

GUROBI
OPTIMIZATION

Gurobi 8 Performance Benchmarks

February 11, 2019

Thank You for Your Interest in Gurobi



The Gurobi Optimizer was designed from the ground up to be the fastest, most powerful solver available for your MIP (MILP, MIQP, and MIQCP), LP and QP problems.

- In industry standard public benchmark tests Gurobi has the...
 - Fastest overall solve times for MIP models
 - Fastest overall solve times for LP models
 - Fastest overall solve times for QP models

And, as problems get harder, our relative performance gets even better.

Two Types of Benchmark Testing



Internal

- Primary Objectives
 - Robustness testing
 - Compare version-to-version improvements
- Test Bank
 - Internal library of over 10,000 models from industry and academia

Public

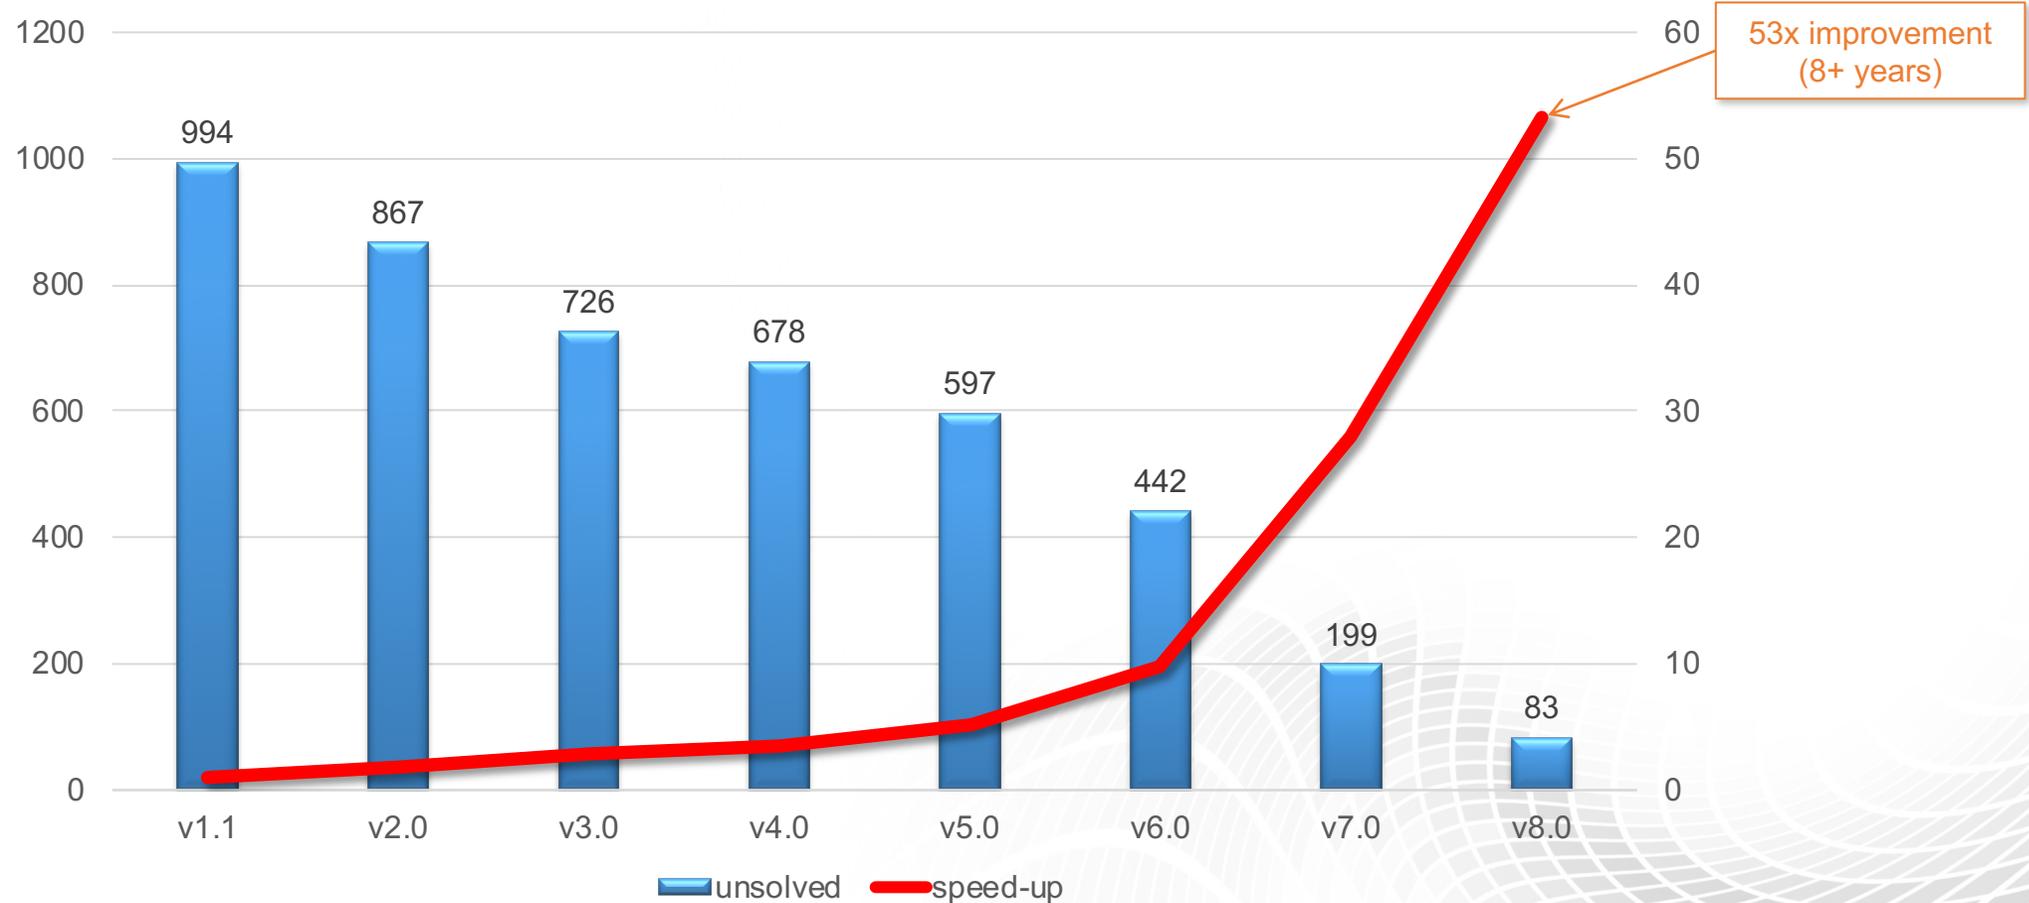
- Primary Objective
 - Competitive benchmarks against other solvers
- Test Bank
 - Based on MIPLIB 2010 and 2017
 - Tests performed by Prof. Hans Mittelmann, Arizona State University

