

**GUROBI**  
OPTIMIZATION

# **Gurobi 7.5 Performance Benchmarks**

# 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.*

\*Industry-standard public benchmarks maintained by [Hans Mittelmann](#) at Arizona State University

# 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

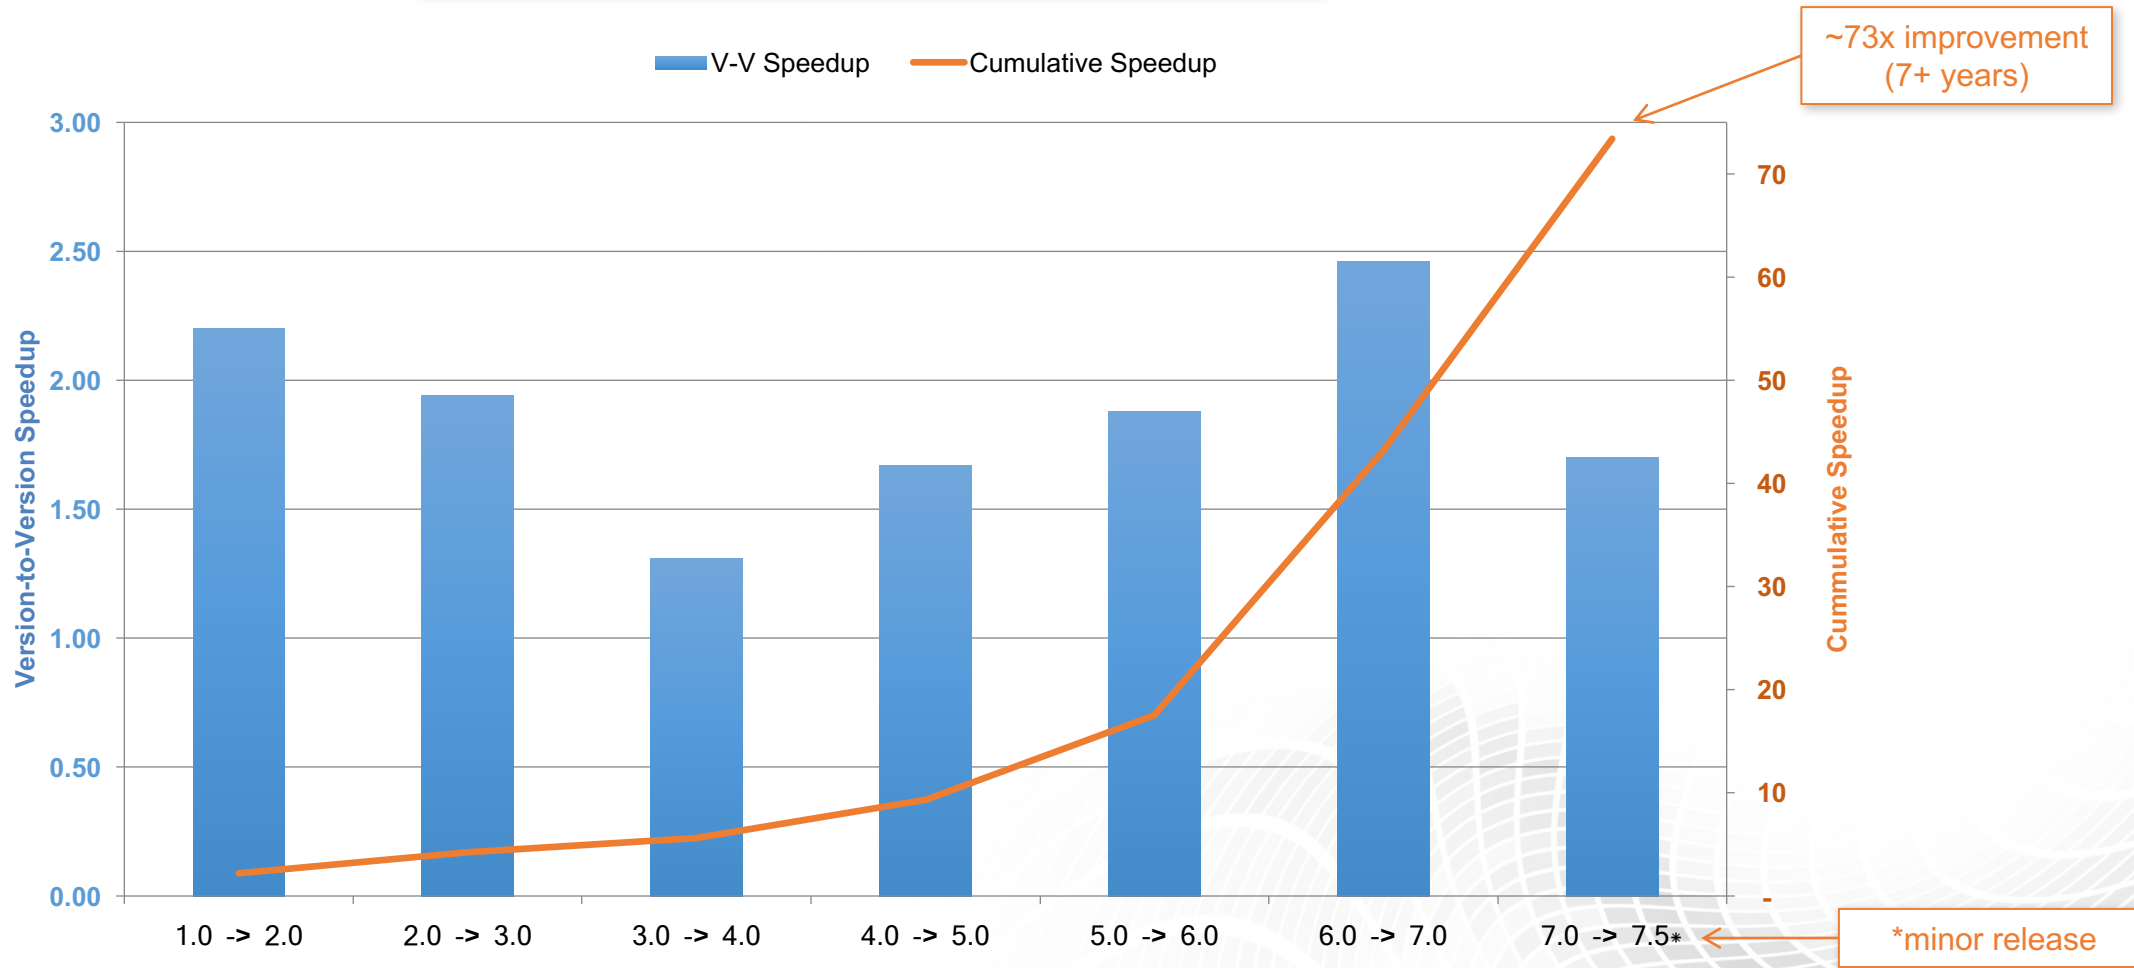
- Primary Objective
  - Competitive benchmarks against other solvers
- Test Bank
  - Maintained by Hans Mittelmann
    - <http://plato.la.asu.edu/bench.html>
  - Based upon MIPLIB 2010

