

The 2026 Land-Doig MIP Competition

Christian Tjandraatmadja

on behalf of the 2026 Land-Doig MIP Competition Committee:

Beste Basciftci, Akif Çördük, Gerald Gamrath, Christopher Hojny,
Jan Kronqvist, Haihao Lu, Christian Tjandraatmadja

2026 Land-Doig MIP Committee



Beste Basciftci
University of Iowa



Akif Çördük
NVIDIA



Gerald Gamrath
COPT



Christopher Hojny
*Eindhoven University
of Technology*



Jan Kronqvist
*KTH Royal Institute
of Technology*



Haihao Lu
*Massachusetts Institute
of Technology*




Christian Tjandraatmadja
(chair)
Google

Thank you for the support!




Provided NVIDIA Brev GPU credits for all registrants who requested
and a total prize pool of ~~\$1500~~ in NVIDIA Brev GPU credits
\$2500



Travel support for winning speaker
and the host of the competition

This year's theme was

GPU-Accelerated Primal Heuristics for MIP



5-minute general-purpose MILP heuristics that leverage GPU, without the use of fully-fledged discrete optimization solvers (LP allowed)

Competition Goals

1

Encourage
timely research in a
developing area of MIP




2

Make the competition
and research in this
area more **accessible**,
especially to students

