Open Energy Benchmark

Daniele Lerede, Enrico Antonini, Jacek Bendig, Kristijan Faust, Fabrizio Finozzi, Goli Vamsi Priya, Siddarth Krishna, Madhukar Mishra, Maximillian Parzen

Open Energy Transition

HiGHS Workshop 2026

01.06.2026

Outline

- Why benchmarking solvers for energy planning?
- The Open Energy Benchmark
- The State of Solvers for Energy Planning
- Assessing HiPO performances
- Future outlook

Why benchmarking solvers for energy planning?

- **Energy planning models are becoming much larger and more complex**
  - Integrate emerging technologies, climate impacts, and socio-economic scenarios
  - Formulated as large-scale LP/MILP problems with millions of variables and constraints
- **Solver performance is a critical bottleneck**
  - Slow solvers limit spatial/temporal resolution and reduce the number of scenarios that can be explored
- **Computational limits reduce planning robustness**
  - Fewer scenarios mean less ability to assess uncertainty in technology costs, policy, and climate change
- **Open-source solvers are improving, but their capabilities remain uncertain**
  - Transparent, reproducible benchmarks are needed to identify strengths, weaknesses, and priorities for development

The Open Energy Benchmark

An **open, transparent, diverse, interactive, representative, automated** benchmark of optimization solvers on representative problems from the energy planning domain.

- **5 solvers** (15 versions): HiGHS, CBC, GLPK, SCIP & Gurobi
- **13 modelling frameworks** (e.g., GenX, TIMES, PyPSA)
- **213 problems** (various scenarios and spatio-temporal resolutions)
- Problem sizes from **600** up to **52M** variables
- **17 contributors**, **14 forks**, and **66 stars** on GitHub

The State of Solvers for Energy Planning
Main assumptions

- All solvers are run with **default options** and solves are started from **one fixed random seed**
- Problems are run on **virtual machines** to increase parallelization, cost efficiency and transparency

  Hardware configuration depends on the benchmark size:

  - **Small and Medium benchmarks** (<1e6 variables): run with a timeout of 1 hour on a machine with 1 core (2 vCPU) and 7 GB RAM
  - **Large benchmark instances:** run with a timeout of 24 hours on a machine with 8 cores (16 vCPU) and 124 GB RAM
- Runtime is obtained as the runtime of the call to `linopy.Model.solve()`

The State of Solvers for Energy Planning

The set of analyzed problems

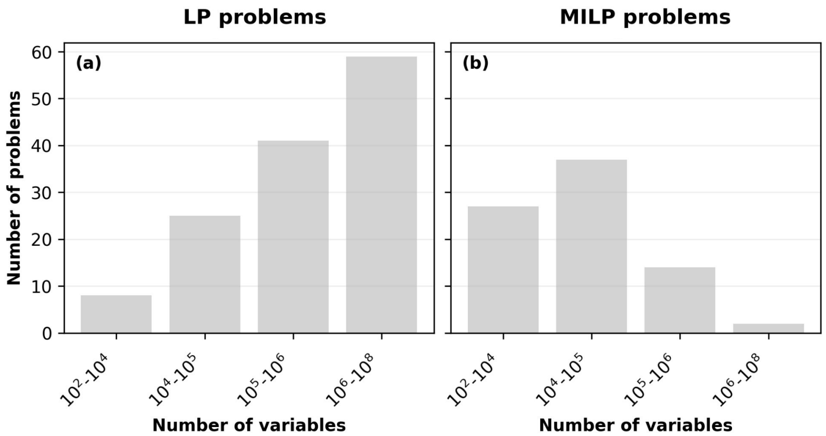

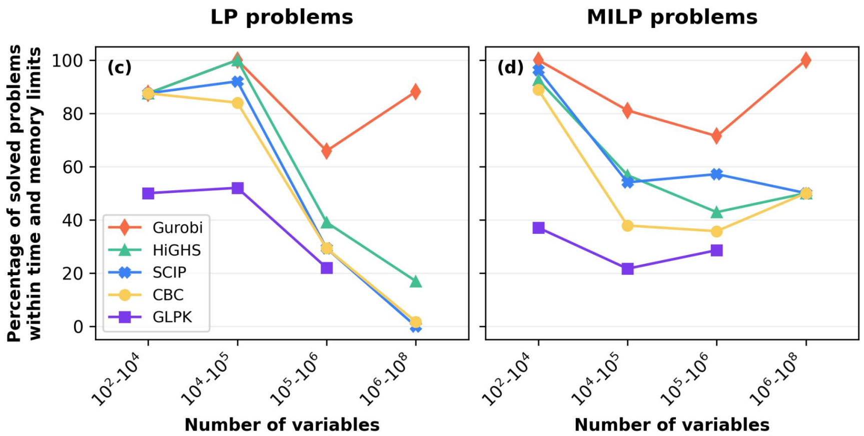

- The set of problems spans over a **wide size range**, though **lacking in large MILP benchmark instances**
- **Open-source solvers** still **fail to compete with Gurobi** in solving **large problem instances**