On the next slides we'll share some specific results as well as results from our own internal testing. Of course, every model is different so we invite you to [try Gurobi for yourself](#) or [contact us](#) with any questions.

Gurobi Keeps Getting Better



Comparison of Gurobi Versions (PAR10)



Time limit: 10000 sec.
Intel Xeon CPU E3-1240 v3 @ 3.40GHz
4 cores, 8 hyper-threads
32 GB RAM

Test set has 5656 models:
- 410 discarded due to inconsistent answers
- 1741 solved in < 100s by all the versions
- 1493 discarded that none of the versions can solve
- speed-up measured on >100s bracket: 2012 models

Broad Performance Improvements in v8.0

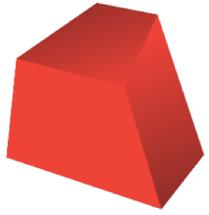


- Consistent with prior releases, the Gurobi Optimizer v8.0 delivers performance improvements over v7.0 across a broad range of model types:
 - **MIP** – **57%** faster overall, 109% faster on models that take >100 seconds to solve
 - **LP**
 - **Concurrent** – **15%** faster overall, 46% faster on models that take >100 seconds to solve
 - **Primal Simplex** – **24%** faster overall, 49% faster on models that take >100 seconds to solve
 - **Dual Simplex** – **32%** faster overall, 82% faster on models that take >100s to solve
 - **Barrier** – **13%** faster overall, 44% faster on models that take >100s to solve
 - **MIQP** – **2.76x** faster overall
 - **MIQCP** – **20%** faster overall
 - **SOCP** – **19%** faster overall

Performance Improvements in v8.1



- Gurobi Optimizer v8.1 improves performance significantly over v8.0 on integer quadratic programs:
 - **MIQP** – **2.8x** faster overall
 - **MIQCP** – **38%** faster overall
 - **LP** – both **dual simplex** and **barrier** are **10%** faster on models that take > 100 seconds to solve



GUROBI
OPTIMIZATION

MILP Competitive Benchmarks

Gurobi 8.1.0 vs. CPLEX 12.8.0 vs. XPRESS 8.5.1
Tests performed by Prof. Hans Mittelmann

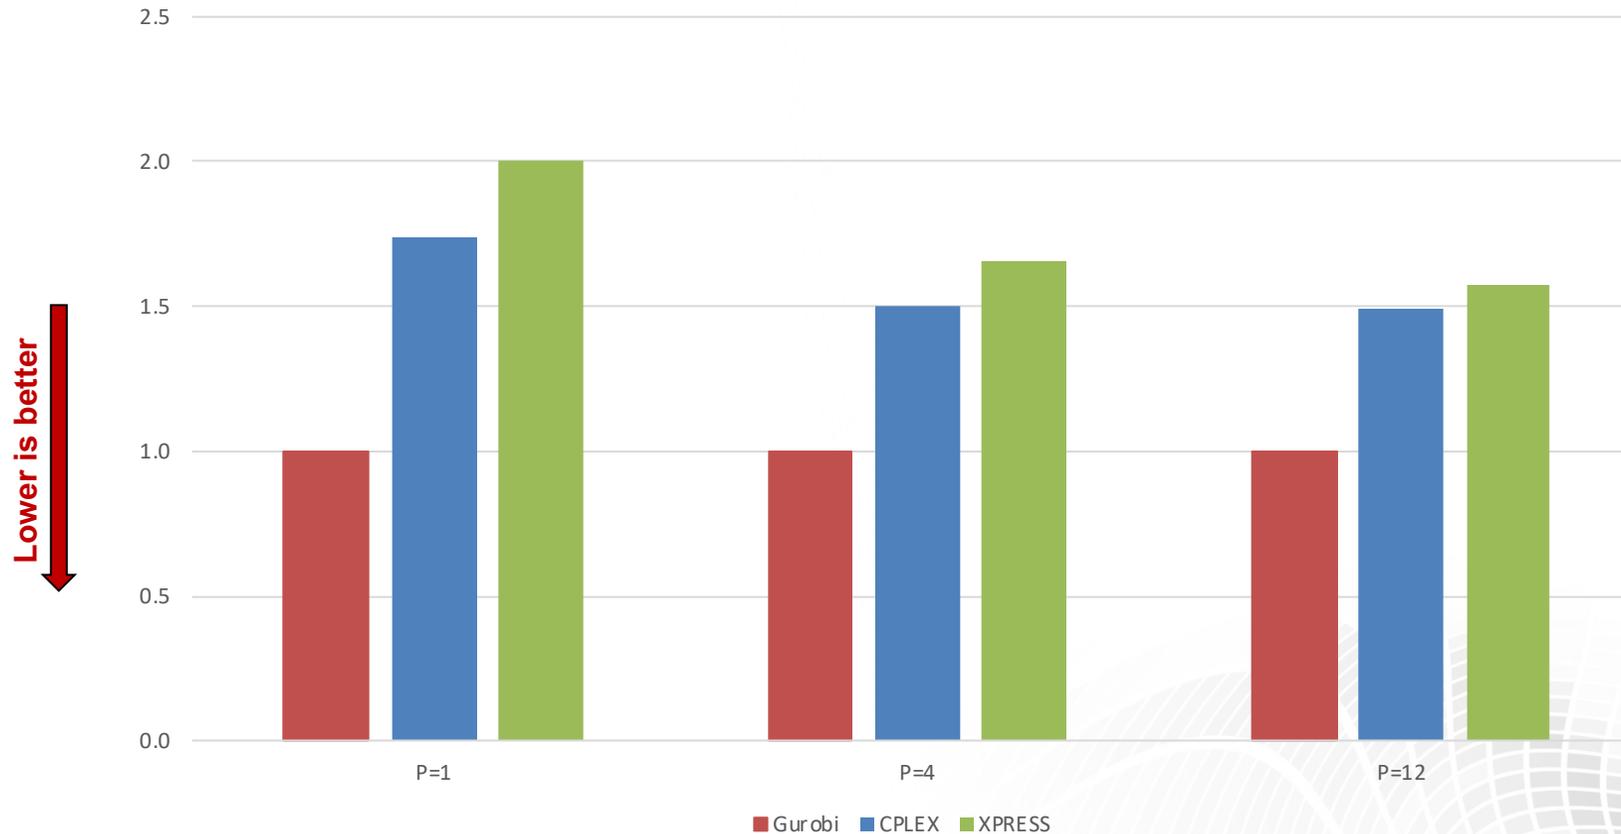
Gurobi is...