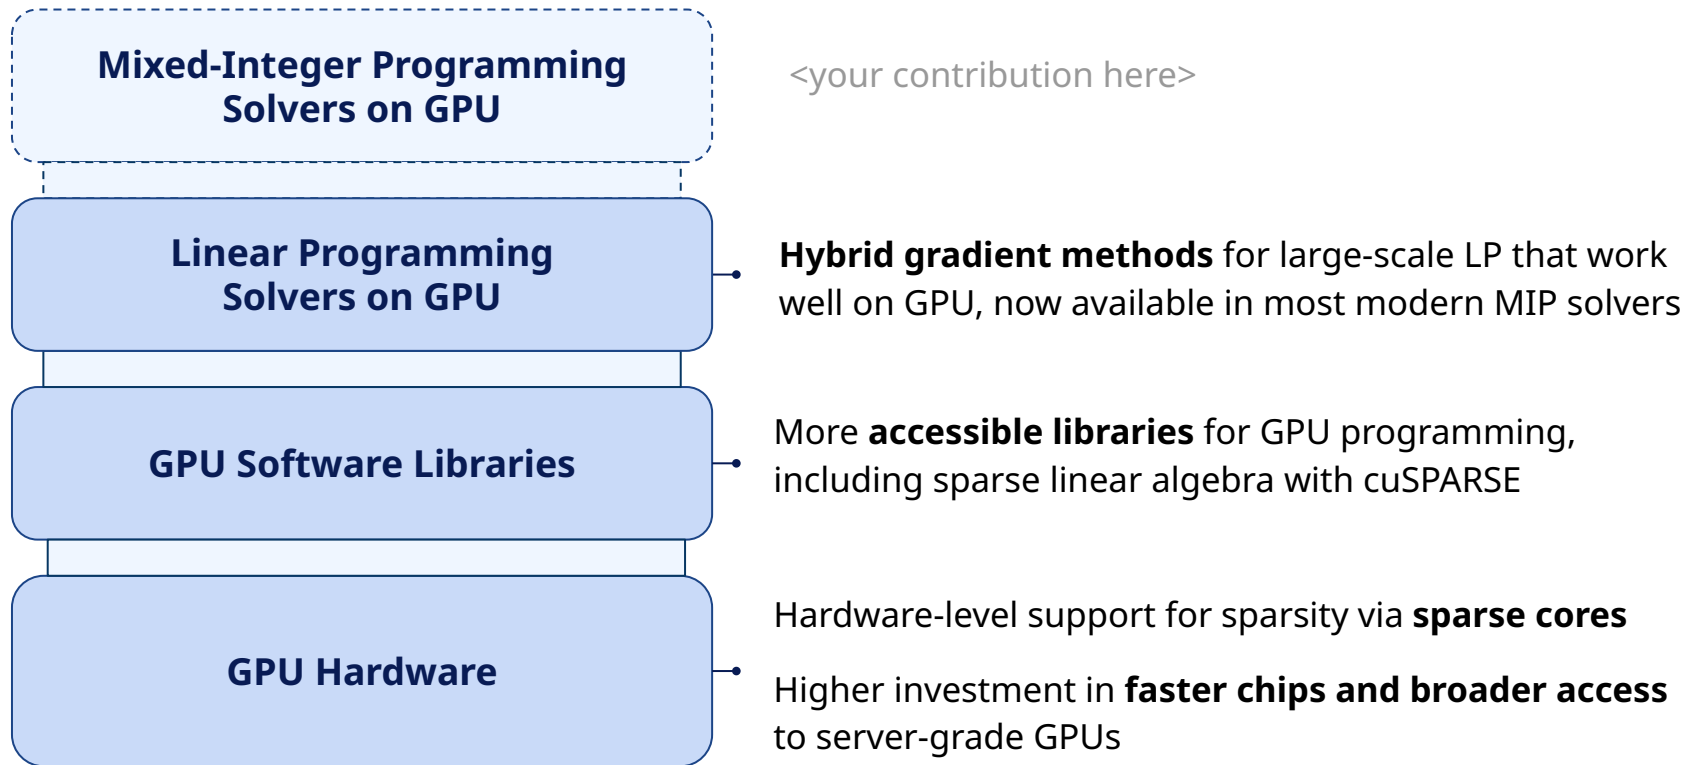


3

Foster methodology
with promising paths
for eventual **adoption**
by MIP solvers



Timely topic: Why is this the right moment for MIP + GPUs?



These foundations also hold for other hardware accelerators (e.g. Google's TPUs), but we are focusing on GPUs in this competition

Accessibility: Learning and compute resources

Committee-contributed tutorials

Feasibility jump on GPU with Python

written by Christian

2-opt on GPU in CUDA with C++

written by Akif

Compute resources

NVIDIA generously provided cloud GPU credits to **all registered teams** who requested them for developing competition code

Learning Material

Tutorial

We provide a customized tutorial where we walk through implementing a local search heuristic on GPU, in particular feasibility jump. This tutorial is in Python and Colab for convenience.

- [2026 MIP Competition Tutorial: Local Search on GPU - Part 1: Introduction + a CPU implementation of feasibility jump](#)
- [2026 MIP Competition Tutorial: Local Search on GPU - Part 2: A GPU implementation of feasibility jump via Numba](#)
- [2026 MIP Competition Tutorial: Local Search on GPU - Part 3: A GPU implementation of feasibility jump via JAX](#)

We also provide a shorter CUDA C++ tutorial with an example of 2-opt.

- [2026 MIP Competition Tutorial: GPU-Accelerated 2-Opt Local Search for MIP in CUDA](#)

External sources

- [This repository](#) is a collection of tutorials in C++ and Python to learn how to use GPUs and a good starting point.
- Material on CUDA:
 - [Talk on getting started with CUDA](#)
 - [Course on CUDA with Python](#)
 - [Course on CUDA with C++](#)
- Profiling/debugging:
 - [Tutorial on profiling and debugging](#)
 - [Webinar on profiling and debugging](#)
 - [Documentation on profiling](#)
- Libraries such as [JAX](#) and [PyTorch](#) are also accessible and efficient ways to perform computation on GPUs in Python.

Opportunity for future adoption: Closing the loop



Decision Optimization

NVIDIA cuOpt

Achieve world-record speed on large-scale problems with millions of constraints and variables—saving time and reducing costs.



GPU Acceleration for Unprecedented Scale

The first commercial optimization solver globally to support GPU acceleration for both first-order method (PDLF) and barrier method.

Maximizes GPU utilization, leveraging the massively parallel architecture of GPUs to deliver breakthrough performance on solving large-scale problems and solves previously unsolvable huge-scale problems arising from academic research and industry.



GPU Acceleration of the Hybrid Gradient Algorithm in FICO Xpress

FICO® Xpress 9.8 comes with a GPU accelerated implementation of the hybrid gradient algorithm, yielding speedups of up to 50x



Gurobi®

Why Gurobi

NEWSROOM

Gurobi Releases Version 13.0 with Improved Performance and New Solving Capabilities

With Gurobi 13.0, users can expect faster solves, new nonlinear capabilities, PDHG with GPU acceleration, and Kubernetes autoscaling.

