Latest Progress in Optimization Software

Hans D Mittelmann

School of Mathematical and Statistical Sciences
Arizona State University

INFORMS Annual Meeting
Seattle, WA
21 October 2024

Outline

Outline

DECISION TREE FOR OPTIMIZATION SOFTWARE

BENCHMARKS FOR OPTIMIZATION SOFTWARE

By Hans Mittelmann (mittelmann at asu.edu)

END OF A BENCHMARKING ERA

For many years our benchmarking effort had included the solvers CPLEX, Gurobi, and XPRESS. Through an action by Gurobi at the 2018 INFORMS Annual Meeting this has come to an end. IBM and FICO demanded that results for their solvers be removed. See here for more details. The resulting void was filled by other developers.

In August 2024 Gurobi decided to withdraw from the benchmarks as well and their results have been removed. See the note at the bottom of the MIPLIB benchmark.

What is on top of the Benchmark Page?

- Links to history and older benchmarks
- Especially details on the elimination of CPLEX, XPRESS, and recently Gurobi
- Hint that full solver logfiles are provided
- Explanation why performance profiles are not used
- Link to Matt Miltenberger's visualization tool

History and Scope of the Benchmarks

- Start about 1998
- Summary presented at INFORMS Annual starting 2002
- 22 benchmarks in these areas
  - TSP, (MI)LP, (MI)SOCP, SDP, convex/nonconvex (MI)QCP, QUBO, (MI)NLP, MPEC
- Codes evaluated: 48, of those 28 actively developed, 14 commercial
- Running times per code: up to 3 weeks

Outline

The Benchmarks at a Turning Point

- Withdrawals
  - Gurobi 8/2024 from all
  - Mindopt/Alibaba 9/2024 from MIP
- Potential use of over-tuning/machine learning
  - in LP/MIP: various codes ?
  - in NLP: Taylor/Huawei ?
- Measures taken
  - one quarter of LP instances undisclosed; should help for MIP as well
  - preprocessing of MIP instances (mathematically equivalent)
  - issues may be addressed by MIPLIB2024 committee
  - in NLP nothing so far

Outline

Selected Benchmarks

These seven benchmarks follow next:

- convex: LPopt, MILP, MIQCP, SDP
- nonconvex: NLP, BQCP, MINLP
- remember: "the great watershed in optimization isn't between linearity and nonlinearity, but convexity and nonconvexity"
- Talk will report some progress; more substantial progress may be announced at meeting: by Gurobi, XPRESS etc

LPopt Benchmark (find optimal basic solution)

- Same own instance selection as in LPfeas benchmark

  CLP-1.17.7, COPT-7.2.0, MindOpt-1.0.0, HiGHS-1.7.2, OptVerse-1.1.0

  MOSEK-10.1.9, GLOP-9.5, SOPLEX-6.0.0, XOPT-0.0.8

- Intel i7-11700K, 3.6GHz, 64GB, 15,000 secs wall clock

```
65 probs    27.2     1  1.97  1.68  5.84  18.2  59.4  90.9  6.60
solved       40    65    63    63    52    51    33    32    52
================================================================
probs        CLP  COPT MDOPT  OPTV MOSEK HiGHS  GLOP  SPLX  XOPT
================================================================
MATLAB uses HiGHS
```

The MILP Benchmarks: Benchmark and Pathological

- Total of 240 resp 45 instances
- Sizes up to 1.5m/1m/43m rows/cols/nonzeros
- AMD Ryzen 9 5900X (12 cores, 128GB), 7,200 resp 10,800 secs wall clock
- Instead of separate tables, combine both benchmarks
- for the original and the modified datasets
- modification through random row and column reordering
- and a quick presolve, both done with SCIP.

