CPLEX Optimization Studio 12.10 Performance improvements

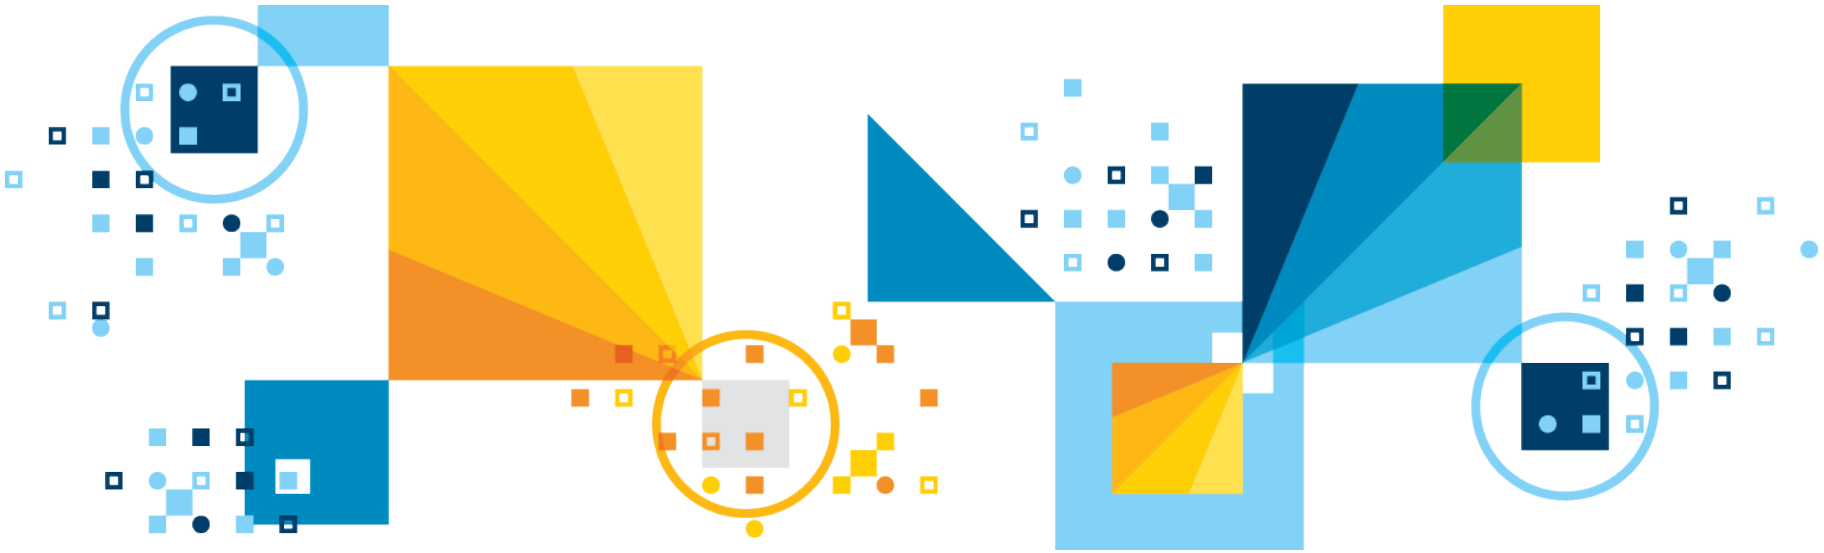

CPO Performance

CPO – Scheduling Performance

- Iterative Diving
  - Fast dives using limited propagation at nodes
  - Used in defaults with >3 threads
  - 1.5 x speedup over testset of 2690 problem instances