- *Fastest to optimality (MIPLIB 2010 benchmark)*
- *Fastest on the new MIPLIB 2017 benchmark*
- *Fastest to feasibility (MIPLIB 2010 feasibility benchmark)*
- *Fastest to infeasibility (MIPLIB 2010 infeasibility benchmark)*

Gurobi is Fastest on MIPLIB 2010 Benchmark



Optimality Benchmark Comparison

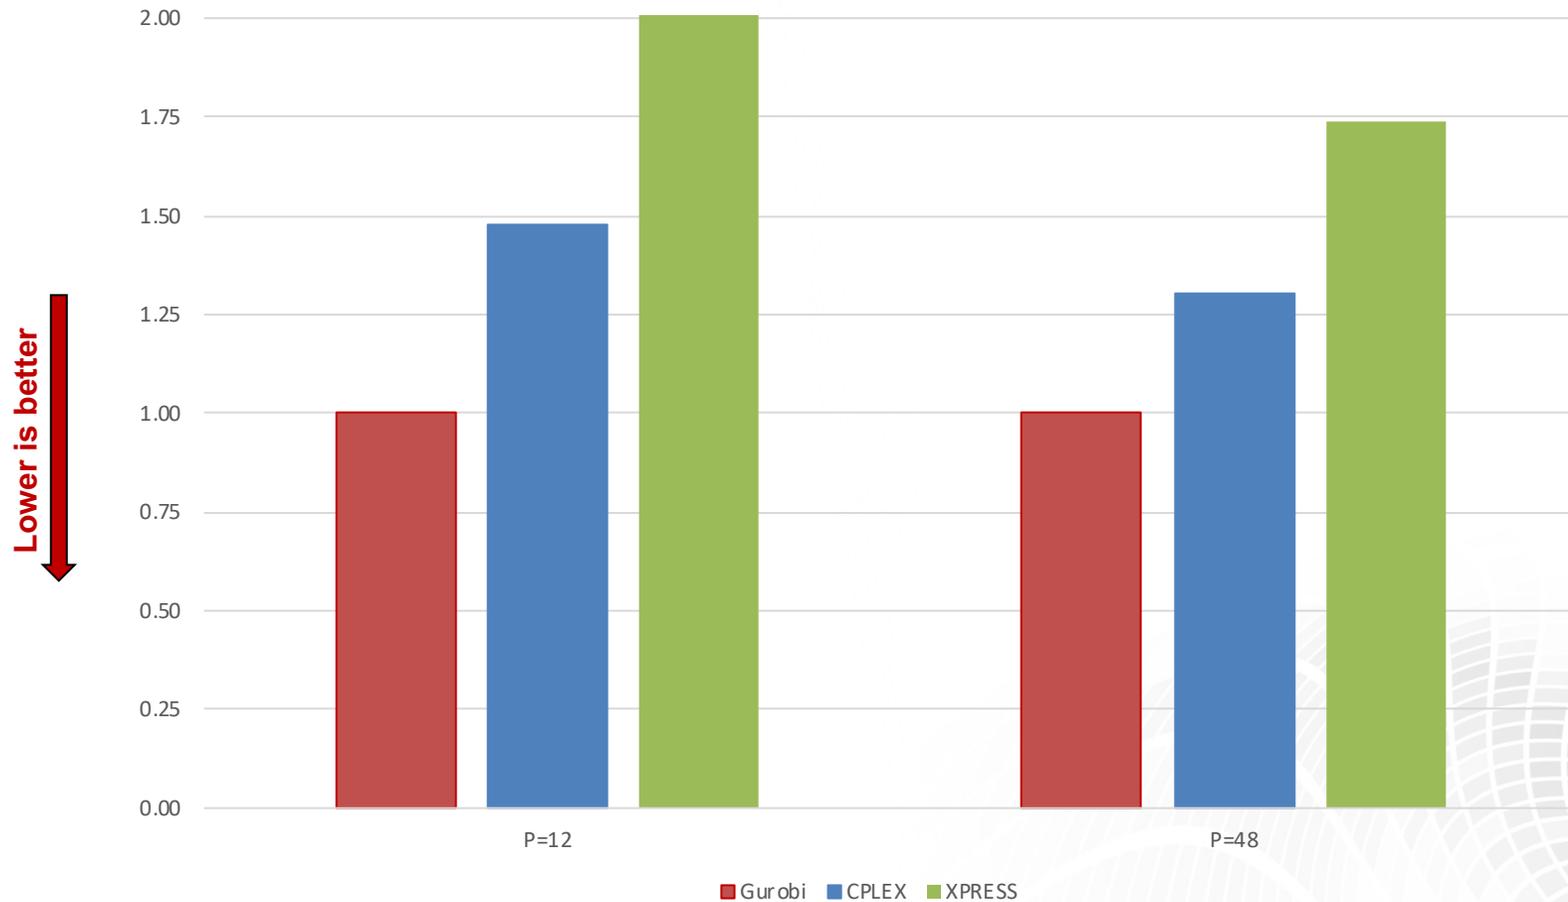


For example, on time to optimality benchmark (87 models) using 4 threads (P=4), CPLEX was 50% slower (1.50) and XPRESS was 66% slower (1.66) than Gurobi.

