PB26
Pseudo-Boolean Competition 2026

organized by Olivier Roussel,
steering committee:
Carlos Ansótegui, Johannes Fichte, Jakob Nordström, Olivier Roussel

SAT'26
July 23rd, 2026

What's new this year?

- We're back to Lisbon, where the story of the PB competition started with Vasco Manquinho 22 years ago!
- We're a bit more formal this year: we have nice certificates for the 3 top solvers in the 2 main tracks.
- Spoiler alert! The winner of the 2 main tracks was given yesterday. Do you remember?
- We have switched to a PAR-2 ranking that considers only final answers (UNSAT, OPT for PBO and SAT for PBS).
- The number of "easy" instances in the selection was limited.
- All the results already known in the competition database were used to check the answers. These are represented by a fake solver called "Best Known Results *oracle solver*" (0 CPU time on each instance).
- New hardware (with some challenges...)
- But as usual, all details are there: **https://www.cril.univ-artois.fr/PB26/**

PB26: some numbers

- 18 solvers submitters, 3 benchmarks submitters,
- around 30 solvers (similar names and same authors),
- 4 new teams
- 70 different solver versions
- more than 6 years of CPU time!
- 2503 instances submitted in the OPT-LIN track
- more than 5 TB of proofs generated, 700 GB of useful proofs (corresponding to an UNSAT/OPT answer)

List of solvers

New comers in bold font.

- AdaptExact (*Shiwei Pan, Wenxin Chu, Minghao Yin*)
- AI448PBSolver (*Yoshiya Imai*)
- AutoAdaptRoundingSat (*Yujiao Zhao, Yujie Yuan, Yiyuan Wang*)
- **AY (***Andrew Yates***)**
- **CPMpy* (***Tias Guns, Thomas Sergeys, Ignace Bleukx, Hendrik Bierlee***)**
- **DivingFalconSolver (***Daiki Tanaka***)**
- Exact (*Jo Devriendt*)
- exact-printemps (*Masahiro Sakai, Yuji Koguma*)
- HitPBO (*Hannes Ihalainen, Dieter Vandesande, André Schidler, Jeremias Berg, Matti Järvisalo*)
- Hybrid-CASHWMaxSATDisjComS-AdaptExact (*Shiwei Pan, Wenxin Chu, Minghao Yin*)
- Hybrid-CASHWMaxSATDisjComS-AdaptRoundingSat (*Yujiao Zhao, Yujie Yuan, Yiyuan Wang*)

List of solvers

- Hybrid-SynLSExact+CASHWMaxSATDisjComS (*Shiwei Pan, Yujie Yuan, Minghao Yin*)
- Hybrid-SynLSRS+CASHWMaxSATDisjComS (*Yujiao Zhao, Xince Sun, Yiyuan Wang*)
- Local-PB (*Peng Lin, Yi Chu, Mengchuan Zou, Shaowei Cai*)
- NaPS (*Masahiko Sakai and Hidetomo Nabeshima*)
- OR-Tools CP-SAT (*Laurent Perron*)
- Picat(kissat) (*Neng-Fa Zhou*)
- PRINTEMPS (*Yuji Koguma, Masahiro Sakai*)
- **Qiuqi (***Yue Zhang***)**
- roundingsat (*Markus Anders, Benjamin Bogø, Xiamin Chen, Wietze Koops, Pinyan Lu, Jakob Nordström, Alexander Tesch, Marc Vinyals, Qingzhao Wu*)
- Sat4j (*Daniel Le Berre*)

List of solvers

- SCIP-NaPS (*Masahiko Sakai and Hidetomo Nabeshima*)
- scip-printemps (*Masahiro Sakai, Yuji Koguma*)
- SynLSExact (*Shiwei Pan, Yujie Yuan, Minghao Yin*)
- SynLSRS (*Yujiao Zhao, Xince Sun, Yiyuan Wang*)
- UWrMaxSat-SCIP (*Marek Piotrów*)
- WMaxCDCL (*Jialu Zhang, Chu-Min Li, Sami Cherif, Shuolin Li, and Jinghu Liang*)

We have 2 new solvers generating proofs.

Many thanks to all submitters!