The State of Solvers for Energy Planning

Runtime analysis

- **HiGHS (default)** is consistently **the fastest among OS solvers** and most competitive option for LPs, but struggles above 10^6 variables
- **MILP problems** remain significantly **more challenging for all open-source alternatives**

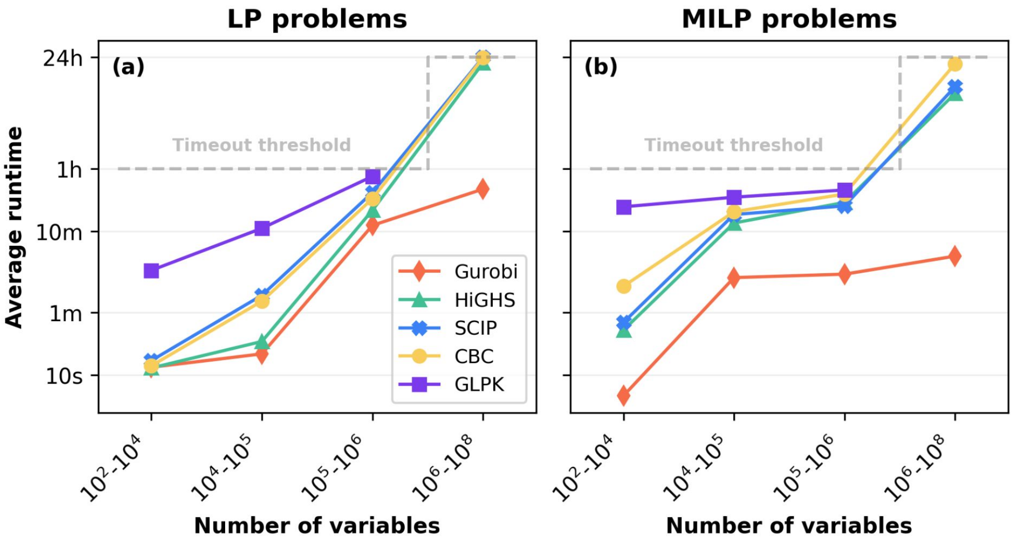

The State of Solvers for Energy Planning
Slowdown analysis

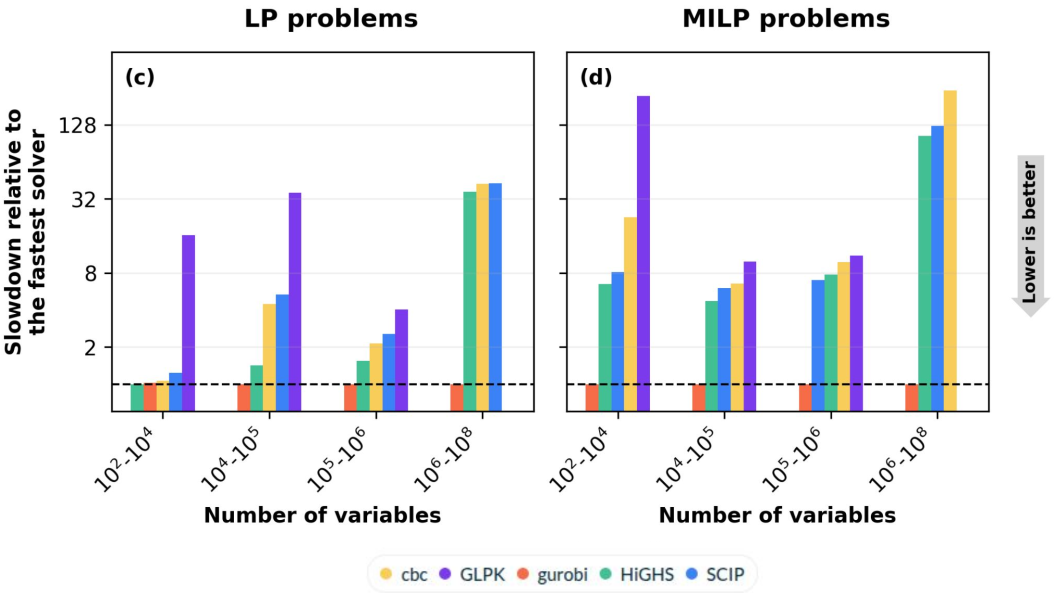

LP problems:

- Performance of OS solvers diverges rapidly beyond 10^6 vars
- **HiGHS is ~1.5× slower than Gurobi** for S/M problems, while other OS solvers exhibit worse scaling

MILP problems:

- **Performance gap** already **significant at small scales**

The State of Solvers for Energy Planning

Solver performances over time

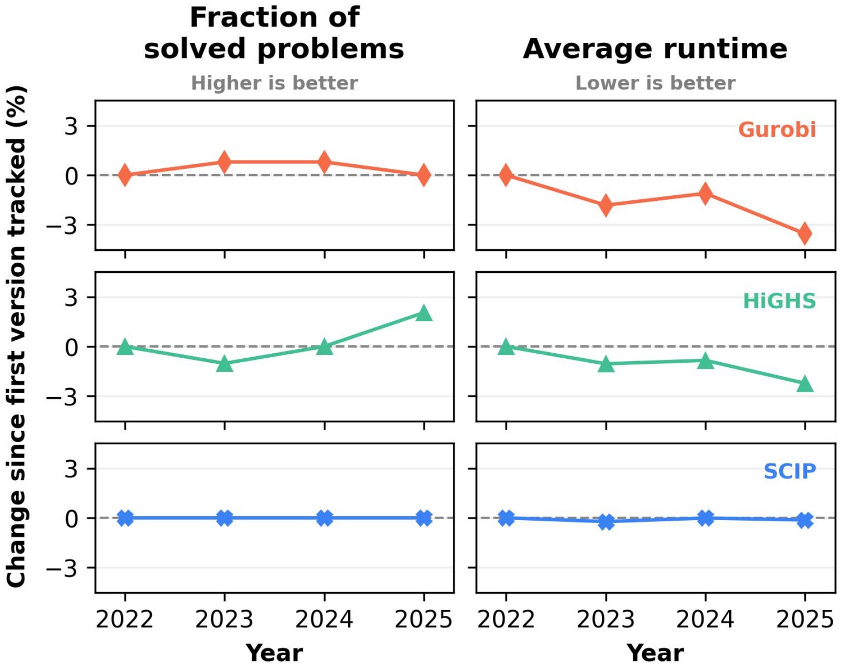

- SCIP has maintained stable performance metrics over the past few years
- **HiGHS showed improvements** in both average runtime and the fraction of successfully solved problems, **matching the relative pace of Gurobi's speed improvements**

Assessing HiPO performances
Success rates over the full set of problems

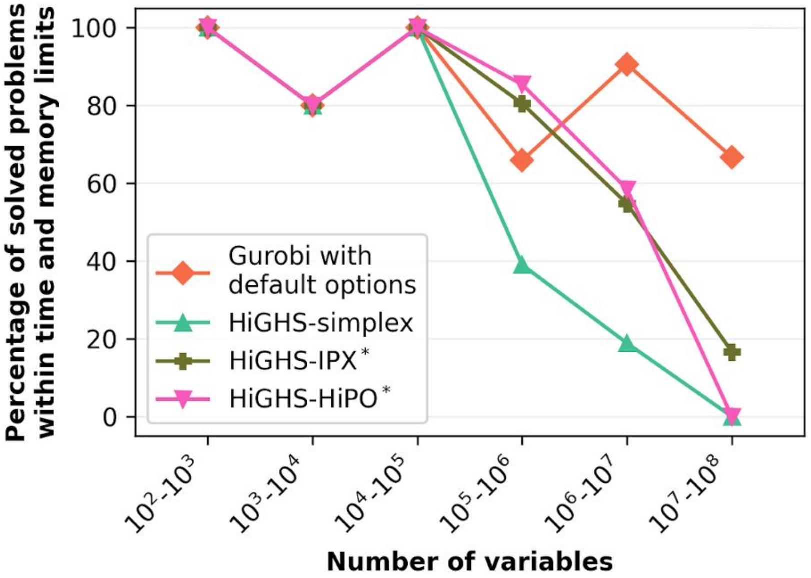

* Solvers configured to return feasible solutions, not necessarily basic

Different HiGHS solver options - including **HiPO** - were **assessed against Gurobi***:

- **$10^5 \div 10^6$ variables**: **85% success rate** for HiPO **(best among HiGHS options)**
- **$1 \div 10$ million variables: 59% success rate** for HiPO **(best among HiGHS options)**
- **> 10 million variables: HiPO cannot solve any problem** (same as simplex)

* Gurobi default runs different algorithms in parallel and always returns basic solutions, whereas HiGHS IPM implementations skip crossover and return interior feasible solutions