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.

# Continual Performance Improvements



Nearly a 2x average improvement per major release



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


## Gurobi Version-to-Version Pairs

Test set has 3740 models:  
- 229 discarded due to inconsistent answers  
- 908 discarded that none of the versions can solve  
- speed-up measured on >100s bracket: 1227 models

# Continual Performance Improvements



Number of unsolved models




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

Test set has 4378 models:  
- 305 discarded due to inconsistent answers  
- 1081 discarded that none of the versions can solve  
- speed-up measured on >100s bracket: 1486 models

# Broad Performance Improvements in v7.5



- Consistent with prior releases, the Gurobi Optimizer v7.5 delivers performance improvements over v7.0 across a broad range of model types:
  - MIP – 32% faster (70% faster on models that take >100 seconds to solve)
  - LP
    - Concurrent – 15% faster (47% faster on models that take >100 seconds to solve)
    - Dual Simplex – 19% faster (42% faster on models that take >100s to solve)
    - Barrier – 9% faster (34% faster on models that take >100s to solve)
  - MIQP – 220% faster (too few models >100s to compare)
  - SOCP – 17% (too few models >100s to compare)



**GUROBI**  
OPTIMIZATION

# MILP Competitive Benchmarks

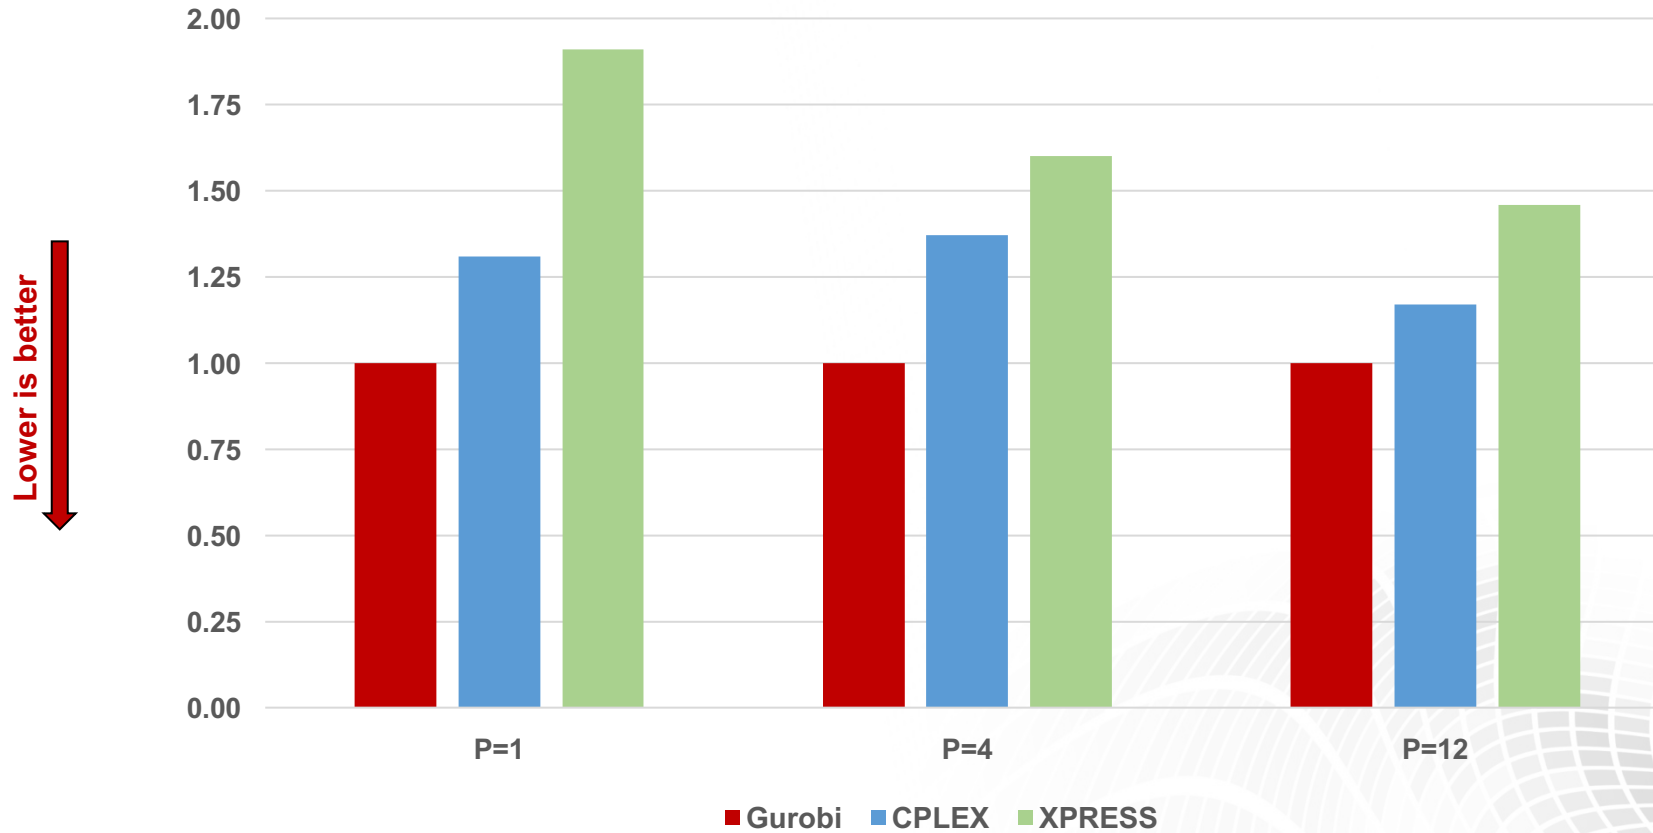
Gurobi 7.5 vs. CPLEX 12.7.1 vs. XPRESS 8.2.1  
(Maintained by Hans Mittelmann)