Competition details

11 submissions (41 registrants)

Allowed resources:

- GPU: H100 (L40 was recommended for development)
- CPU: 8 cores and 64GB RAM (max. 8 threads)
- Time limit of 5 minutes
- No “fully-fledged” discrete optimization solvers; continuous solvers are ok

Submissions were evaluated for:

- Computational performance (primal integral and primal gap)
- Methodological and engineering innovations

Instances

A **diverse** benchmark of MILP instances that includes **large-scale instances**

- 50 instances in public test set
 - 18 problem classes
- 60 instances in hidden set
 - 23 problem classes: 18 from test set + 5 additional
 - Hidden instances are slightly larger than test set and receive higher weight in evaluation

Instance sources: submissions to next MIPLIB + MIPLIB 2017 + committee-contributed

All instances are **already presolved** (via default Gurobi presolve)

Instance classes

Public test set (also in Hidden set)

01-02: Vehicle routing with crew scheduling (bab*)

03-05: Airplane boarding

06-08: Bus crew scheduling

09-10: Linear classification

11-13: Wagon ordering for trains

14-15: Telecommunication network configuration

16-18: Support vector machine classification

19-21: Length-constrained cycle partition

22-24: Maintenance scheduling

25-27: N-Queens in higher dimensions

28-30: Rectangle packing

31-33: Open pit mining (rmine*)

34-36: Sports scheduling

37-38: Sum graphs

39-41: Time-dependent TSP

42-44: Security-constrained unit commitment

45-47: Binarized neural network verification

48-50: Railway network design

Hidden set only

Physician scheduling

Optimal classification tree

Train routing and scheduling

Error correcting

Capacitated vertex separator

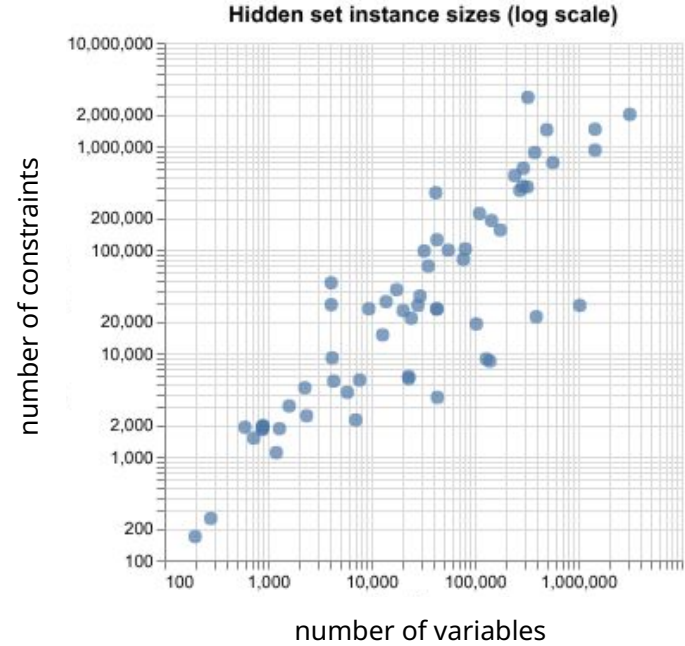
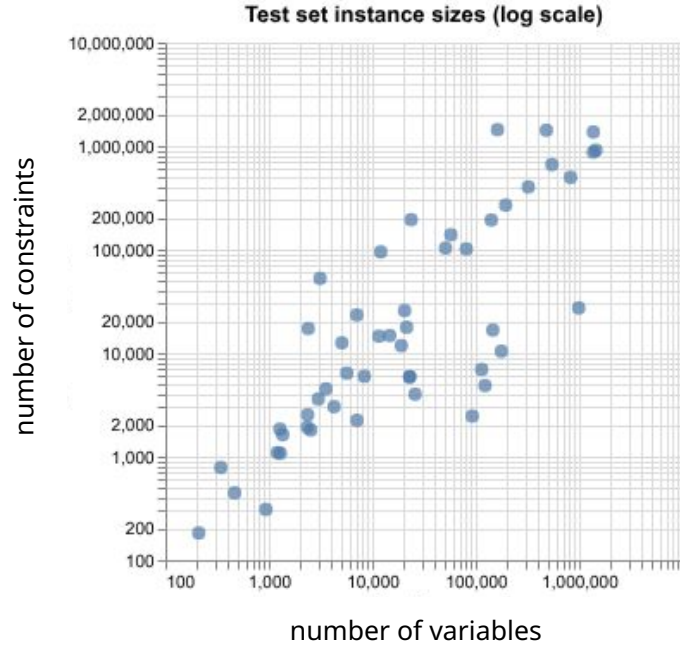
Color legend