Gurobi is Fastest on the New MIPLIB 2017 Benchmark



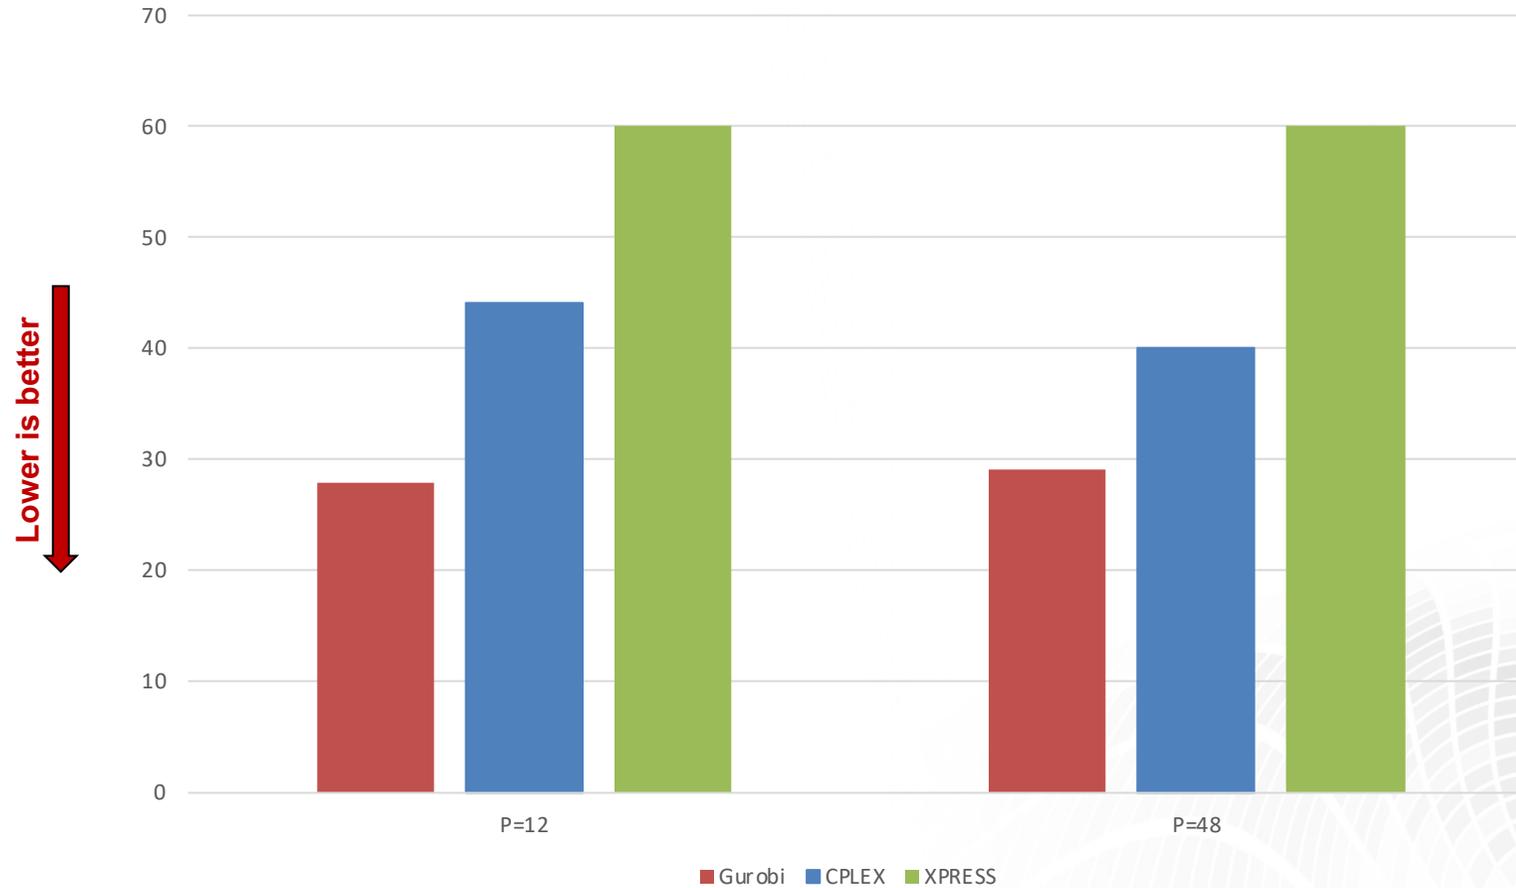
MIPLIB 2017 Benchmark Comparison



Gurobi Solves More Models in MIPLIB 2017 Benchmark



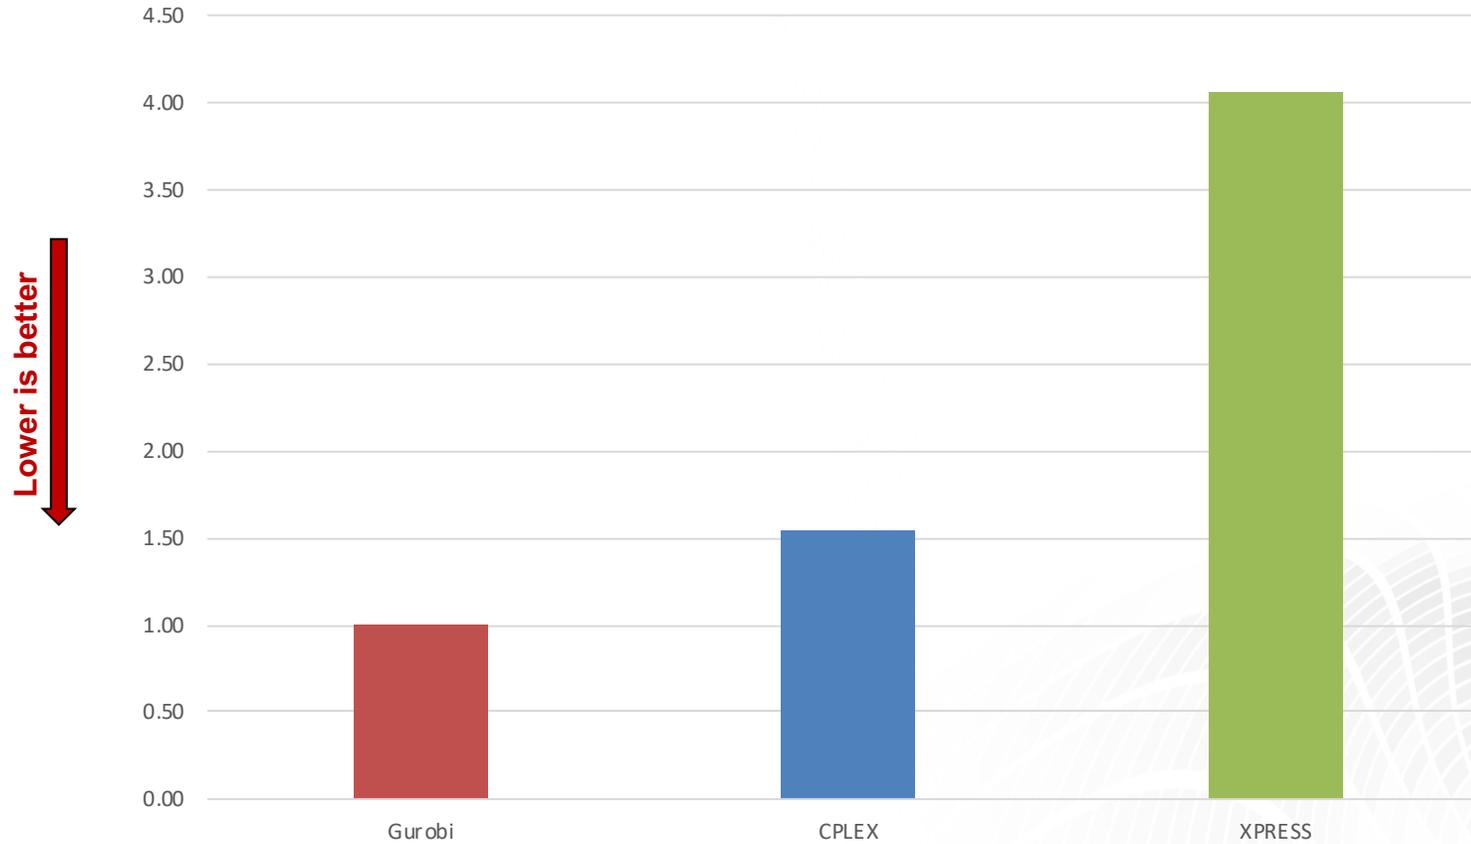