## ***Gurobi is...***

- *Fastest to optimality*
- *Fastest on the broadest test set*
- *Fastest to feasibility*
- *Fastest to detect infeasibility*

# Gurobi is Fastest Optimality


## Optimality Benchmark Comparison



For example, on this Mittelmann speed to optimality benchmark (87 models) using 12 threads (P=12), CPLEX was 17% slower (1.17) and XPRESS was 46% slower (1.46) than Gurobi.




# Gurobi is Fastest on the Broadest Test Set



# Gurobi is Fastest to Feasibility




Feasibility Benchmark Comparison (P=4)




# Gurobi is Fastest to Detect Infeasibility



Infeasibility Benchmark Comparison (P=8)






**GUROBI**  
OPTIMIZATION


# LP Competitive Benchmarks

Gurobi 7.5 vs. CPLEX 12.7.1 vs. XPRESS 8.2.1 vs. Mosek 8.0.0.74  
(Maintained by Hans Mittelmann)

*Gurobi has the fastest solve times*

# Gurobi is Fastest Across All LP Benchmarks






**GUROBI**  
OPTIMIZATION


# QP Competitive Benchmarks

Gurobi 7.5 vs. CPLEX 12.7.1 vs. XPRESS 8.2.1 vs. Mosek 8.0.0.64  
(Maintained by Hans Mittelmann)