Next MIPLIB

MIPLIB 2017

Committee-contributed

Instance sizes



Review process

Due to time and resource limitations, we performed a two-round review process:

Round 1:

- All submissions were run once and each report was reviewed by 3 committee members
- Based on both results, we jointly selected a subset of submissions for the second round

Round 2:

- Submissions were rerun two more times for hidden set to reduce impact of compute variability
- Round 2 submissions were ranked by committee members independently in terms of both computational results and innovations
- Rankings were aggregated via Borda count, deciding winner and honorable mentions

Technical details

- We use the **MIPLIB solution checker** to check the top 160 solutions in each run
 - Constraint tolerance of $1e-6$ and integer feasibility tolerance of $1e-5$
 - Solutions rarely failed the checker; we discarded those without penalty
- We **excluded all instances where no submissions found a feasible solution** (1 in test set, 5 in hidden set)
- Runs that **fail to find a feasible solution receive a primal integral of 360 and a primal gap of 120%** (i.e. 1.2x on having always a 100% gap).
 - Given that instances lean difficult, we did not wish to penalize submissions too much for failing to find a feasible solution. That said, the relative ranking did not change with a 2x penalty.
- Metrics are aggregated across instances via **shifted geometric mean**, with a shift of 1 for primal integral and 1% for primal gap
- Primal integral and gap are with respect to a **strong reference value** (typically commercial solver run for hours, or provided by source)

Results

Bold+Underline: Best value
Underline: Second best value

Hidden set instances (55 instances)

	Primal Integral	Primal Gap (%)	# Insts. w/ Sol. Found	# Insts. within 1% of Comp. Best
Team 11 (Win.)	<u>54.95</u>	<u>12.91</u>	<u>51.3</u>	<u>19.3</u>
Team 2 (Hon.)	<u>63.84</u>	<u>15.18</u>	<u>50.0</u>	14.7
Team 9 (Hon.)	73.97	16.46	48.0	<u>16.3</u>
Team 3 (Round 2)	91.46	17.79	48.3	11.7
Team 5 (Round 2)	85.40	18.75	45.3	13.3
Team 1	312.79	100.80	16	0
Team 4	244.96	70.18	37	0
Team 6	211.58	61.98	29	0
Team 7	225.80	69.90	34	0
Team 8	170.82	53.62	36	3
Team 10	171.38	45.24	35	3

Test set instances (49 instances)

Primal Integral	Primal Gap (%)	# Insts. w/ Sol. Found
<u>52.69</u>	<u>11.95</u>	<u>47</u>
70.05	14.34	<u>44</u>
<u>64.82</u>	14.39	41
87.84	15.23	<u>44</u>
75.28	<u>14.09</u>	<u>44</u>
200.96	48.78	36
239.95	55.41	31
196.71	49.83	34
231.51	72.00	37
109.71	30.91	41
173.10	46.66	37

Standard deviations for multiple round-2 runs omitted for clarity and will be up on the website; not significant enough to alter the qualitative picture
 Recall that the number of instances only includes those where at least one submission found a solution

Results with cuOpt and Gurobi references

All instances (110 instances) W/T/L = Better by 1e-5 / Within 1e-5 / Worse by 1e-5

	Primal Integral			Primal Gap			# Insts. w/ Sol. Found
	Shifted Geomean	vs cuOpt (W/T/L)	vs Gurobi (W/T/L)	Shifted Geomean	vs cuOpt (W/T/L)	vs Gurobi (W/T/L)	
Team 11 (Win.)	59.83	55 / 8 / 47	51 / 3 / 56	13.81	45 / 23 / 42	30 / 8 / 72	98
Team 2 (Hon.)	72.98	54 / 10 / 46	36 / 5 / 69	16.61	47 / 21 / 42	31 / 9 / 70	92
Team 9 (Hon.)	75.80	42 / 13 / 55	37 / 5 / 68	17.24	35 / 25 / 50	31 / 11 / 68	89
Team 3 (R2)	98.12	35 / 10 / 65	32 / 5 / 73	19.63	35 / 22 / 53	30 / 9 / 71	92
Team 5 (R2)	87.46	34 / 11 / 65	35 / 5 / 70	18.45	33 / 29 / 48	31 / 11 / 68	90
cuOpt v26.04	76.02	-	41 / 5 / 64	14.10	-	32 / 12 / 66	95
Gurobi 13	38.98	64 / 5 / 41	-	5.86	66 / 12 / 32	-	104

