Performance of Optimization Software - an Update

CREST Workshop

Chuo University, Tokyo

23 March 2012

H. D. Mittelmann

School of Math and Stat Sciences

Arizona State University

Services we provide

- Guide to Software: **"Decision Tree"**
- http://plato.asu.edu/guide.html
- Software Archive
- Software Evaluation: **"Benchmarks"**
- Archive of Testproblems
- Web-based Solvers **(1/3 of NEOS)**

We maintain the following NEOS solvers (8 categories)

```
Combinatorial Optimization  * CONCORDE [TSP Input]
Global Optimization         * ICOS [AMPL Input]
Linear Programming
  * bpmpd [AMPL Input][LP Input][MPS Input][QPS Input]
Mixed Integer Linear Programming
  * FEASPUMP [AMPL Input][CPLEX Input][MPS Input]
  * SCIP [AMPL Input][CPLEX Input][MPS Input] [ZIMPL Input]
  * qsopt_ex [LP Input][MPS Input] [AMPL Input]
Nondifferentiable Optimization   * condor [AMPL Input]
Semi-infinite Optimization       * nsips [AMPL Input]
Stochastic Linear Programming    * bnbs [SMPS Input]
                                 * DDSIP [LP Input][MPS Input]
```

We maintain the following NEOS solvers (cont.)

Semidefinite (and SOCP) Programming

* csdp [MATLAB_BINARY Input][SPARSE_SDPA Input]
* penbmi [MATLAB Input][MATLAB_BINARY Input]
* pensdp [MATLAB_BINARY Input][SPARSE_SDPA Input]
* sdpa [MATLAB_BINARY Input][SPARSE_SDPA Input]
* sdplr [MATLAB_BINARY Input][SDPLR Input][SPARSE_SDPA Input]
* sdpt3 [MATLAB_BINARY Input][SPARSE_SDPA Input]
* sedumi [MATLAB_BINARY Input][SPARSE_SDPA Input]

Overview of Talk

- **Current and Selected(*) Benchmarks**
  - Parallel LP benchmarks
  - MILP benchmark (MIPLIB2010)
  - Feasibility/Infeasibility Detection benchmarks (MIPLIB2010)
  - slightly pathological MILP cases
  - MI(QC)QP benchmark
- Conclusions

COMBINATORIAL OPTIMIZATION

Concorde-TSP with different LP solvers (8-16-2011)

LINEAR PROGRAMMING

Benchmark of serial LP solvers (11-9-2011)
* Benchmark of parallel LP solvers (2-25-2012)
* Parallel Barrier Solvers on Large LP/QP problems (3-2-2012)
Large Network-LP Benchmark (commercial vs free) (2-25-2012)

MIXED INTEGER LINEAR PROGRAMMING

* MILP Benchmark - MIPLIB2010 (2-22-2012)
* MILP cases that are slightly pathological (2-29-2012)
* Feasibility Benchmark (12-10-2011) (MIPLIB2010)
* Infeasibility Detection for MILP Problems (2-28-2012) (MIPLIB2010)

SEMIDEFINITE/SQL PROGRAMMING

Several SDP-codes on SDP problems with free variables (4-1-2008)
Several SDP codes on problems from SDPLIB (4-10-2008)
SQL problems from the 7th DIMACS Challenge (8-8-2002)
Newer SDP/SOCP-codes on the 7th DIMACS Challenge problems(4-7-2008)
Several SDP codes on sparse and other SDP problems (6-22-2011)
SOCP (second-order cone programming) Benchmark (3-2-2012)

NONLINEAR PROGRAMMING

Benchmark of commercial and other (QC)QP Solvers (11-16-2011)
AMPL-NLP Benchmark, IPOPT, KNITRO, LOQO, PENNLP, SNOPT & CONOPT (11-9-2011)

MIXED INTEGER NONLINEAR PROGRAMMING

* MI(QC)QP Benchmark (3-9-2012)

PROBLEMS WITH EQUILIBRIUM CONSTRAINTS

MPEC Benchmark (2-26-2012)

Important features of all our benchmarks

- Statistics of problems (dimensions etc)
- Links to codes given
- Links to test problems given
- Links to full logfiles given
- Same selection for commercial/free codes

Reasons for updates

- New version of commercial software
  - CPLEX-12.4, GUROBI-4.6.1, KNITRO-7.2.1, MOSEK-6.0.0.126/135
- New versions of free software
  - CBC, CLP, SCIP, SYMPHONY
  - BONMIN, COUENNE, IPOPT, FEASPUMP2
- More multicore hardware