*Gurobi has the fastest solve times*

# Gurobi is Fastest Across All QP Benchmarks



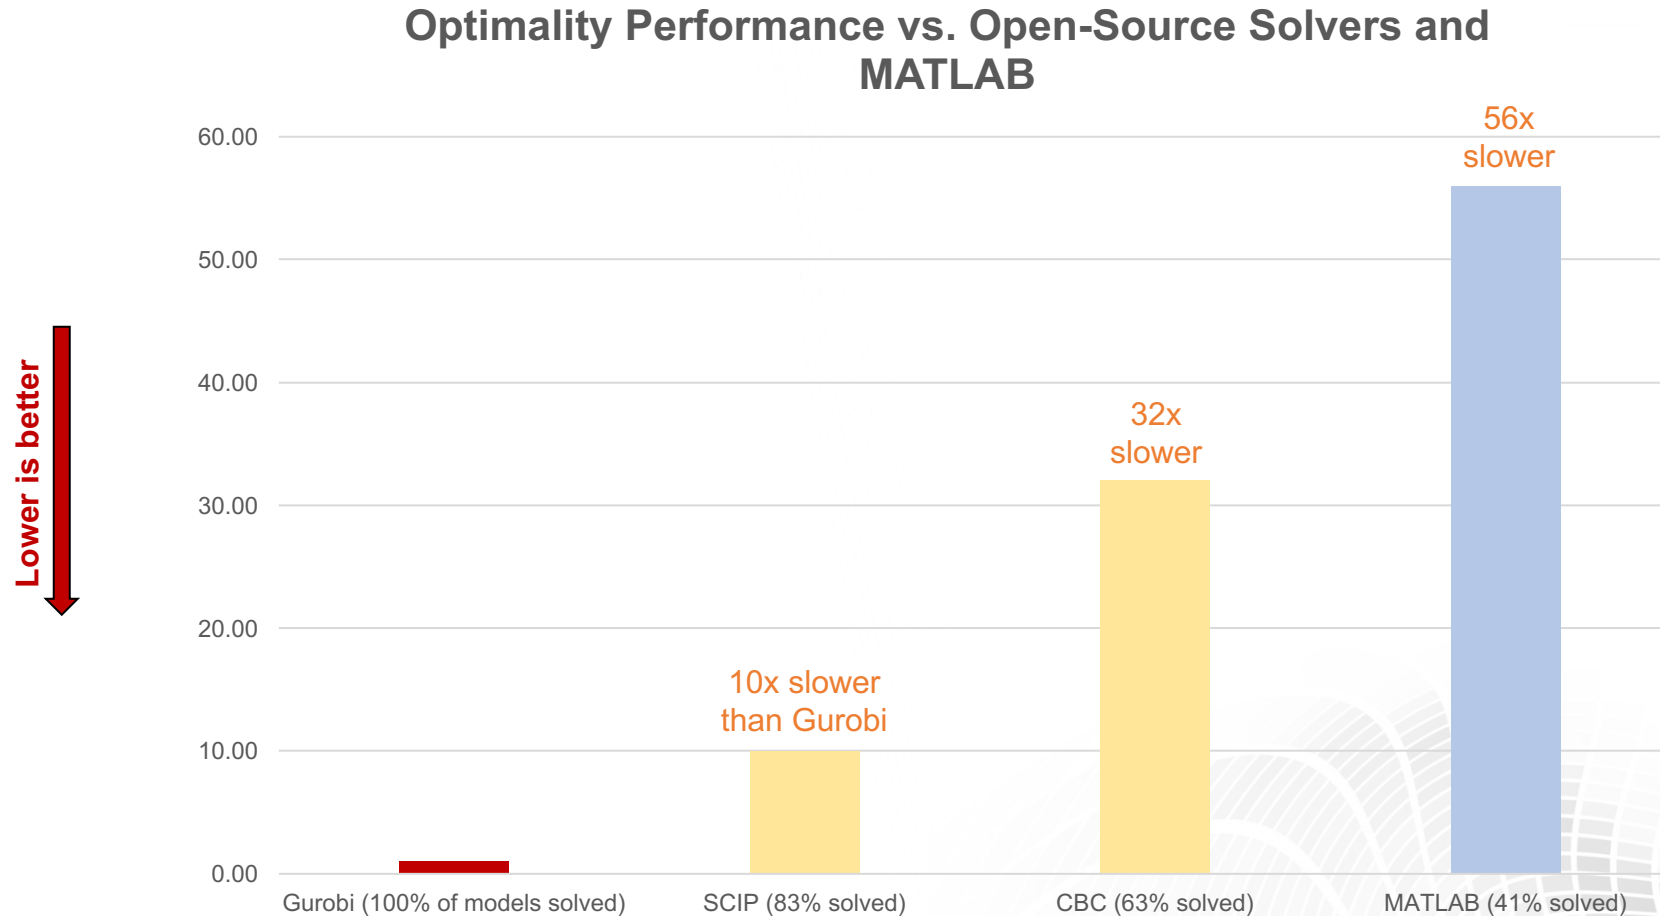


**GUROBI**  
OPTIMIZATION


# Open Source Solvers and MATLAB Benchmarks



# Benchmarks – Open Source MIP and MATLAB



- ▶ Gurobi, SCIP, CBC and MATLAB data from current MIPLIB test results (<http://plato.asu.edu/ftp/milpc.html>). LPSolve and GLPK are not included in benchmark results as they solved too few models (five and zero respectively) to calculate useful performance comparisons.
- ▶ Test set includes 87 models. 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
    - 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, obviously 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](http://www.gurobi.com).
2. You can request a free commercial evaluation license by contacting us at: [info@gurobi.com](mailto:info@gurobi.com).
3. We are happy to assist with benchmarking your own models