Pseudo-Boolean Constraints

- **Linear** (LIN) pseudo-Boolean (PB) constraint = sum of integer $\times$ a literal $\{\geq, \leq, =\}$ constant
  Example: $3x_1 - 3x_2 + 2\bar{x}_3 + \bar{x}_4 + x_5 \geq 5$
- **Non-linear** (NLC) pseudo-Boolean (PB) constraint = sum of integer $\times$ a product of literals $\{\geq, \leq, =\}$ constant
  Example: $3x_1\bar{x}_2 - 3x_2x_4 + 2\bar{x}_3 + \bar{x}_4 + x_5x_6x_7 \geq 5$
- As an example, PB allows compact encodings of:
  - cardinalities: $x_1 + x_2 + x_3 \geq 2$
  - adder (C=A+B): $2c_1 + c_0 = 2a_1 + a_0 + 2b_1 + b_0$
  - knapsack:
    $max : 5x_1 + 10x_2 + 2x_3;$
    $5x_1 + 8x_2 + x_3 \leq 10$
  - integer factorization (X*Y=3): $x_0y_0 + 2x_0y_1 + 2x_1y_0 + 4x_1y_1 = 3$
- Cutting-planes proof system stronger than resolution: PHP easily solved in polynomial time

Benchmark categories (1)

Classification based on the linearity of constraints

- LIN: All constraints are linear
- NLC: At least one constraint is non linear (contains products of literals)

Classification based on the objective function

- DEC: No objective function to optimize (decision problem). The solver must simply find a solution.
- OPT: An objective function is present. The solver must find a solution with the best possible value of the objective function.

Classification based on the existence of hard/soft clauses (generalization of MaxSAT with a top cost as in WCSP)

- SOFT: All constraints may be violated if needed (no hard constraint). Minimize the weight of unsatisfied constraints.
- PARTIAL: At least one hard constraint.

No classification on the size of integers, but solvers that cannot deal with large integers must answer NS (no support). The limit is solver specific. Solvers expected to support at least 64 bits integers.

Tracks

Listed by decreasing number of registered solvers:

- **OPT-LIN: 57 solvers, 585 instances selected**
- **DEC-LIN: 47 solvers, 492 instances selected**
- OPT-NLC: 20 solvers, 63 instances selected
- DEC-NLC: 20 solvers, 10 instances selected
- PARTIAL-LIN: 14 solvers, 208 instances selected
- SOFT-LIN: 14 solvers, 60 instances selected
- **DEC-LIN-CERT and OPT-LIN-CERT: 8 solvers, same selection as in DEC-LIN/OPT-LIN, solvers also ranked in DEC-LIN/OPT-LIN**

New instances

All new instances belong to the OPT-LIN category.

- Daiki Tanaka: 156 instances
  - market split instances
- Jo Devriendt: 58 instances
  - encoding of the small and medium sized instances from the old TSPLIB
- Thomas Sergeys: 2289 instances
  - translation of JSPLib, PSPLib and MIPLib instances.
  - RCP instances and several MIP instances had the same names as older benchmarks. They were probably duplicates, possibly with a different encoding. For safety, they were not considered this year.

Many thanks to all submitters!

Selection of instances

- By default, selection of a fixed number of instances selected per domain (instances with similar characteristics)
- 15 instances per domain in the DEC-LIN track
- 10 instances randomly selected per domain in the other tracks
- If a domain doesn't contain enough instances, they are all selected.
- At most 5 % of the whole instance set selected per submitter.
- Duplicates removed (syntactical identification).
- **[new] At most 20% of easy instances.**
- Instances submitted this year were selected randomly.

Easy instances

- No perfect definition of an easy instance!
- Choice of a limited set of solvers to evaluate the hardness: 3 best contributors to the VBS in PB25 + the last 2 generic solvers in the ranking.
- Instances solved by **all** these reference solvers in less that 60 seconds are considered easy.
- 7430 easy instances out of 20906 in the DEC-LIN track (35,5 %)
- 2352 easy instances out of 15268 in the OPT-LIN track (15.4 %)
- The 20 % limit is satisfied for OPT-LIN
- In the DEC-LIN track, the goal was to select 3 easy instances and 12 harder instances in each domain (when possible)