Solving the MIPLIB 2010 Problems

- In late 2010 we had solved a number of instances
- In early 2012 we decided to solve those problems that a commercial solver with default options could solve in a day or less. We solved 13 cases plus 4 feasible solves.
- We used three platforms
  - Intel Xeon X5680 (3.33 GHz, 32GB, 12 cores)
  - Intel Xeon X5690 (3.47 GHz, 48GB, 6 cores)
  - Intel i7-2600 (3.4 GHz, 16GB, 4 cores)

instance	cores	time	category
b2c1s1	12/12	116575	hard
buildingenergy	1/4	895	easy
maxgasflow	6/6	66726	hard
opm2-z10-s2	12/12	10131	hard
opm2-z11-s8	12/12	18286	hard
opm2-z12-s14	4/4	49164	hard
opm2-z12-s7	12/12	72492	hard
r80x800	12/12	283303	hard
rmatr200-p10	4/4	19644	hard
satellites3-40-fs	4/6	32952	hard
satellites3-40	4/4	21205	hard
transportmoment	12/12	30	easy
wnq-n100-mw99-14	12/12	28124	hard

Overview of Talk

- **Current and Selected(*) Benchmarks**
  - **Parallel LP benchmarks**
  - MILP benchmark (MIPLIB2010)
  - Feasibility/Infeasibility Detection benchmarks (MIPLIB2010)
  - slightly pathological MILP cases
  - MI(QC)QP benchmark
- Observations and Conclusions

25 Feb 2012

Benchmark of parallel LP solvers

H. Mittelmann <mittelmann@asu.edu>

This benchmark was run on a Linux-PC (2.67 GHz Intel Core 2 Quad).
The MPS-datafiles for all testcases are in one of (see column "s")

miplib.zib.de/ [1]
plato.asu.edu/ftp/lptestset/ [2]
www.netlib.org/lp/data/ [3,7]
www.sztaki.hu/~meszaros/public_ftp/lptestset/
(MISC[4], PROBLEMATIC[5], STOCHLP[6], INFEAS[8])

The (dual) simplex, barrier, and concurrent methods were tested of:

CPLEX-12.4.0.0	CPLEX
GUROBI-4.6.1	www.gurobi.com/
MOSEK-6.0.0.126	www.mosek.com
XPRESS-7.2.1:	XPRESS

Scaled geometric mean of runtimes

	4.11	2.78	4.43	3.27	1.05	1.10	1.49	1.69	1.78	1	1.25
s	problem	CPX-S	GRB-S	MSK-S	XPR-S	CPX-B	GRB-B	MSK-B	XPR-B	CPX-A	GRB-A	MSK-A

2 Mar 2012

Parallel Barrier Solvers on Large LP/QP problems

H. Mittelmann (mittelmann@asu.edu)

Logiles at plato.asu.edu/ftp/barrier_logs/

CPLEX-12.4.0.0: CPLEX
GUROBI-4.6.0: GUROBI
MOSEK-6.0.0.126: MOSEK
XPRESS-7.2.1: XPRESS
IPOPT-3.10.1: IPOPT

The barrier methods (w/o crossover) of the above solvers were run on
a 3.47 GHz Intel Xeon X5690 (6 cores, 48GB) on large LP problems from here.
Times given are elapsed times in seconds.

2 Mar 2012

Parallel Barrier Solvers on Large LP/QP problems

problem	CPLEX	GUROBI	MOSEK	XPRESS	IPOPT
in	46454	190	879	309	
pde_10	2271	747	3128	15906	
pde_1	2516a	2033	6625	fail	
pde_200	299a	528	1237	1528	
pde_20	1896a	2350	13503	3186	
pde_2	5269a	16200	72254	9889	
qap_2	slow	fail	slow	1581	
srd300	10091	fail	18705	>60000	
zib01	6010a	8251	9142	7113	
L2CTA3D	3271	6459	5673	2169	
bdry2_0	1925a	995	5123a	2412	4641
bdry2_1	10977a	3919	fail	10111	31089
cont5_2_0	512	1113	27	564	1186
cont5_2_1	1746	4118a	57a	1545	3270
cont5_2_2	7979	fail	127	4181	11579
twod_00	2499	8650	12525a	2848	9348
twod_0	12391	47586	fail	12841	41616