Number of Models not Solved



Gurobi is Fastest to Feasibility

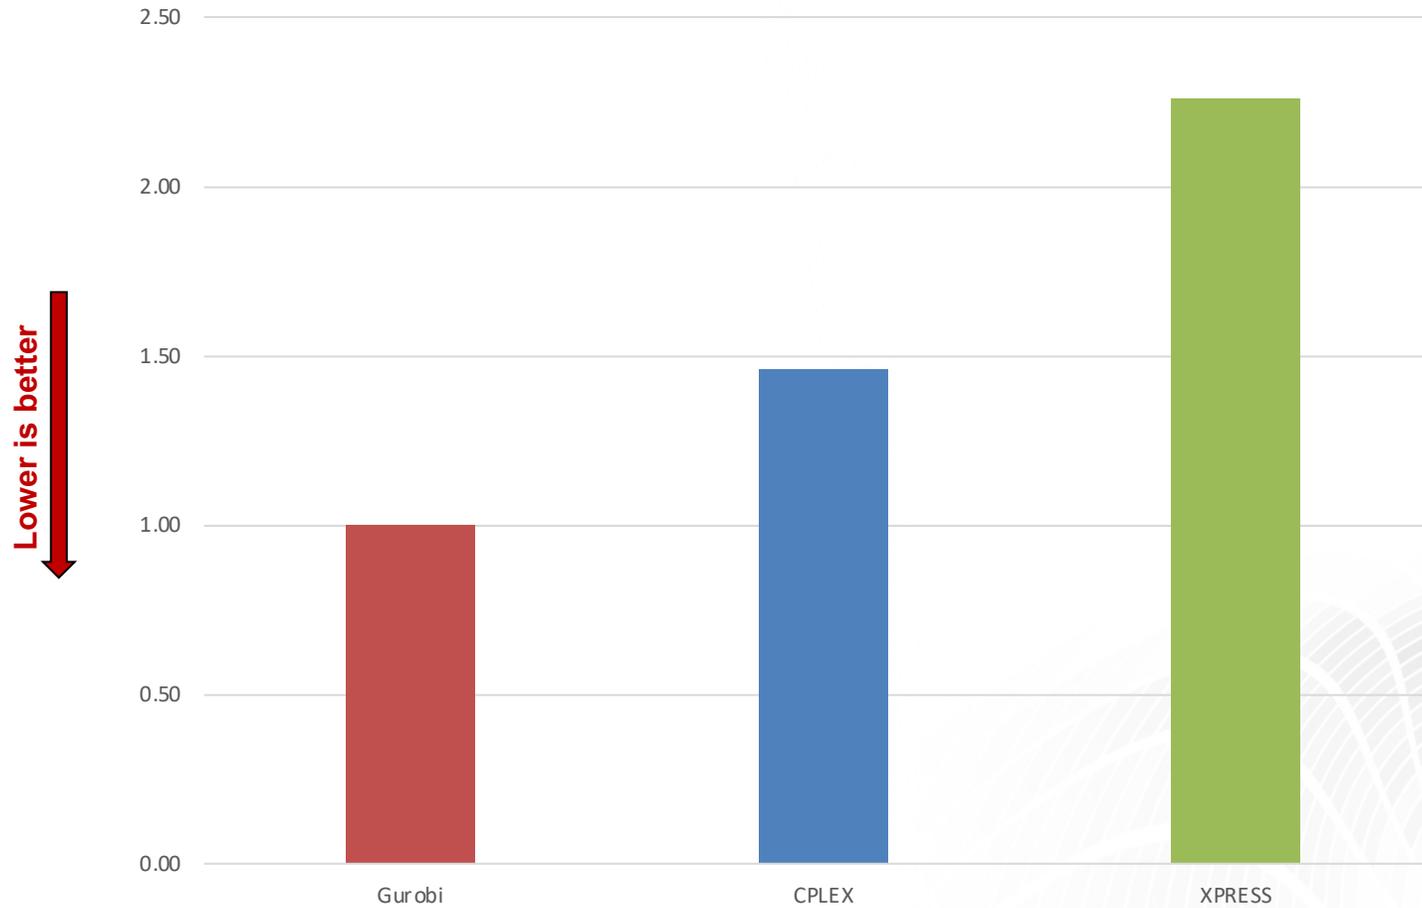


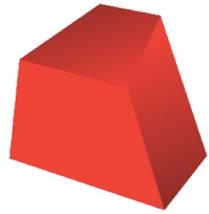
Feasibility Benchmark Comparison (P=4)



Gurobi is Fastest to Detect Infeasibility