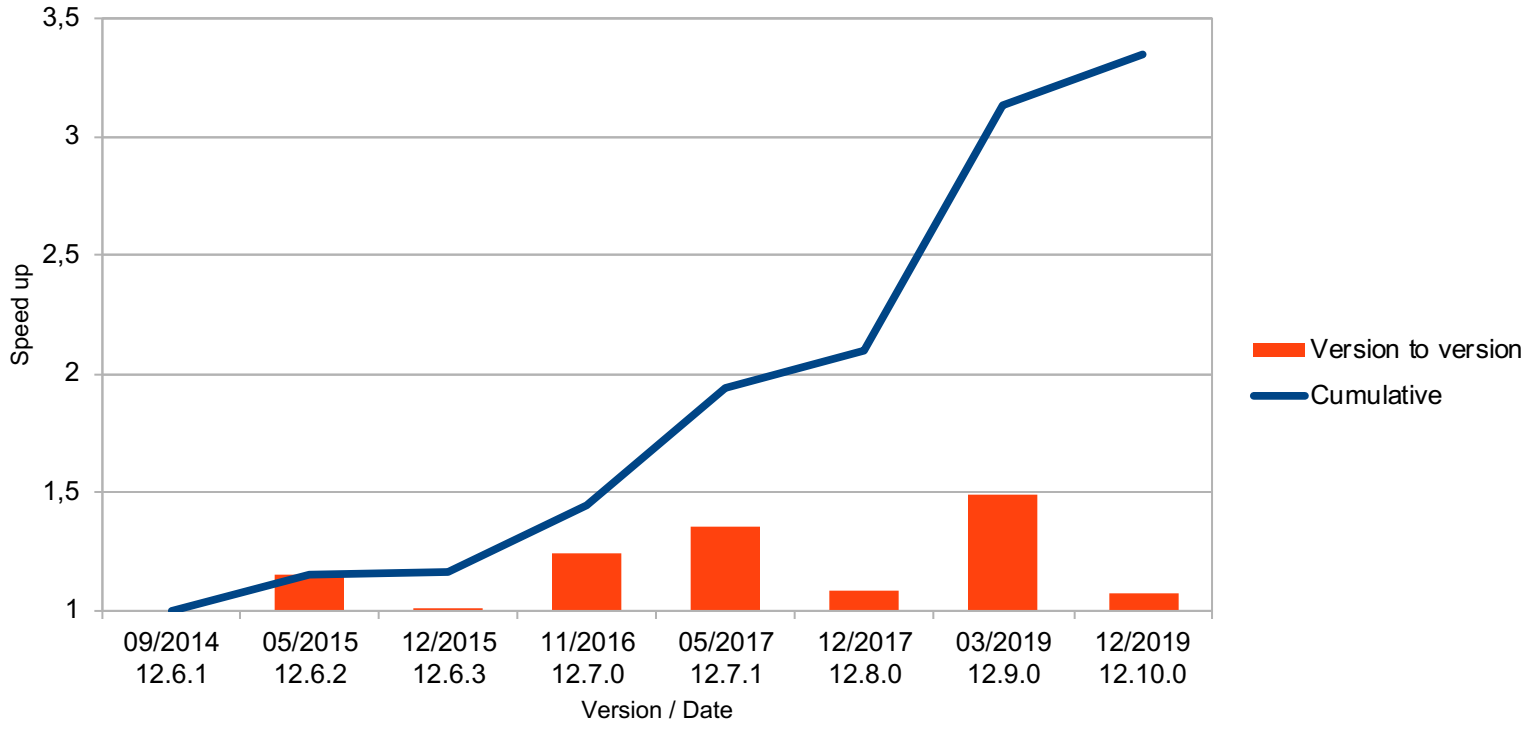

CPO – Optimizer Performance

- Strategy selection via reinforced learning
  - CPO comprises a portfolio of search and propagation strategies
  - Automatically and dynamically pick the “best” one to use on a particular problem instance via reinforced learning
  - 1.5x overall performance improvement
  - More for feasibility problems (no objective):