Evaluation environment

kindly provided by CRIL, University of Artois, France

- Each solver was given a time limit of 1 hour CPU time (1 hour wall-clock for the few parallel solvers, which ran on 16 cores $\Longrightarrow$ 16 hours CPU time))
- The size limit for proofs of unsatisfiability/optimality was set to 100 GB.
- VeriPB+CakePB was used to verify the proofs and were allowed to run for 5 hours (CPU time)
- We used the same cluster for all experiments: nodes with two AMD EPYC 9355 processors, 3.55GHz, 32 cores (64 hyperthreads), 256 MB cache, DDR5-6400, 1TiB RAM.
- For sequential solvers, 16 jobs were run in parallel (*4 cores/job, 62 GB/job)*. This configuration was tested on a sample of benchmarks and instances and induces a slowdown of at most 10% compared to a solver running alone on the host.
- For parallel solvers: 4 jobs in parallel, *16 cores/job, 248 GB/job*

Verification of results

- The usual verifications are performed:
  - the models given by the solvers are checked
  - the answers given by the different solvers on a given instance are checked for consistency
  - **[new] a fake solver named "Best Known Results** *oracle solver*" **was introduced to represent the results already known locally.**
- The unsatifiability/optimality proofs in the *-CERT tracks where checked by VeriPB+CakePB. Answers OPTC and UNSC (C for Certified).
- Solvers giving a wrong answer in a category are disqualified in that category. Up to 3 submissions of bug fix were allowed.
- Not all UNSATISFIABLE and OPTIMUM FOUND answers could be checked and therefore some results should be taken with caution.

Official rankings

Ranking of solvers and Virtual Best Solver (VBS)

The two main rankings (DEC-LIN and OPT-LIN):

- are based on the PAR-2 score
- consider only final answer (UNSAT, OPT for optimization and SAT for decision problems)
- consider only sequential solvers (it wouldn't make sense to include parallel solvers in a PAR-2 ranking without additional limits)
- consider all instances (including those unsupported by some solvers)

Several other rankings are available on the web site.

The Virtual Best Solver (VBS)

- is the virtual solver obtained by combining the best results of all submitted solvers.
- **can be obtained by running in parallel all submitted solvers**
- represents the current state of the art (SOTA)
- is a reference for the evaluation of the other solvers

All instances, including those unsupported?

- Solvers may answer 'No Support' when they know at parse time they can't solve an instance. This is meant to deal with the problem of big integers support.
- It is expected that the set of unsupported instances will be more or less the same for all solvers with no full support.
- DivingFalconSolver is not a PB solver. It can only solve the Market Split problem. It answers 'No Support' on all but 9 instances and solves 5 of them (6 for the parallel configuration).
- **It doesn't make sense to generate a ranking on the set of instances supported by all solvers!**
- CPMpy also generates $\sim$100 'No Support' answers. All the other solvers generate less than 25 'No Support' answer.
- The web site gives the rankings on the set of instances supported by all solvers but DivingFalconSolver and CPMpy

Official DEC-LIN Ranking

Rank	Solver	PAR-2	#solved	Detail	%inst.	%VBS
	Total number of instances: 492					
	Virtual Best Solver (VBS)	-	464	183 SAT, 281 UNS	94%	100%
1	**Local-PB**	704,150	399	155 SAT, 244 UNS	81%	86%
2	**Hybrid-SynLSExact+CASHWMaxSATDisjComS** *v1*	745,047	398	158 SAT, 240 UNS	81%	86%
3	**roundingsat-satsuma-MIR**	746,148	393	144 SAT, 249 UNS	80%	85%
4	roundingsat-satsuma-MIR (cert)	751,560	392	143 SAT, 249 UNS	80%	84%
5	Hybrid-SynLSExact+CASHWMaxSATDisjComS *v2*	755,223	397	157 SAT, 240 UNS	81%	86%
6	Hybrid-SynLSRS+CASHWMaxSATDisjComS *v1*	765,389	398	159 SAT, 239 UNS	81%	86%
7	Hybrid-SynLSRS+CASHWMaxSATDisjComS *v2*	778,888	397	159 SAT, 238 UNS	81%	86%
8	roundingsat-satsuma-liftedcover	789,326	389	141 SAT, 248 UNS	79%	84%
9	roundingsat-satsuma-liftedcover (cert)	791,937	389	141 SAT, 248 UNS	79%	84%
10	cpmpy_ortools (1 core)	813,229	383	156 SAT, 227 UNS	78%	83%
11	roundingsat-satsuma-saturation	826,065	383	140 SAT, 243 UNS	78%	83%
12	SynLSExact	826,586	383	153 SAT, 230 UNS	78%	83%
13	roundingsat-satsuma-saturation (cert)	832,519	382	139 SAT, 243 UNS	78%	82%