Problem classes where at least one team consistently outperformed Gurobi in primal integral:

- Time-dependent TSP
- Railway network design
- N-Queens in higher dimensions
- Rectangle packing
- Open pit mining (rmine)



Values are only for reference and not a true comparison:

- cuOpt and Gurobi are **exact solvers** and must spend effort on optimality and bounds, whereas submissions are pure heuristics
- Gurobi is running on a **strict subset of CPU-only compute** on 8 threads
- Solver/compute variability could be better accounted for

Acknowledgement

The committee would like to acknowledge the strong performance of the runner-ups in round 2

Team 3: Stefano Gualandi, Paul Juenger, Andrea Lodi, Vinit Ranjan, and Bartolomeo Stellato
“First-Order Business in MIP: GPU-Accelerated Primal Heuristics for Mixed-Integer Linear Programs”

Team 5: Jean Pauphilet and Yupeng Wu
“PaNGEA: Parallel Node Generation and Exploration Algorithm”

Team 3 (Round 2)	91.46	17.79	48.3	11.7	87.84	15.23	<u>44</u>
Team 5 (Round 2)	85.40	18.75	45.3	13.3	75.28	<u>14.09</u>	<u>44</u>
Team 1	312.79	100.80	16	0	200.96	48.78	36
Team 4	244.96	70.18	37	0	239.95	55.41	31
Team 6	211.58	61.98	29	0	196.71	49.83	34
Team 7	225.80	69.90	34	0	231.51	72.00	37
Team 8	170.82	53.62	36	3	109.71	30.91	41
Team 10	171.38	45.24	35	3	173.10	46.66	37

Standard deviations for multiple round-2 runs omitted for clarity and will be up on the website; not significant enough to alter the qualitative picture
Recall that the number of instances only includes those where at least one submission found a solution

Prizes

\$1500 in NVIDIA Brev GPU credits for the winner

\$500 in NVIDIA Brev GPU credits for each of the honorable mentions (2x)




brev.nvidia.com

Cloud platform for NVIDIA GPUs

+ Opportunity for fast-track submissions to
Mathematical Programming Computation



Mathematical
Optimization Society



**2026 LAND-DOIG MIP COMPETITION
HONORABLE MENTION AWARD**

proudly presented on May 19, 2026 to


**Peng Lin, Zhe Wang,
Kewu Yang, and Shaowei Cai**

for the work entitled

**cuLocalMIP: A GPU-Oriented Local Search
Framework for Mixed Integer Programming**

MIP Competition Committee

Beste Basciftci, Akif Çördük, Gerald Gamrath, Christopher Hojny,
Jan Kronqvist, Haihao Lu, Christian Tjandraatmadja



**2026 LAND-DOIG MIP COMPETITION
HONORABLE MENTION AWARD**

proudly presented on May 19, 2026 to


Yiran Zhu

for the work entitled

**GPU-MIP: A GPU-Native Local Search Solver
for Mixed-Integer Programming**

MIP Competition Committee

Beste Basciftci, Akif Çördük, Gerald Gamrath, Christopher Hojny,
Jan Kronqvist, Haihao Lu, Christian Tjandraatmadja



2026 LAND-DOIG MIP COMPETITION
FIRST PLACE AWARD

proudly presented on May 19, 2026 to

**Timo Berthold, Ambros Gleixner, Alexander Hoen,
Nils-Christian Kempke, Thorsten Koch, Gioni Mexi,
Sebastian Pokutta, and Gennesaret Tjusila**

for the work entitled

CHAP - Coordinating Heuristics Across Platforms

MIP Competition Committee

Beste Basciftci, Akif Çördük, Gerald Gamrath, Christopher Hojny,
Jan Kronqvist, Haihao Lu, Christian Tjandraatmadja