MOSEK version 9 (work in progress)

July 4, 2018

Erling D. Andersen

www.mosek.com

Mosek

- A software package.
- Solves large-scale sparse linear, quadratic and conic optimization problems.
- Stand-alone as well as embedded.
- Version 1 release in 1999.
- Version 8 is released Fall 2016.
- Version 9 planned for release Fall 2018.

For details about interfaces, trials, academic license etc. see

`https://mosek.com`.

Improvements in version 9

An overview

- Conic optimizer
  - Added support for exponential and power cones.
  - Added functions for doing elementary projections on cones.
- Mixed-integer optimizer.
  - Added support for exponential and power cones.
  - Improved performance.
- Fusion(a modelling API for conic optimization problems).
  - Improved performance.
  - A lot of polishing and bug fixing.
- MATLAB and R toolboxes.
  - Added interface affine conic constraints.
- .NET core package. (Is just released for v8).

Deprecated features

- Advanced sensitivity analysis.
  - The traditional analysis based on the basis is still available.
- General convex optimization.

$$\begin{array}{lll} \min & f(x) & \\ \text{subject to} & g(x) & \leq & 0, \end{array}$$

  except for explicit quadratic functions.
- Dropped the Fusion interface for MATLAB.
  - Use affine conic constraints instead (Fusion light).

Generic conic optimization problem

Primal form

$$
\begin{array}{lll}
\text{minimize} & \sum_k (c^k)^T x^k & \\
\text{subject to} & \sum_k A^k x^k & = \quad b, \\
 & x^k \in \mathcal{K}^k, & \quad \forall k,
\end{array}
$$

where

- $c^k \in \mathbb{R}^{n^k}$,
- $A^k \in \mathbb{R}^{m \times n^k}$,
- $b \in \mathbb{R}^m$,
- $\mathcal{K}^k$ are convex cones.

The 5 cones

MOSEK v9 will support the 5 cone types:

- Linear.
- Quadratic.
- Semidefinite.
- **Exponential**.
- **Power**.

- Almost all convex problems appearing in practice can be formulated using those 5 cones.
- See my blog post from 2010 about a lunch with Stephen Boyd at Stanford:
  - `http://erlingdandersen.blogspot.com/2010/11/which-cones-are-needed-to-represent.html`
- Until now we simply did not have a satisfactory algorithm handling the nonsymmetric cones.

The power cone

The power cone:

$$\mathcal{K}_{pow}(\alpha) := \left\{ (x, z) : \prod_{j=1}^{n} x_j^{|\alpha_j|} \geq \|z\|^{\sum_{j=1}^{n} |\alpha_j|},\ x \geq 0 \right\}.$$

Examples ($\alpha \in (0,1)$):

$$\begin{aligned}
(t, 1, x) \in \mathcal{K}_{pow}(\alpha, 1-\alpha) \quad &\Leftrightarrow \quad t \geq |x|^{1/\alpha},\ t \geq 0, \\
(x, 1, t) \in \mathcal{K}_{pow}(\alpha, 1-\alpha) \quad &\Leftrightarrow \quad x^{\alpha} \geq |t|, x \geq 0, \\
(x, t) \in \mathcal{K}_{pow}(e) \quad &\Leftrightarrow \quad \left(\prod_{j=1}^{n} x_j \right)^{1/n} \geq |t|,\ x \geq 0.
\end{aligned}$$

More examples that can modelled using the power cone from Chares [1]:

- p-norm:

$$t \geq \|x\|_p \,.$$

- l_p cone:

$$\left\{ (x, t, s) : \; \sum_{j=1}^{n} \left(\frac{1}{p_i} \left(\frac{|x_j|}{t} \right)^{p_j} \right) \leq \frac{s}{t}, \, t \geq 0 \right\}$$

where $p > 0$.

Dual power cone

- Is self-dual using a redefined inner-product.
- But is not homogeneous.
- And there not symmetric.

