1.0000000647e+00	8	16	32.1
taha1b	7.7331369251e-01	8	26	184.8
trto5	1.2791923555e+04	7	59	695.4
vibra5	1.6566666010e+02	7	64	1460.0
yalsdp	1.7921272862e+00	7	22	112.8

Stage 4: The new mixed integer conic optimizer

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

The optimizer

Summary and conclusions

- Build a new mixed integer conic optimizer from scratch.
- Branch and bound plus usual cuts and tricks.
- Is parallelized.
- Run-to-run deterministic.
- Free-of-charge.
- Tuned for conic quadratic problems.
- Not so fast for pure linear problems.
- Applicable to convex QO and QCQO.

Computational results

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

The optimizer

Summary and conclusions

```
27 Sep 2013  =====
              MISOCP and large SOCP Benchmark
              =====
              Hans D. Mittelmann (mittelmann@asu.edu)
```

Logfiles for these runs are at: plato.la.asu.edu/ftp/socp_logs/

```
MOSEK-7.0.0.85 www.mosek.com/
CPLEX-12.5.1 CPLEX
GUROBI-5.5.0 GUROBI
XPRESS-7.6 XPRESS
```

...

problem	CPLEX	GUROBI	MOSEK	XPRESS
pp-n100-d10	6228	4562	t	903
uflquad-nopsc-20-100	1135	667	261	117
uflquad-nopsc-20-150	6453	980	1031	590
uflquad-nopsc-30-100	t	902	1622	276
uflquad-nopsc-30-150	t	2959	3628	772
uflquad-nopsc-30-200	t	5211	t	f
uflquad-psc-30-150	25	1011	4	1
uflquad-psc-30-200	89	5572	5	3
uflquad-psc-30-300	292	t	18	19

Introduction

The tour

Stage 1: The interior-point optimizer

Stage 2: Computational results

Stage 3: Semidefinite optimization

Stage 4: The new mixed integer conic optimizer

The optimizer

Summary and conclusions

- Improved the linear algebra. Particularly for dense problems.
- Can handle semidefinite optimization problems.
- New mixed integer conic optimizer.
- Version 7 is a significant improvement over version 6.
- Slide are at <http://mosek.com/resources/presentations/>.