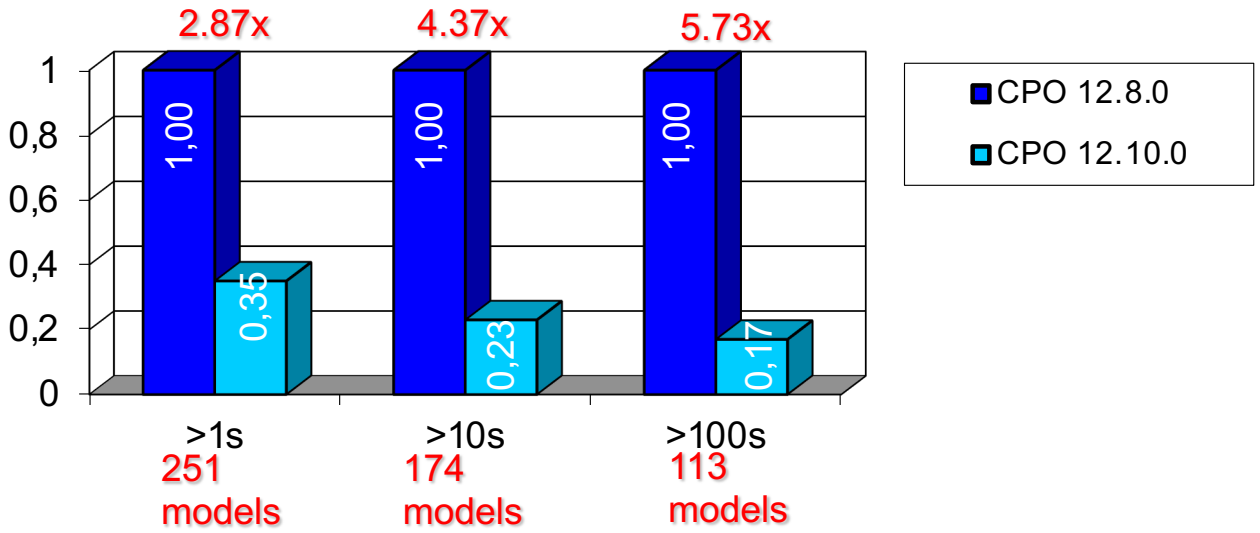

Feasibility Problems

LP Performance Improvements

CPLEX 12.9.0 vs.12.10.0: LP performance improvement

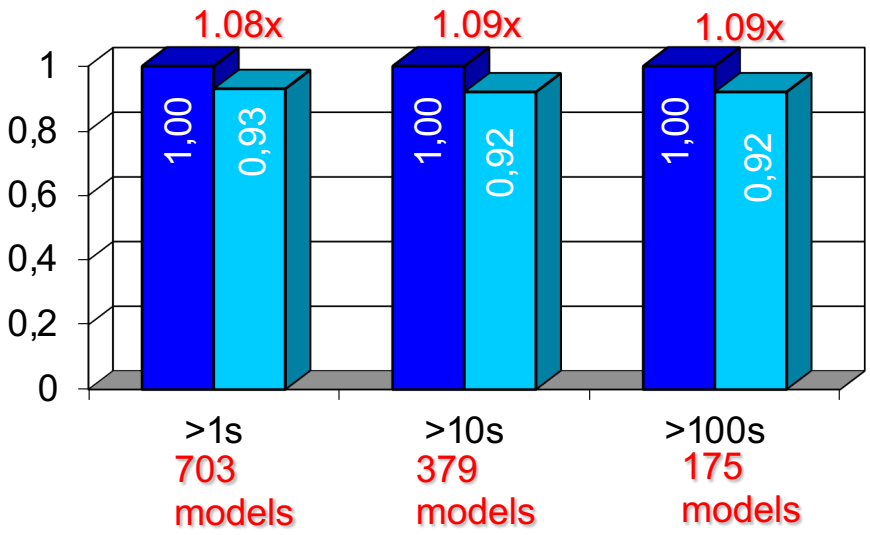

Deterministic concurrent LP (12 threads)

- Improved code generation
  - Switched to icc 19.0.4.243
  - Compiler pragmas to better exploit SIMD instructions
    Helped by some code reorganization
- Algorithmic improvements
  - Better handling of superbasics in crossover
  - Better handling of presolve in concurrent optimizer

Date: January 2, 2020
Testset: LP: 2203 models
Machine: Intel(R) Xeon(R) CPU E5-2667 v4 @ 3.20GHz, 64 GB RAM, 12 threads, deterministic
Timelimit: 10,000 sec