Infeasibility Benchmark Comparison (P=4)





GUROBI
OPTIMIZATION

LP Competitive Benchmarks

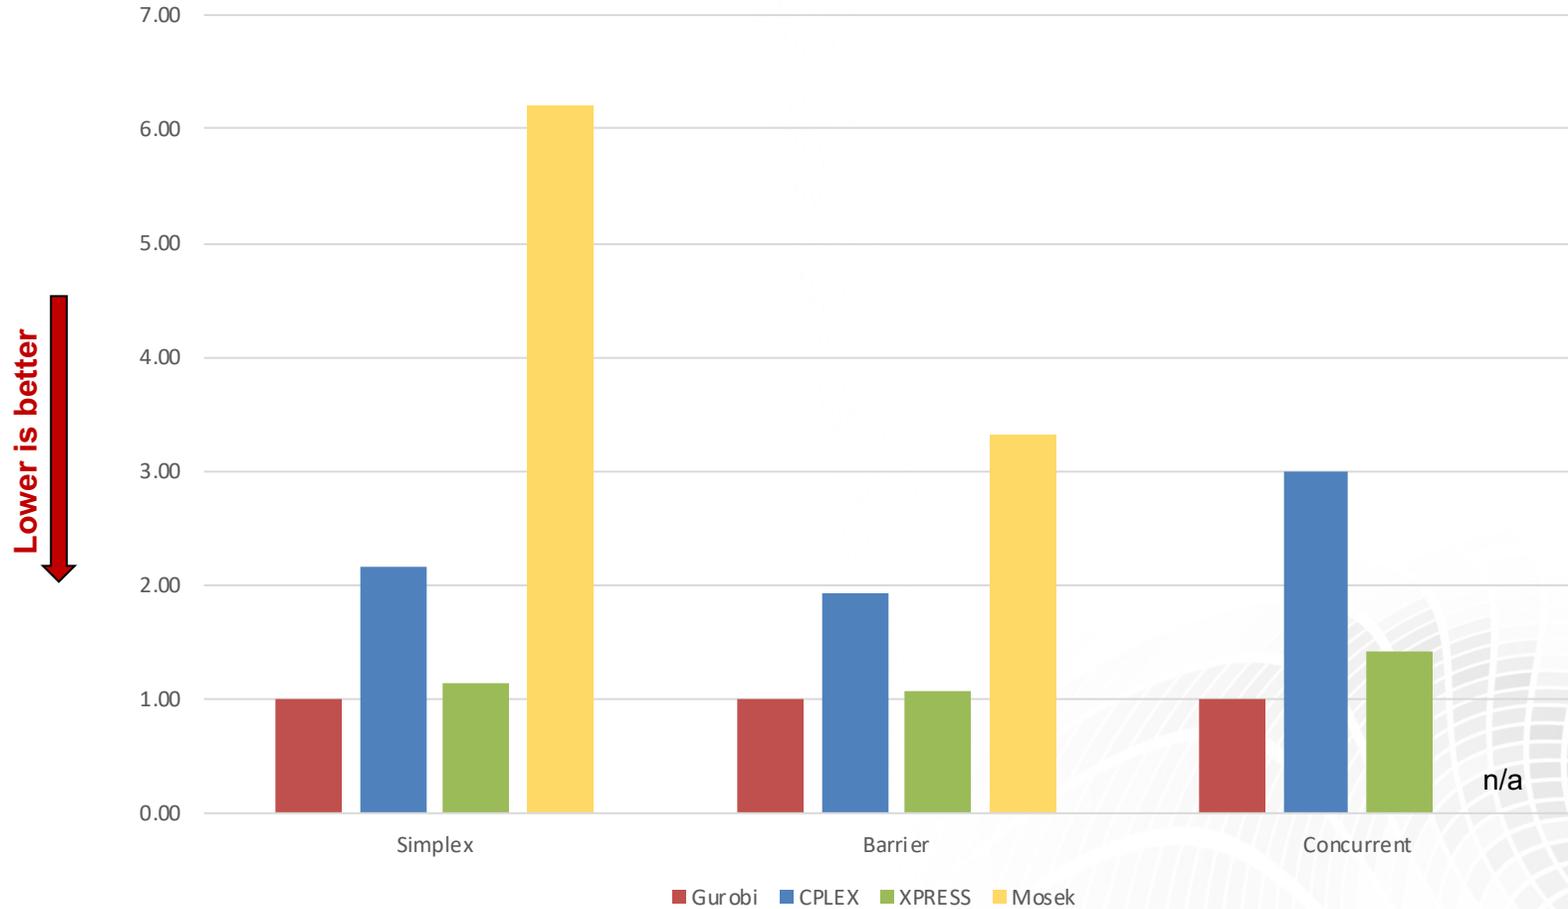
Gurobi 8.1.0 vs. CPLEX 12.8.0 vs. XPRESS 8.5.1 vs. Mosek 8.1.0.x
Tests performed by Prof. Hans Mittelmann