Official DEC-LIN Ranking (2)

Rank	Solver	PAR-2	#solved	Detail	%inst.	%VBS
14	UWrMaxSat-SCIP *2.0.4*	853,710	383	154 SAT, 229 UNS	78%	83%
15	SCIP-NaPS	871,182	385	159 SAT, 226 UNS	78%	83%
16	NaPS *1.04*	878,203	384	158 SAT, 226 UNS	78%	83%
17	SynLSRS	882,090	377	149 SAT, 228 UNS	77%	81%
18	Exact (cert)	982,664	362	140 SAT, 222 UNS	74%	78%
19	Exact	1,001,279	359	138 SAT, 221 UNS	73%	77%
20	Exact_symbreak	1,013,752	357	133 SAT, 224 UNS	73%	77%
21	Exact_symlocal	1,024,730	355	134 SAT, 221 UNS	72%	77%
22	cpmpy_highs (1 core)	1,172,470	336	132 SAT, 204 UNS	68%	72%
23	AI448PBSolver	1,187,165	331	120 SAT, 211 UNS	67%	71%
24	cpmpy_highs (1 core)	1,206,841	330	129 SAT, 201 UNS	67%	71%
25	exact-printemps (inc, 1 core)	1,215,086	326	140 SAT, 186 UNS	66%	70%
26	OR-Tools CP-SAT (1 core)	1,238,779	324	141 SAT, 183 UNS	66%	70%
27	AY	1,343,663	315	120 SAT, 195 UNS	64%	68%
28	Picat(kissat)	1,385,690	307	150 SAT, 157 UNS	62%	66%
29	cpmpy_pindakaas_cadical	1,484,996	294	131 SAT, 163 UNS	60%	63%

Official DEC-LIN Ranking (3)

Rank	Solver	PAR-2	#solved	Detail	%inst.	%VBS
30	scip-printemps (inc, 1 core)	1,512,417	286	127 SAT, 159 UNS	58%	62%
31	AY-cert *(cert)*	1,621,803	273	115 SAT, 158 UNS	55%	59%
32	AY	1,712,560	289	112 SAT, 177 UNS	59%	62%
33	Sat4j Resolution	1,739,486	258	115 SAT, 143 UNS	52%	56%
34	Sat4j Resolution CERT	1,743,731	258	115 SAT, 143 UNS	52%	56%
35	Sat4j CuttingPlanes	2,002,352	219	68 SAT, 151 UNS	45%	47%
36	Sat4j CuttingPlanes CERT	2,002,704	219	68 SAT, 151 UNS	45%	47%
37	PRINTEMPS (inc, 1 core)	2,668,849	124	122 SAT, 2 UNS	25%	27%

Official OPT-LIN ranking (1)

Rank	Solver	PAR-2	#solved	Detail	%inst.	%VBS
Total number of instances: 585						
	Virtual Best Solver (VBS)	-	444	425 OPT, 19 UNS	76%	100%
1	**Local-PB**	1,607,474	369	353 OPT, 16 UNS	63%	83%
2	**UWrMaxSat-SCIP** *2.0.4*	1,758,090	348	331 OPT, 17 UNS	59%	78%
3	**SCIP-NaPS**	1,854,663	340	323 OPT, 17 UNS	58%	77%
4	UWrMaxSat-SCIP *2.0.3*	1,855,028	334	317 OPT, 17 UNS	57%	75%
5	AutoAdaptRoundingSat	1,894,517	331	317 OPT, 14 UNS	57%	75%
6	Hybrid-CASHWMaxSATDisjComS-AdaptRndingSat *v2*	1,911,915	354	338 OPT, 16 UNS	61%	80%
7	WMaxCDCL-S9	1,919,248	333	319 OPT, 14 UNS	57%	75%
8	Hybrid-CASHWMaxSATDisjComS-AdaptExact *v1*	1,924,261	331	315 OPT, 16 UNS	57%	75%
9	Hybrid-CASHWMaxSATDisjComS-AdaptExact *v2*	1,926,804	331	315 OPT, 16 UNS	57%	75%
10	Hybrid-CASHWMaxSATDisjComS-AdaptRndingSat *v1*	1,927,025	351	335 OPT, 16 UNS	60%	79%
11	Qiuqi (1 core)	1,928,386	328	313 OPT, 15 UNS	56%	74%
12	roundingsat-satsuma-opt	1,935,528	322	309 OPT, 13 UNS	55%	73%
13	roundingsat-opt	1,944,048	321	308 OPT, 13 UNS	55%	72%