Assessing HiPO performances

Relative improvements in runtime

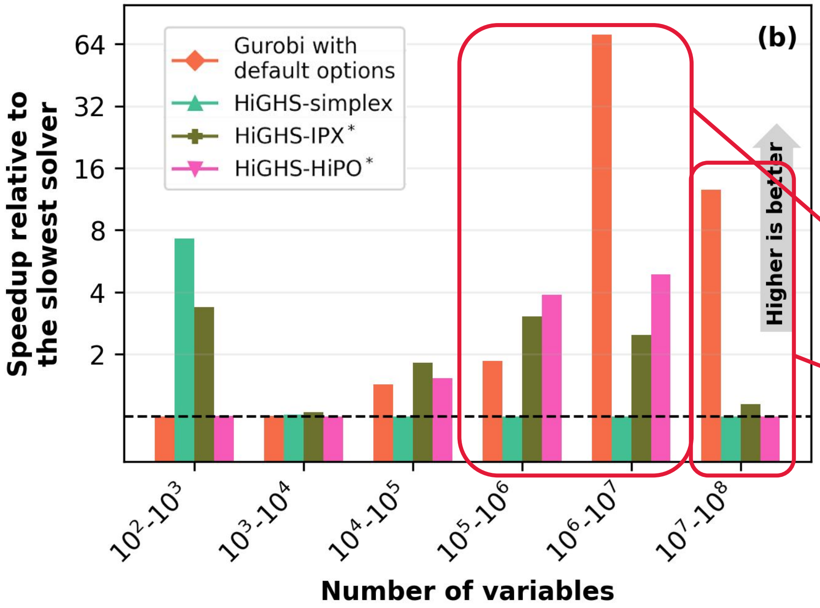

- **Simplex** and **IPX both** **perform better than HiPO only on XS problems**
- **HiPO** guarantees **better runtime performances as problem size scales up**

HiPO in average 5× faster* than simplex

Gurobi still unmatched on XL problems

*runtimes calculated using a shifted geometric mean (unsolved problems are penalized with a factor of 2)

Assessing HiPO performances

PyPSA-DE spatial resolution-scaling study

Assessed problem: PyPSA-DE power model (overnight capacity expansion to 2050, no transmission expansion, LP formulation, 1h resolution)

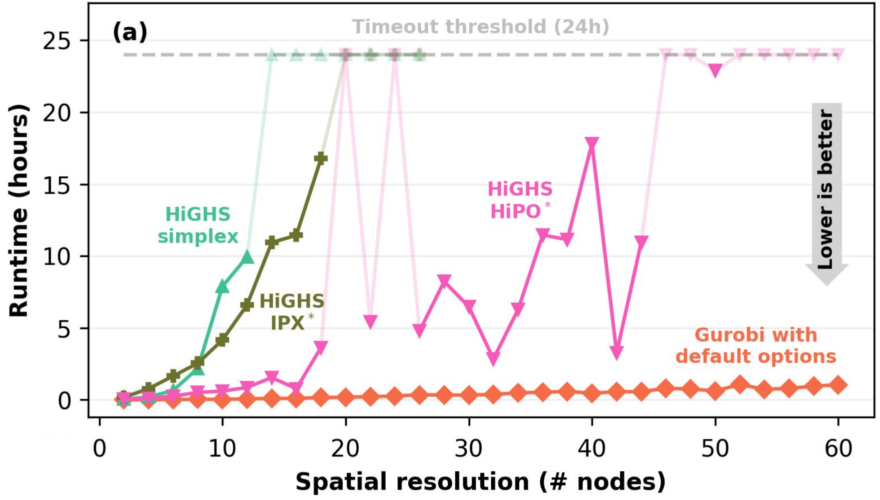

- **HiGHS-simplex** hits a **computational wall above 12 nodes**
- **HiGHS-IPX** performs slightly better before **timing out above 18 nodes**
- **HiGHS-HiPO** demonstrates **remarkable resilience* up to 50 nodes (9 M vars)**, but runtimes cannot compare to Gurobi

* Poor scaling appears to explain failures in some intermediate-size instances that deviate from the overall trend.

Assessing HiPO performances
Conclusions

- **HiPO** significantly **pushes the boundary of the scale of solvable problems** with open-source tools
- Modelers can now solve a **1-hour resolution, national-scale electricity model with 50 nodes** in less than 1 day using an OS solver (more than double the resolution of the most refined model solved using other open-source methods)
- **Across all problems in our benchmark suite**, **HiPO increases the number of problems that can be solved** by open-source solvers **by 7%**, completing complex models with millions of variables up to 15 times faster than the best alternative

Future outlook

- **New benchmark campaign (2026)**

  Planned for the **end of 2026** with additional LP and MILP problems and **updated solver versions** (potentially considering **other solver configurations** and **multiple random seeds**)

- **Ongoing open call for benchmark problems**

  Continued community-driven **expansion of the benchmark suite** through contributions of realistic and challenging optimization problems from different modelling frameworks and applications

- **Shared governance and sustainability**

  Discussion toward a governance model involving universities, non-profit organizations, solver and model developers to ensure neutrality, openness, reproducibility, and long-term sustainability

Online resources

Open Energy Benchmark website

https://openenergybenchmark.org

Open call for benchmarks

https://openenergybenchmark.org/blog/open_call#why-contribute