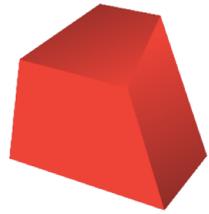
Gurobi has the fastest solve times

Gurobi is Fastest Across All LP Benchmarks



LP Solver Benchmark Comparison





GUROBI
OPTIMIZATION

QP Competitive Benchmarks

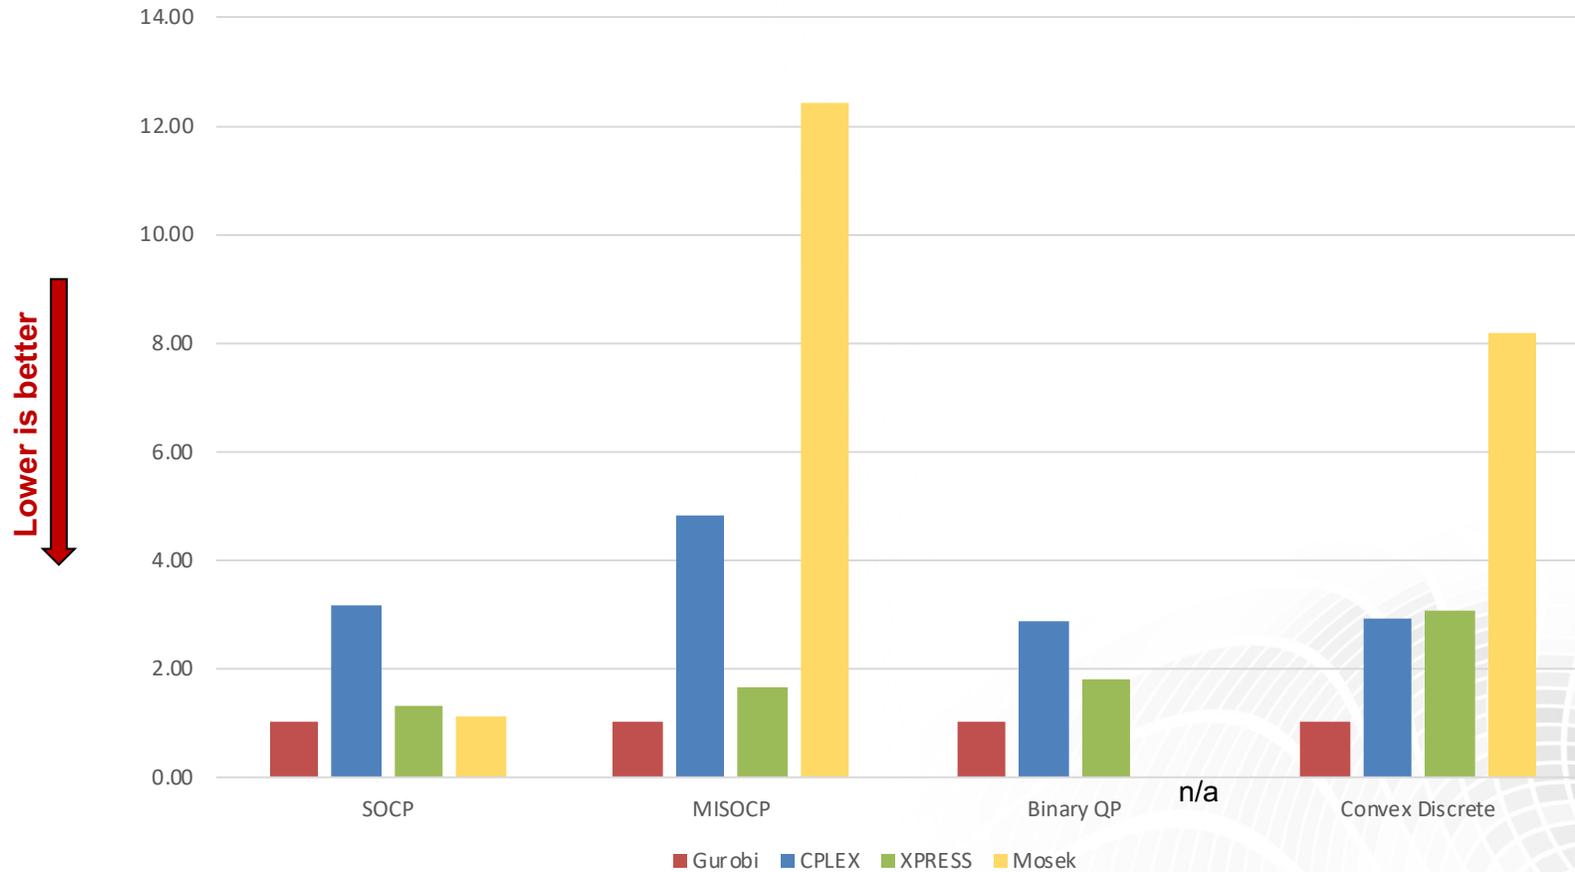
Gurobi 8.1.0 vs. CPLEX 12.8.0 vs. XPRESS 8.5.1 vs. Mosek 8.1.0.x
Tests performed by Prof. Hans Mittelmann