	COPT	SCIP	SCIPC	HiGHS	OPTV	XOPT	LEOPT	TAYLR	
scaled	1.11	22.4	13.0	11.8	1	7.90	12.4	2.72	PATH
solved	41	20	28	29	41	31	27	34	
scaled	1	20.3	11.7	11.3	2.59	9.90	11.2	3.57	P_PATH
solved	41	22	25	23	33	24	24	30	

	COPT	SCIP	SCIPC	HiGHS	OPTV	XOPT	LEOPT	TAYLR	
scaled	1.05	9.03	7.71	7.27	1	4.62	3.04	1.78	MIPLIB
solved	220	139	152	159	222	188	179	227	
scaled	1	8.76	7.07	6.79	1.89	5.05	6.69	3.39	MODIFIED
solved	217	128	145	151	203	160	150	181	

Convex Discrete QPLIB Benchmark

- Total of 31 instances
- Sizes up to 26k/65k/200k variables/constraints/nonzeros
- AMD Ryzen 9 5900X, 12 cores, 128GB , 2 hours wall clock

	MOSEK	KNITRO	BARON	BONMIN	SCIP	MNTAUR	SHOT	COPT
mean	9.66	12.3	4.22	52.6	21.0	61.6	1.03	1
solved	18	16	20	7	14	8	25	24

Several SDP-codes on sparse and other SDP problems

- Total of 75 instances, Own instance selection
- Sizes up to 100k/100k/450m variables/constraints/nonzeros
- AMD Ryzen 9 5900X, 12 cores, 128GB , 40,000 secs wall clock

scaled geomean	1	5.21	3.22	10.5	5.14	28.9	7.86	1.44
count of "a"	6	5	2	17	13	2	11	12
solved of 75	75	70	73	61	69	62	70	75

problem	COPT	CSDP	MOSEK	SDPA	SDPT3	SeDuMi	HDSDP	MDOPT

"a": reduced accuracy

AMPL-NLP Benchmark

- Total of 47 instances
- Sizes up to 140k/260k constraints/variables
- AMD Ryzen 9 5900X (12 cores, 128GB), 7,200 secs wall clock

scaled geom mean	10.6	1.28	94.4	82.4	10.9	34.0	1
solved	46	47	30	27	45	36	47

47 problems	IPOPT$	KNITR	SNOPT	CONPT	WORHP	MATLAB	TAYLOR
scaled geom mean	8.29	1	73.6	64.3	8.49	26.5	
solved	46	47	30	27	45	36	

$ Higher performance of IPOPT possible for different LA routines

Binary Non-Convex QPLIB Benchmark

- Total of 128 instances
- Sizes up to 36k/120k/5k variables/total constraints/nonlin constraints
- AMD Ryzen 9 5900X (12 cores, 128GB), 3,600 secs wall clock

mean	5.87	31.7	63.1	74.9	1	10.3
solved	66	34	16	6	91	69

prob#	BARON	SCIP	ANTIGONE	COUENNE	SHOT	RAPOSa

BARON/SCIP with CPLEX; SHOT, RAPOSa with Gurobi

Mixed Integer Nonlinear Programming Benchmark

- Total of 87 instances; run through GAMS
- Sizes up to 100k/100k/5k variables/lin constraints/nonlin constraints
- AMD Ryzen 9 5900X (12 cores, 128GB), 7,200 secs wall clock
- Instances from MINLPLib

	ANTIGONE	BARON	LINDO	SCIP	SHOT
geomean	19.0	1.0	22.3	4.6	2.3
solved	52	76	28	65	53

#: SHOT can only solve 66/87 instances

Outline

Conclusions

Some Takeaways

- transition to post-MIPLIB2017 regime partially achieved
- regretfully Mindopt withdrew from MILP
- even convex MIQCP is still challenging
- regretfully Gurobi withdrew from MIQCP
- BARON continues to lead in MINLP; where are Gurobi, XPRESS?
- GPUs still to be included

THANK YOU

Questions?

Slides of talk at
https://plato.asu.edu/talks/informs2024.pdf