"a": reduced accuracy

2 Mar 2012

Parallel Barrier Solvers on Large LP/QP problems

problem	constraints	variables	nonzeros	MPS-file
in	1526203	1449074	6813128	561 MB
pde_10	13225799	4414198	39669001	1.96 GB
pde_1	26993999	9005998	80970001	4.1 GB
pde_200	5067399	3382596	15197001	824 MB
pde_20	13225799	8824196	39669001	2.2 GB
pde_2	26993999	18005996	80970001	4.5 GB
qap_2	57360	14400	203040	70.7 MB
srd300	1101467	2052911	59170079	1.6 GB
zib01	5887041	12471400	49877768	2.5 GB
L2CTA3D	210000	10000000	30000000	1.8 GB
bdry2_0	4001998	4003997	19991999	
bdry2_1	12253498	12256997	61235999	
cont5_2_0	1959681	1961081	11749904	
cont5_2_1	3999656	4001656	23986126	
cont5_2_2	8999641	9002641	53979959	
twod_00	505284	511446	4856046	
twod_0	989604	999306	9586066	

Overview of Talk

- **Current and Selected(*) Benchmarks**
  - Parallel LP benchmarks
  - **MILP benchmark (MIPLIB2010)**
  - Feasibility/Infeasibility Detection benchmarks (MIPLIB2010)
  - slightly pathological MILP cases
  - MI(QC)QP benchmark
- Conclusions

22 Feb 2012 Mixed Integer Linear Programming Benchmark (MIPLIB2010)

The following codes were run on the MIPLIB2010 benchmark set with the MIPLIB2010 scripts on an Intel Xeon X5680 (32GB, Linux, 64 bits, 2*6 cores), with one, four and twelve threads. (deterministically) and a time limit of 1 hour. These are up dated and extended versions of the results produced for the MIPLIB2010 paper.

CPLEX-12.4.0.0: CPLEX
GUROBI-4.6.1: GUROBI
ug[SCIP/cpx/spx]: Parallel development version of SCIP (SCIP+CPLEX/SOPLEX/CLP on 1 thread)
CBC-2.7.6: CBC
XPRESS-7.2.1: XPRESS
GLPK-4.47: GLPK
LP_SOLVE-5.5.2: LP_SOLVE

Table for single thread, Result files per solver, Log files per solver

Table for 4 threads, Result files per solver, Log files per solver

Table for 12 threads, Result files per solver, Log files per solver

Statistics of the problems can be obtained from the MIPLIB2010 webpage.

++
Geometric means of times

All non-successes are counted as max-time. The fastest solver is scaled to 1.
The second line lists the number of problems (87 total) solved.

threads	CBC	CPLEX	GLPK	GUROBI	LPSOLVE	SCIPC	SCIPL	SCIPS	XPRESS
1	10.3	1.45	22.11	1	19.4	3.76	6.4	5.33	1.28
solved	41	73	3	77	5	63	52	57	74

threads	CBC	CPLEX	FSCIPC	FSCIPS	GUROBI	XPRESS
4	9.89	1	5.83	8.43	1	1.23
solved	52	84	66	63	85	79

threads	CBC	CPLEX	FSCIPC	FSCIPS	GUROBI	XPRESS
12	12.3	1.1	8.64	12.7	1	1.33
solved	56	84	71	65	87	83

Overview of Talk

- **Current and Selected(*) Benchmarks**
  - Parallel LP benchmarks
  - MILP benchmark (MIPLIB2010)
  - **Feasibility/Infeasibility Detection benchmarks (MIPLIB2010)**
  - slightly pathological MILP cases
  - MI(QC)QP benchmark
- Conclusions

10 Dec 2011 === Feasibility Benchmark ===
H. Mittelmann (mittelmann@asu.edu)

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

MILP problems from MIPLIB2010 were solved for a feasible point

The following codes were run on an Intel i7-2600 (3.4 GHz, 16GB, Linux, 4 cores) with 4 threads:

CPLEX-12.4.0.0: CPLEX
FEASPUMP2: as implemented for interactive use at NEOS (utilizes CPLEX)
GUROBI-4.6.1: GUROBI
XPRESS-7.2.1: XPRESS
CBC-2.7.5: CBC

Times given are elapsed times in seconds. A time limit of 1 hr was imposed.
Geometric means of the times are listed. For objective values see logfiles.

problem(30 tot)	CPLEX	FP2	GUROBI	XPRESS	CBC
geometric mean	6.97	14.3	1	14.85	50.9
problems solved	22	17	27	19	11

26 Feb 2012 Infeasibility Detection for MILP Problems

The following codes were run on the infeasible problems from MIPLIB2010 with the MIPLIB2010 scripts

CPLEX-12.4.0.0: CPLEX
GUROBI-4.6.1: GUROBI
ug[SCIP/spx/cpx]: Parallel development version of SCIP
CBC-2.7.4: CBC
XPRESS-7.2.1: XPRESS

Table for 12 threads, Result files per solver, Log files per solver

Statistics of the problems can be obtained from the MIPLIB2010 webpage.