The exponential cone

The exponential cone

$$\begin{aligned}\mathcal{K}_{\exp} \quad := \quad & \{(x_1, x_2, x_3) : \; x_1 \geq x_2 e^{\frac{x_3}{x_2}}, \; x_2 \geq 0\} \\ & \cup \{(x_1, x_2, x_3) : \; x_1 \geq 0, x_2 = 0, \; x_3 \leq 0\}\end{aligned}$$

Applications:

$$\begin{aligned}(t, 1, x) \in \mathcal{K}_{\exp} \quad &\Leftrightarrow \quad t \geq e^x, \\ (t, 1, \ln(a)x) \in \mathcal{K}_{\exp} \quad &\Leftrightarrow \quad t \geq a^x, \\ (x, 1, t) \in \mathcal{K}_{\exp} \quad &\Leftrightarrow \quad t \leq \ln(x), \\ (1, x, t) \in \mathcal{K}_{\exp} \quad &\Leftrightarrow \quad t \leq -x\ln(x), \\ (y, x, -t) \in \mathcal{K}_{\exp} \quad &\Leftrightarrow \quad t \geq x\ln(x/y), \text{(relative entropy)}.\end{aligned}$$

Challenge

- Do you have a convex set that can be modelled in say AMPL or GAMS that is not representable using the 5 cones?
- Then please contact MOSEK support.
  - The MOSEK modelling wizards Michael or Henrik is likely to prove you wrong.

- Has been extended to handle 3 dimensional power cones and exponential cones.
- Reuse the presolve, the efficient linear algebra from the existing conic optimizer. One code path!
- Algorithm based on work of: Tuncel [5], Myklebust and T. [2].
- Related work: Skajaa and Ye [4], Serrano [3].
- Future: Will add the n dimensional power cone and p norm cones.

Mixed integer conic optimizer

- Has been extended to handle to the nonsymmetric cones.
- Work-in-progress: Outer approximation algorithm for solution of the relaxations.

Fusion

- A modelling orientated API for build affine conic constraints.
- Alternative to CVX, CVXPY etc.
- Performance improvements.
  - Handling of big linear expressions are streamlined.
  - Adding many semi-definite variables has been tuned. Think about a bunch of 3 dimensional variables.
- Reduced memory fragmentation.
- Polishing.

MATLAB and R toolboxes

- Homogenized the functionality of the 2 toolboxes.
- Added support for the nonsymmetric cone types.
- Added support for affine conic constraints:

$$Fx + g \in \mathcal{K}.$$

Previously:

$$\begin{array}{rcl} Fx + g - s & = & 0 \\ s & \in & K. \end{array}$$

  - Simplifies model building.
  - Particularly for overlapping of variables.

Exponential/power cone optimization

- Hardware: Intel based server. (Xeon Gold 6126 2.6 GHz, 12 core)
- MOSEK: Version 9 (alpha). Highly experimental!
- Threads: 8 threads is used in test to simulate a typical user environment.
- All timing results t are in wall clock seconds.
- Test problems: Public (e.g `cblib.zib.de`) and customer supplied.

Optimized problems

Name	# con.	# cone	# var.	# mat. var.
task_dopt3	1600	26	376	2
task_dopt16	1600	26	376	2
entolib_a_bd	26	4695	14085	0
entolib_ento2	26	4695	14085	0
task_dopt10	1600	26	376	2
task_dopt17	1600	26	376	2
entolib_a_36	37	7497	22491	0
entolib_ento3	28	5172	15516	0
task_dopt12	1600	26	376	2
task_dopt21	1600	26	376	2
entolib_a_25	37	6196	18588	0
entolib_ento26	28	7915	23745	0
entolib_ento45	37	9108	27324	0
entolib_a_26	37	9035	27105	0
entolib_ento25	28	10142	30426	0
entolib_a_16	37	8528	25584	0
entolib_a_56	37	9702	29106	0
exp-ml-scaled-20000	19999	20000	79998	0
entolib_entodif	40	12691	38073	0
exp-ml-20000	19999	20000	79998	0
patil3_conv	418681	413547	1264340	0
c-diaz_test_c47	164404	160000	519810	0