MILP Performance Improvements

CPLEX MILP performance evolution

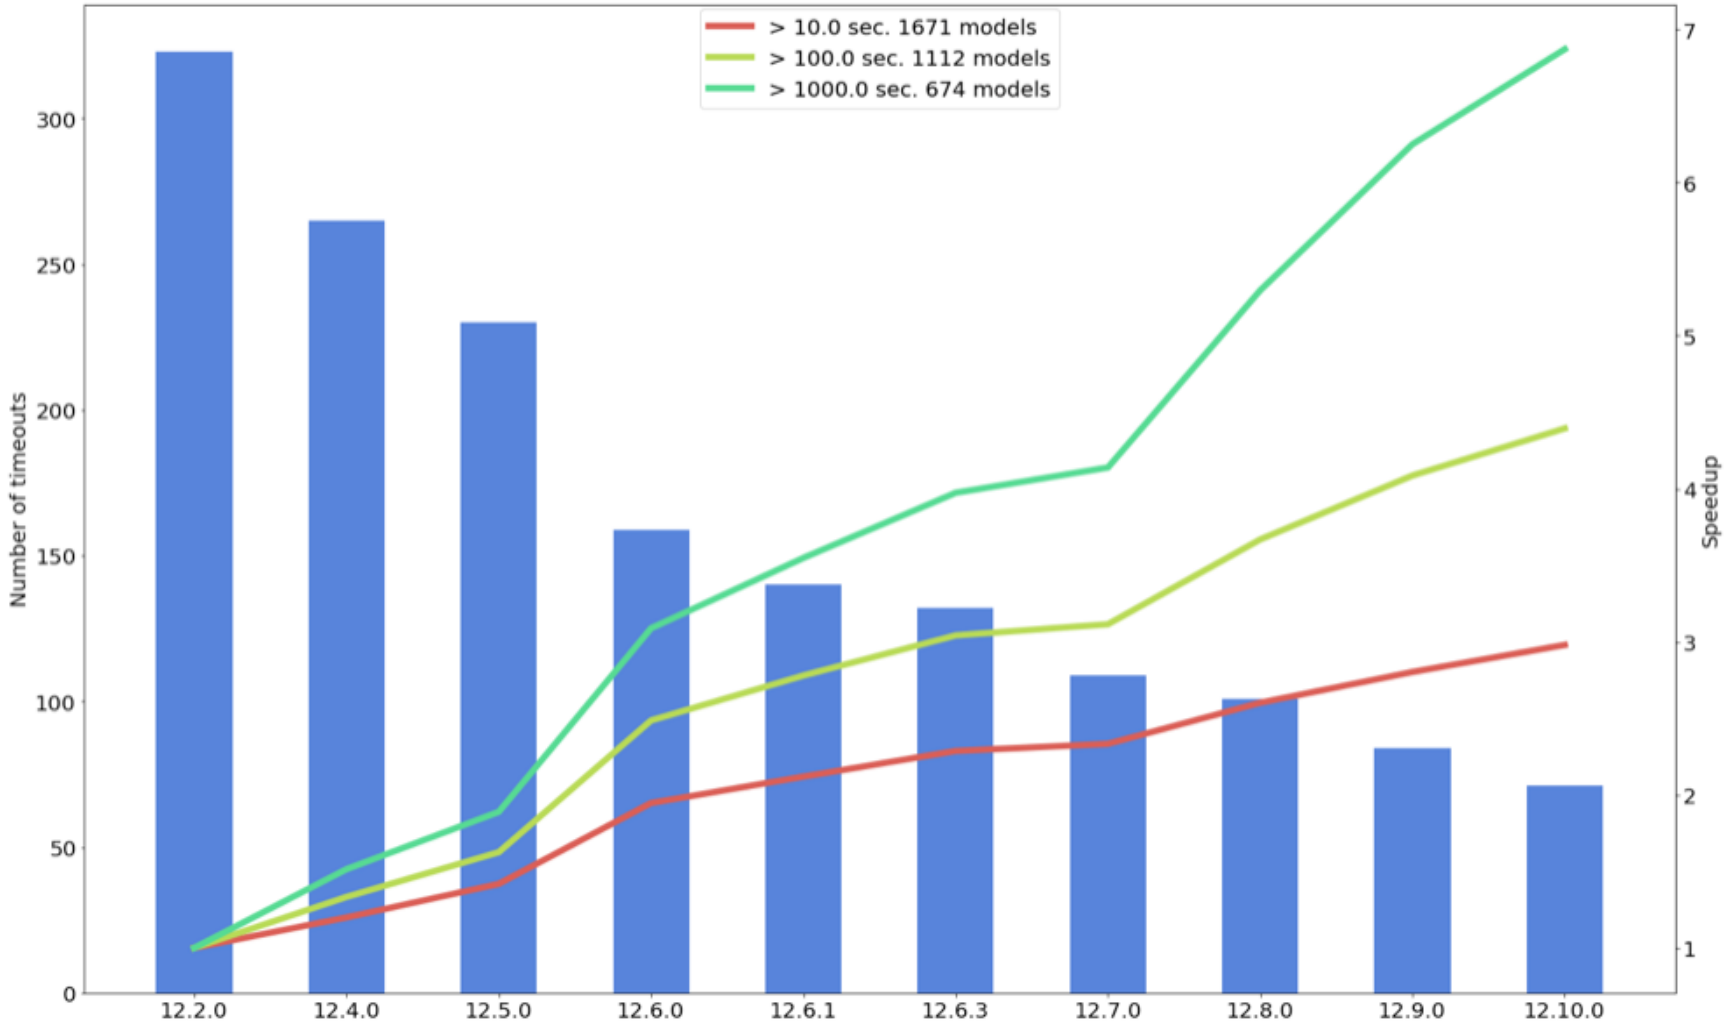

Date: January 2, 2020
Testset: MILP: 5118 models
Machine: Intel(R) Xeon(R) CPU E5-2667 v4 @ 3.20GHz, 64 GB RAM, 12 threads, deterministic
Timelimit: 10,000 sec