Official OPT-LIN Ranking (2)

Rank	Solver	PAR-2	#solved	Detail	%inst.	%VBS
14	NaPS *1.04*	1,945,754	323	306 OPT, 17 UNS	55%	73%
15	roundingsat-satsuma-opt (cert)	1,954,918	319	306 OPT, 13 UNS	55%	72%
16	roundingsat-opt (cert)	1,963,255	318	305 OPT, 13 UNS	54%	72%
17	Exact	2,024,347	311	296 OPT, 15 UNS	53%	70%
18	Exact_symlocal	2,033,160	310	295 OPT, 15 UNS	53%	70%
19	HitPBO	2,046,729	308	299 OPT, 9 UNS	53%	69%
20	Exact_symbreak	2,049,355	307	292 OPT, 15 UNS	52%	69%
21	AutoAdaptRoundingSat *v2*	2,052,435	327	313 OPT, 14 UNS	56%	74%
22	AdaptExact	2,060,956	306	291 OPT, 15 UNS	52%	69%
23	AdaptExact *v2*	2,061,389	306	291 OPT, 15 UNS	52%	69%
24	Exact (cert)	2,125,312	297	282 OPT, 15 UNS	51%	67%
25	cpmpy_ortools (1 core)	2,172,999	289	276 OPT, 13 UNS	49%	65%
26	cpmpy_ortools (1 core)	2,187,297	287	274 OPT, 13 UNS	49%	65%
27	roundingsat-ihs	2,209,593	283	270 OPT, 13 UNS	48%	64%
28	roundingsat-ihs (cert)	2,214,194	283	270 OPT, 13 UNS	48%	64%
29	cpmpy_highs (1 core)	2,281,375	277	266 OPT, 11 UNS	47%	62%

Official OPT-LIN Ranking (3)

Rank	Solver	PAR-2	#solved	Detail	%inst.	%VBS
30	scip-printemps (inc, 1 core)	2,314,974	264	249 OPT, 15 UNS	45%	59%
31	cpmpy_highs (1 core)	2,321,453	270	259 OPT, 11 UNS	46%	61%
32	WMaxCDCL	2,414,370	258	248 OPT, 10 UNS	44%	58%
33	exact-printemps (inc, 1 core)	2,438,733	247	233 OPT, 14 UNS	42%	56%
34	HitPBO-PL *(cert)*	2,467,609	248	234 OPT, 14 UNS	42%	56%
35	OR-Tools CP-SAT (1 core)	2,588,681	230	215 OPT, 15 UNS	39%	52%
36	Picat(kissat)	2,592,556	235	220 OPT, 15 UNS	40%	53%
37	AY	2,632,333	226	214 OPT, 12 UNS	39%	51%
38	Sat4j Resolution	3,131,533	153	141 OPT, 12 UNS	26%	34%
39	Sat4j Resolution CERT	3,136,626	152	140 OPT, 12 UNS	26%	34%
40	Sat4j CuttingPlanes	3,166,251	148	137 OPT, 11 UNS	25%	33%
41	Sat4j CuttingPlanes CERT	3,166,728	148	137 OPT, 11 UNS	25%	33%
42	AY	3,226,224	150	138 OPT, 12 UNS	26%	34%
43	AY-cert *(cert)*	4,140,000	10	10 UNS	2%	2%
44	PRINTEMPS (inc, 1 core)	4,140,002	10	10 UNS	2%	2%
45	DivingFalconSolver (1 core)	4,179,606	5	5 OPT	1%	1%

FLoC Olympics 2026

Local-PB is the winner of the two main tracks of the competition (DEC-LIN and OPT-LIN), which is a significant achievement.