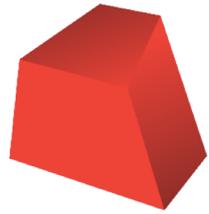
Gurobi has the fastest solve times

Gurobi is Fastest Across All QP Benchmarks



Quadratic Model Solve Times

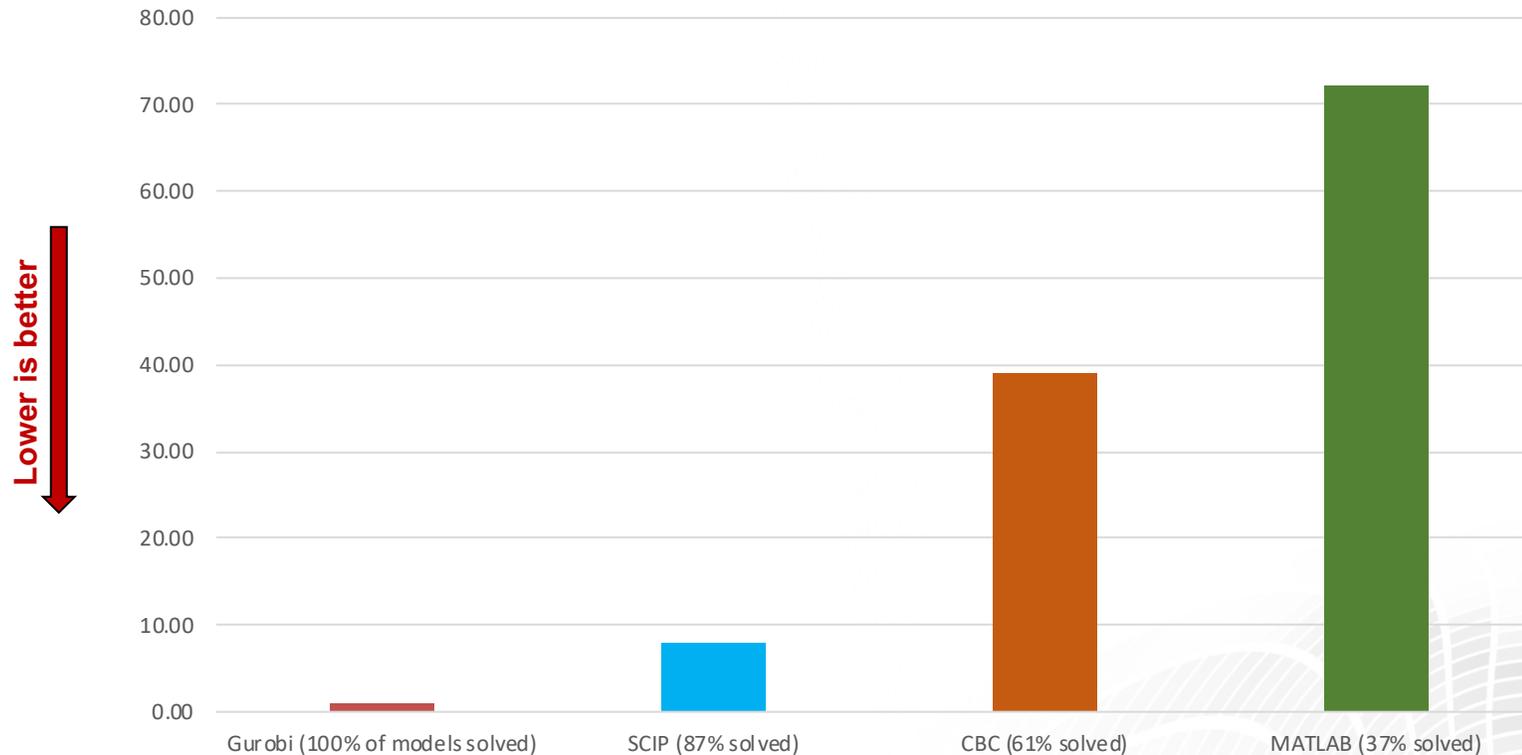




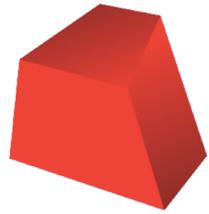
GUROBI
OPTIMIZATION

Open Source Solvers and MATLAB Benchmarks

Optimality vs. Open-source Solvers and MATLAB



- ▶ LPSolve and GLPK are not included here as they solved too few models (seven and two respectively) to calculate useful performance comparisons.
- ▶ Open-source and MATLAB solver performance is worse than shown as unsolved models are treated as solved at max time limit.



GUROBI
OPTIMIZATION

Pushing Performance Even Higher

Taking advantage of Gurobi's
Parameter Tuning and Distributed Optimization capabilities

Tuning can have a significant positive impact on performance results

- Test Set: MIPLIB 2010 benchmark, 87 models
 - Default tuning run with TuneTrial=1, each model separately
 - It uses 10X of default solving time
 - Two tuning runs, one with a single machine, one with 5 machines
- Results: (> 1 means faster)
 - Mean improvement from the best settings:
 - A single machine: **1.68X**
 - 5 machines: **2.52X**

Gurobi gives you industry-leading out-of-the-box performance. However, you can take that performance up even higher by tuning Gurobi's parameters for your model(s).

To help you do that we provide an automatic tuning tool you can run on just one machine (the 1.68X performance improvement you see above across the test set), or on a number of machines (the 2.52X improvement in the five machine example above).

Note, the performance gain on your particular model could be higher or lower than the test results above. We are always happy to assist our commercial users in tuning and evaluating Gurobi's performance.

Using distributed optimization can further improve Gurobi's performance



- MIPLIB 2010 (87 models)
 - Note: This test set was not designed for testing distributed optimization. Because of this, the results below understate the potential gains.
 - Models that take >1 second to solve

Machines	Distributed
4	1.43X
8	1.53X

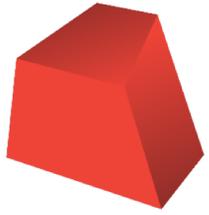
- Models that take >100 seconds to solve

Machines	Distributed
4	2.09X
8	2.87X

Models suited for distributed optimization can see significantly greater speed-ups

- Model: *seymour*
 - Hard set covering model from MIPLIB 2010
 - 4944 constraints, 1372 (binary) variables, 33K non-zeroes

Machines	Nodes	Time (s)	Speedup
1	476,642	9,267s	-
16	1,314,062	1,015s	9.1X
32	1,321,048	633s	14.6X



GUROBI
OPTIMIZATION

Isn't it time you considered upgrading to Gurobi?

1. You can get a free academic license at www.gurobi.com.
2. You can request a free commercial evaluation license by contacting us at: info@gurobi.com.
3. We are happy to assist with benchmarking your own models.