CPLEX 12.9.0 vs.12.10.0: MILP performance improvement

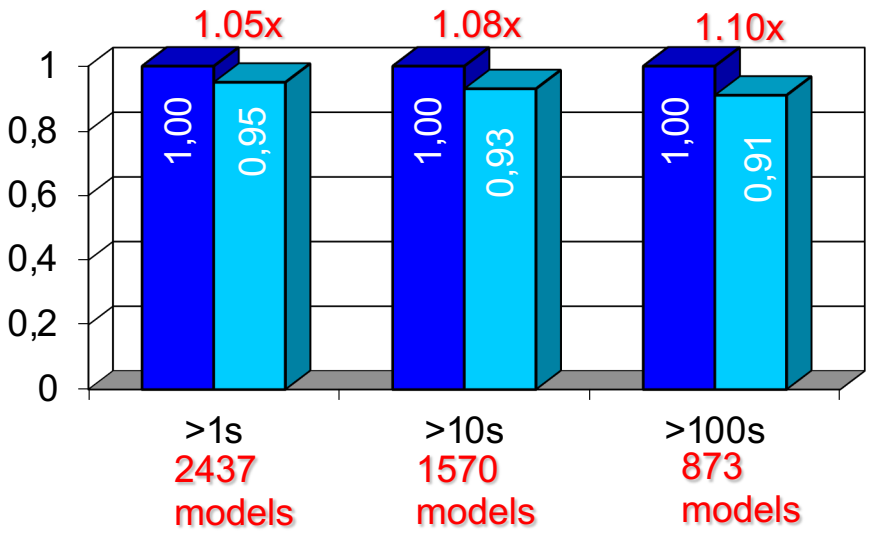

CPLEX 12.9.0

CPLEX 12.10.0

Time limits:
57 / 37

Deterministic parallel MILP (12 threads)

- MILP improvements summary
  - Dynamic search: Improved restarts: 12% on 7% affected models
  - Improved handling of indicator variables: 20% on 0.01% affected models
  - Heuristics: 1-2% overall
  - Cuts:
    - improvements of zero-half cut separator for models with relatively dense binary constraints: 10% (on 5% affected models)
    - Farkas Cuts from infeasible node LPs: 2% overall

Date: January 2, 2020
Testset: MILP: 5118 models
Machine: Intel(R) Xeon(R) CPU E5-2667 v4 @ 3.20GHz, 64 GB RAM, 12 threads, deterministic
Timelimit: 10,000 sec

MIQP Performance Improvements

CPLEX 12.9.0 vs.12.10.0: MIQP performance improvement

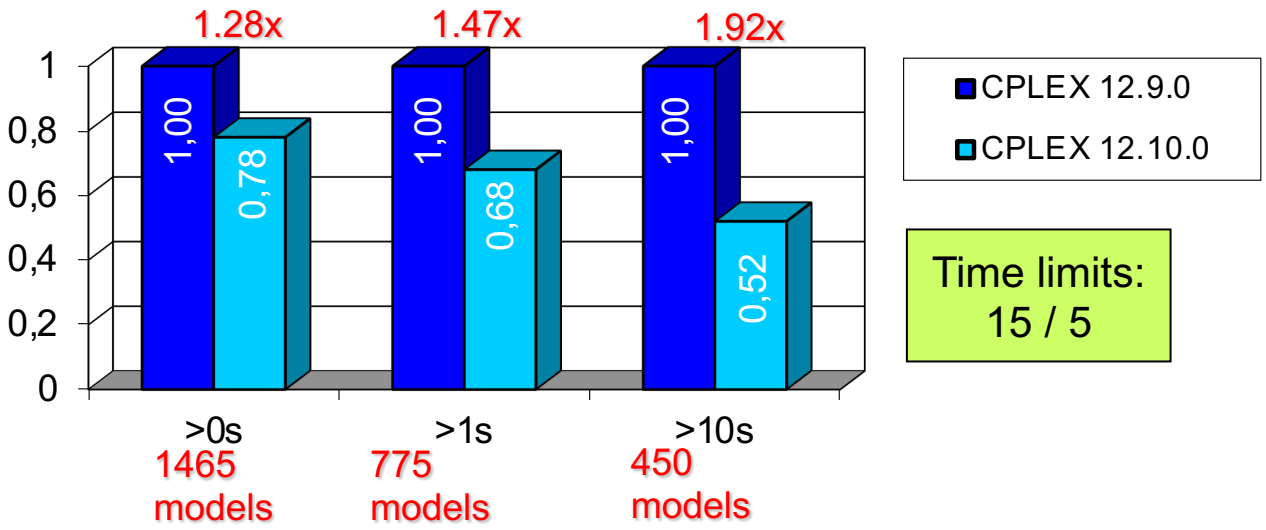

Deterministic MIQP (12 threads)