For this reason, it was decided to give a special award at the FLoC Olympics to this solver.

The team Peng LIN, Yi CHU, Mengchuan ZOU, Shaowei CAI is therefore the recipient of the

”PB26: Grand Slam award”

And parallel solvers?

DEC-LIN ranking on #solved, max CPU time for any solver=3601 s, all instances

Rank	Solver	PAR-2	#solved	Detail	%inst.	%VBS
Total number of instances: 492						
Virtual Best Solver (VBS)		-	464	183 SAT, 281 UNS	94%	100%
1	**cpmpy_ortools (16 cores)**	698,414	401	158 SAT, 243 UNS	82%	86%
2	Local-PB	704,150	399	155 SAT, 244 UNS	81%	86%
3	**OR-Tools CP-SAT (16 cores)**	711,279	398	157 SAT, 241 UNS	81%	86%
4	Hybrid-SynLSExact+CASHWMaxSATDisjComS *v1*	745,047	398	158 SAT, 240 UNS	81%	86%
5	Hybrid-SynLSRS+CASHWMaxSATDisjComS *v1*	765,389	398	159 SAT, 239 UNS	81%	86%
6	Hybrid-SynLSExact+CASHWMaxSATDisjComS *v2*	755,223	397	157 SAT, 240 UNS	81%	86%
7	Hybrid-SynLSRS+CASHWMaxSATDisjComS *v2*	778,888	397	159 SAT, 238 UNS	81%	86%
8	roundingsat-satsuma-MIR	746,148	393	144 SAT, 249 UNS	80%	85%
...						

Note: the comparison is not perfect because the ratio CPU time/wall-clock time is different from 16 (8 for CPMpy, 11.6 for OR-Tools)

And parallel solvers?

OPT-LIN ranking on #solved, max CPU time for any solver=3601 s, all instances

Rank	Solver	PAR-2	#solved	Detail	%inst.	%VBS
Total number of instances: 585						
	Virtual Best Solver (VBS)	-	444	425 OPT, 19 UNS	76%	100%
1	**cpmpy_ortools (16 cores)**	1,575,859	376	360 OPT, 16 UNS	64%	85%
2	**OR-Tools CP-SAT (16 cores)**	1,593,076	374	356 OPT, 18 UNS	64%	84%
3	Local-PB	1,607,474	369	353 OPT, 16 UNS	63%	83%
4	Hybrid-CASHWMaxSATDisjComS-AdaptRndingSat *v2*	1,911,915	354	338 OPT, 16 UNS	61%	80%
5	Hybrid-CASHWMaxSATDisjComS-AdaptRndingSat *v1*	1,927,025	351	335 OPT, 16 UNS	60%	79%
6	UWrMaxSat-SCIP *2.0.4*	1,758,090	348	331 OPT, 17 UNS	59%	78%
7	SCIP-NaPS	1,854,663	340	323 OPT, 17 UNS	58%	77%
		...				
13	Qiuqi (1 core)	1,928,386	328	313 OPT, 15 UNS	56%	74%
14	**Qiuqi (16 cores)**	1,918,472	327	312 OPT, 15 UNS	56%	74%

Note: the comparison is not perfect because the ratio CPU time/wall-clock time is different from 16 (8.4 for CPMpy, 13.1 for OR-Tools, 9.4 for Qiuqi)

Rankings in the CERT tracks

DEC-LIN-CERT (only certified answers)

Rank	Solver	PAR-2	#solved	Detail
Total number of instances: 492				
1	roundingsat-satsuma-MIR (cert)	786,137	387	143 SAT, 244 UNSATC
2	roundingsat-satsuma-saturation (cert)	871,200	376	139 SAT, 237 UNSATC
3	roundingsat-satsuma-liftedcover (cert)	873,845	376	141 SAT, 235 UNSATC
4	Exact (cert)	1,015,290	356	140 SAT, 216 UNSATC
5	AY-cert *(cert)*	1,621,803	273	115 SAT, 158 UNSATC
6	Sat4j Resolution CERT	1,743,731	258	115 SAT, 143 UNSATC
7	Sat4j CuttingPlanes CERT	2,002,704	219	68 SAT, 151 UNSATC