Exponential/power cone optimization

Result

Name	P. obj.	# sig. fig.	# iter	time(s)
task_dopt3	1.5283637983e+01	10	15	0.6
task_dopt16	1.3214502950e+01	8	13	0.5
entolib_a_bd	-1.1355545199e+01	6	32	0.4
entolib_ento2	-1.1355545199e+01	6	32	0.4
task_dopt10	1.4373686806e+01	9	15	0.6
task_dopt17	1.6884197007e+01	8	26	1.0
entolib_a_36	2.1889809271e+00	5	39	1.1
entolib_ento3	-6.5140869257e+00	4	40	0.6
task_dopt12	2.3128665910e+01	8	18	0.7
task_dopt21	2.5769935873e+01	9	21	0.8
entolib_a_25	-7.9665952286e+00	5	42	0.7
entolib_ento26	-1.1578799246e+01	5	36	0.7
entolib_ento45	-8.7979548118e+00	4	45	1.1
entolib_a_26	-7.6599669507e+00	5	36	0.9
entolib_ento25	-7.2816620638e+00	5	47	1.1
entolib_a_16	-4.7675480178e+00	5	53	1.1
entolib_a_56	-8.2843376213e+00	5	39	1.0
exp-ml-scaled-20000	-3.3117384598e+00	6	57	6.8
entolib_entodif	-6.3553400228e+00	5	44	1.4
exp-ml-20000	-1.9741149790e+04	7	91	8.7
patil3_conv	-1.0538487417e+00	5	84	106.7
c-diaz_test_c47	1.7662819461e-02	5	55	46.0

- We use a subset of Mittelmann's benchmark-set + customer provided problems.
- MOSEK on 4 threads, MATLAB up to 20 threads.
- Problems are categorized as **failed** if
  - MOSEK returns unknown solution status.
  - SeDuMi returns `info.numerr==2`.
  - SDPT3 returns `info.termcode` $\notin [0, 1, 2]$ and `norm(info.dimacs,Inf)>1e-5`.

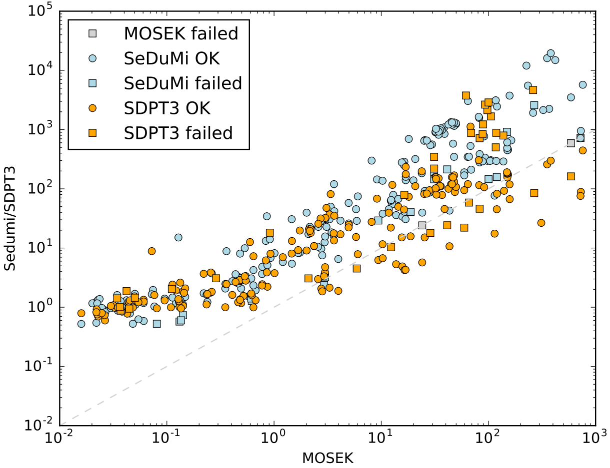

Solution time for MOSEK v8.0.0.42 vs SeDuMi/SDPT3 on 234 problems.

	small			medium		
	MOSEK	SeDuMi	SDPT3	MOSEK	SeDuMi	SDPT3
Num.	127	127	127	63	63	63
Firsts	116	1	10	47	0	16
Total time	218.2	1299.5	843.6	2396.0	32709.8	5679.5

	large			fails		
	MOSEK	SeDuMi	SDPT3	MOSEK	SeDuMi	SDPT3
Num.	44	44	44	6	27	47
Firsts	31	0	13			
Total time	9083.8	119818.1	35268.2			

- MOSEK is fastest on average with fewest failures.
- SeDuMi is almost always slower than MOSEK.

The take home message

- Version 9 supports the power and exponential cones.
  - A breaktrough!
- V9 has received a lot polishing.
- Affine conic constraints is a nice (syntactic) addition to the Matlab and R toolbox.

[1] Peter Robert Chares.
Cones and interior-point algorithms for structed convex optimization involving powers and exponentials.
PhD thesis, Ecole polytechnique de Louvain, Universitet catholique de Louvain, 2009.

[2] T. Myklebust and L. Tunçel.
Interior-point algorithms for convex optimization based on primal-dual metrics.
Technical report, 2014.

[3] Santiago Akle Serrano.
Algorithms for unsymmetric cone optimization and an implementation for problems with the exponential cone.
PhD thesis, Stanford University, 2015.

[4] Anders Skajaa and Yinye Ye.
A homogeneous interior-point algorithm for nonsymmetric convex conic optimization.
Math. Programming, 150:391–422, May 2015.

[5] L. Tunçel.
Generalization of primal-dual interior-point methods to convex optimization problems in conic form.
Foundations of Computational Mathematics, 1:229–254, 2001.