- Automatic decision whether Q should be linearized or kept as is
  - Use Machine Learning to learn a classifier [Bonami, Lodi, Zapperlon, 2018]
- Problem instances from training set excluded

Date: January 2, 2020
Testset: MIQP: 1465 models
Machine: Intel(R) Xeon(R) CPU E5-2667 v4 @ 3.20GHz, 64 GB RAM, 12 threads, deterministic
Timelimit: 10,000 sec

Benders Performance Improvements

CPLEX 12.9.0 vs.12.10.0: Benders B&C improvements

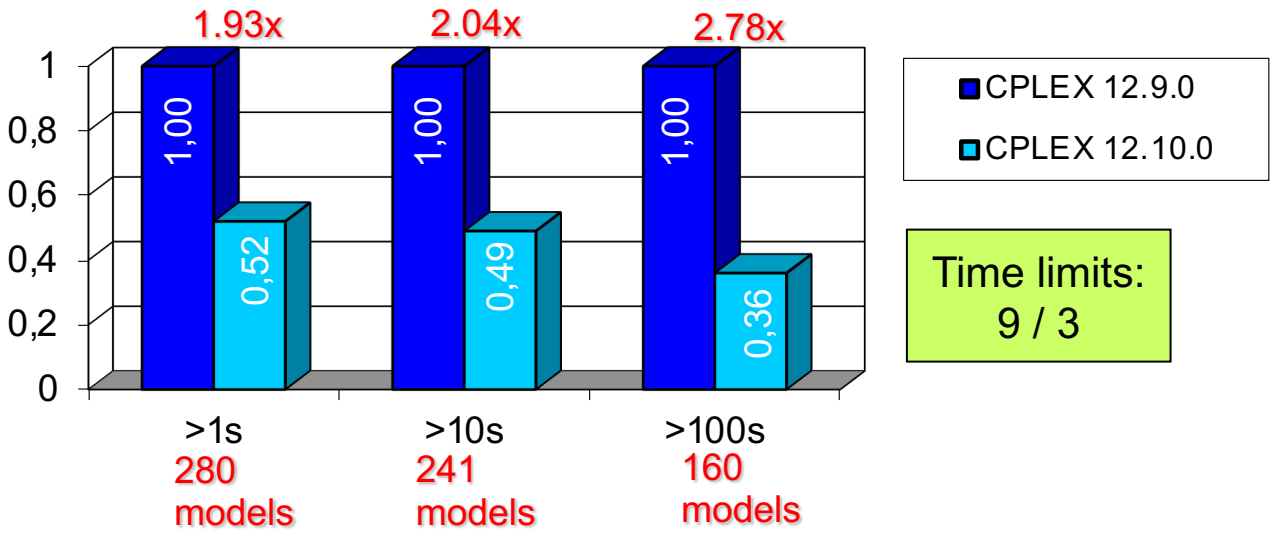

Benders MILP branch-and-cut (12 threads)

- Benders improvements summary
  - Benders specific presolve: 2%
  - Improved purging of Benders cuts: 2-3%
  - Exploitation of General Bound Constraints: 5%
    - 66x speedup on Partial Set Covering Location Problems

Date: January 2, 2020
Testset: Benders MILP: 374 models
Machine: Intel(R) Xeon(R) CPU E5-2667 v4 @ 3.20GHz, 64 GB RAM, 12 threads, deterministic
Timelimit: 10,000 sec

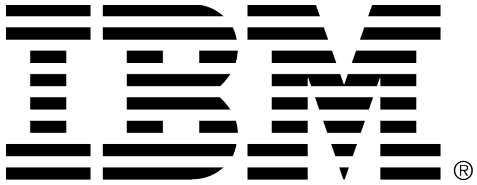

Legal Disclaimer

-
- The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.
- References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.
- Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Other company, product, or service names may be trademarks or service marks of others.