++
Geometric means of times (19 instances)

All non-successes are counted as max-time. The fastest solver is scaled to 1.

	CBC	CPLEX	FSCIPC	FSCIPS	GUROBI	XPRESS
	79.3	1.68	17.7	30.5	1.00	4.74
solved	6	16	13	11	18	15

Overview of Talk

- **Current and Selected(*) Benchmarks**
  - Parallel LP benchmarks
  - MILP benchmark (MIPLIB2010)
  - Feasibility/Infeasibility Detection benchmarks (MIPLIB2010)
  - **slightly pathological MILP cases**
  - MI(QC)QP benchmark
- Conclusions

29 Feb 2012

MILP cases that are slightly pathological

H. Mittelmann (mittelmann@asu.edu)

CPLEX-12.4.0.0: CPLEX to be added
GUROBI-4.6.1: GUROBI
ug[SCIP/spx]: FSCIP-Parallel development version of SCIP
CBC-2.7.5: CBC
XPRESS-7.2.1: XPRESS
SCIP-2.1.1: serial SCIP with CPLEX

These codes were run with the MIPLIB2010 scripts in default mode on an Intel Xeon X5680 (32GB, Linux, 64 bits, 2*6 cores) on problems from here. Times given are elapsed CPU seconds. Time limit 3 hrs. Available memory 24GB. This benchmark is not giving a representative impression of the relative performance of the codes.

Table for 12 threads, Result files per solver, Log files per solver

Scaled geometric mean of runtimes and problems solved (24 total)

CBC	CPLEX	FSCIP	GUROBI	SCIP	XPRESS
11	1	5.15	1.38	7.4	3.39
4	17	13	20	14	10

Overview of Talk

- **Current and Selected(*) Benchmarks**
  - Parallel LP benchmarks
  - MILP benchmark (MIPLIB2010)
  - Feasibility/Infeasibility Detection benchmarks (MIPLIB2010)
  - slightly pathological MILP cases
  - **MI(QC)QP benchmark**
- Conclusions

9 Mar 2012

Mixed Integer (QC)QP Benchmark

H. Mittelmann (mittelmann@asu.edu)

Logiles at plato.asu.edu/ftp/miqp_logs/

The MPS-datafiles are in plato.asu.edu/ftp/miqp/ and the AMPL files in plato.asu.edu/ftp/ampl_files/miqp_ampl/ and egon.cheme.cmu.edu/ibm/files/

The following codes were run in default mode on a 2.66GHz Intel Core2 Quad, For accuracy reached, see logfiles. In the last columns are results for the QCQPs obtained by rewriting the QPs as: min t, subject to quadratic obj <= t plus constraints. SCIP does this transformation itself. SCIP uses relgap=0;

CPLEX-12.4.0.0: CPLEX
Bonmin-1.5.3: projects.coin-or.org/Bonmin (Bonmin: hybrid algorithm. with Cbc)
Couenne-0.4.2: projects.coin-or.org/Couenne
FilMINT: currently only at NEOS, (run locally)
GUROBI-4.6.1: gurobi.com
SCIP-2.1.1: scip.zib.de (with CPLEX and IPOPT)
XPRESS-7.2.1: XPRESS

Times given are user times in seconds. Time limit 10,800 seconds.

QP
==

Scaled geometric means of runtimes

53.9	239	2.68	261	1	14.8	1.67

problem	Bonm	Couen	CPLEX	FilMI	GUROBI	SCIP	XPRESS

QCQP
====

Scaled geometric means of runtimes

12.9	18.7	2.67	32.7	1	

problem	Bonmin	Couenne	CPLEX	FilMINT	XPRESS

Observations and Conclusions

- Observation: TTB (**Tuning To the Benchmarks**3
  - CPLEX-12.3 -> CPLEX-12.4preview
    - better for MIPLIB benchmarks
    - worse for pathological ones
  - MOSEK-6.0.0.122 -> 6.0,0.126 Changelog
    - Improved the interior point optimizers for linear problems for certain hard problems.

- Conclusions: **Declare Winners?**
  - **MIPLIB-bench**: CPLEX/Gurobi-XPRESS
  - **MIPLIB-feas**: Gurobi-CPLEX-XPRESS, FEASPUMP close
  - **MIPLIB-infeas**: Gurobi, CPLEX close, XPRESS slow
  - **Pathological**: Gurobi-CPLEX, SCIP,FSCIP respectable
  - **MIQP**: Gurobi
  - **MIQCQP**: XPRESS
  - **LP**: CPLEX/Gurobi; **QP/QCQP**: XPRESS/MOSEK

Thank you!