DEC-LIN-CERT ranking (counting uncertified answers)

Rank	Solver	PAR-2	#solved	Detail
Total number of instances: 492				
1	roundingsat-satsuma-MIR (cert)	751,560	392	143 SAT, 5 UNSAT, 244 UNSATC
2	roundingsat-satsuma-liftedcover (cert)	791,937	389	141 SAT, 13 UNSAT, 235 UNSATC
3	roundingsat-satsuma-saturation (cert)	832,519	382	139 SAT, 6 UNSAT, 237 UNSATC
4	Exact (cert)	982,664	362	140 SAT, 6 UNSAT, 216 UNSATC
5	AY-cert *(cert)*	1,621,803	273	115 SAT, 158 UNSATC
6	Sat4j Resolution CERT	1,743,731	258	115 SAT, 143 UNSATC
7	Sat4j CuttingPlanes CERT	2,002,704	219	68 SAT, 151 UNSATC

OPT-LIN-CERT (certified answers only)

Rank	Solver	PAR-2	#solved	Detail
Total number of instances: 585				
1	roundingsat-opt (cert)	1,969,867	317	304 OPTC, 13 UNSATC
2	roundingsat-satsuma-opt (cert)	2,028,457	308	295 OPTC, 13 UNSATC
3	roundingsat-ihs (cert)	2,227,546	281	269 OPTC, 12 UNSATC
4	Exact (cert)	2,342,226	262	247 OPTC, 15 UNSATC
5	HitPBO-PL *(cert)*	2,543,970	234	220 OPTC, 14 UNSATC
6	Sat4j Resolution CERT	3,152,046	149	137 OPTC, 12 UNSATC
7	Sat4j CuttingPlanes CERT	3,166,728	148	137 OPTC, 11 UNSATC

AY has generated proofs that were not validated.

OPT-LIN-CERT ranking (counting uncertified answers)

Rank	Solver	PAR-2	#solved	Detail
Total number of instances: 585				
1	roundingsat-satsuma-opt (cert)	1,954,918	319	11 OPT, 295 OPTC, 13 UNSATC
2	roundingsat-opt (cert)	1,963,255	318	1 OPT, 304 OPTC, 13 UNSATC
3	Exact (cert)	2,125,312	297	35 OPT, 247 OPTC, 15 UNSATC
4	roundingsat-ihs (cert)	2,214,194	283	1 OPT, 269 OPTC, 1 UNSAT, 12 UNSATC
5	HitPBO-PL *(cert)*	2,467,609	248	14 OPT, 220 OPTC, 14 UNSATC
6	Sat4j Resolution CERT	3,136,626	152	3 OPT, 137 OPTC, 12 UNSATC
7	Sat4j CuttingPlanes CERT	3,166,728	148	137 OPTC, 11 UNSATC

AY has generated proofs that were not validated.

The I/O problem in the CERT tracks

- Last year, we had only 4 jobs per node and the proofs could be stored on /tmp. No problem.
- With the new hardware, we want to run 16 jobs at the same time on a node.
- Solvers in the CERT tracks may generate files of up to 100 GB, with a throughput of 50 MB/s.
- Running 16 jobs on a machine, on 10 machines means that in on hour up to 16 GB of proofs could be generated at up to 8 GB/s (in the wort case).
- The file server is not so fast.
- The solution implemented was to write a filesystem that caches the proofs in memory and compresses them on the fly (details on the web site).
- This filesystem uses the resources of 3 jobs (12 cores, 192 GiB) to compress the proofs, but at least that allows to run 13 jobs simultaneously.
- **The solver doesn't have to pay the price of compression.**

Take home message

Benefits of a competition:

- evaluate solvers in the same conditions
- help collecting publicly available benchmarks in one place
- help identifying new solvers and new ideas on the market
- help debug solvers and verifiers

Caveats:

- don't take the rankings too seriously!
- the set of instances does matter (obviously)!
- the information collected cannot be reduced to a simple ranking.
- **there are many different ways to look at the results!**

More information

- All details are on the web site **https://www.cril.univ-artois.fr/PB26/**
- Questions can be addressed to the organizer: roussel @ cril . fr.
- Thanks to all participants and submitters!
- Feedback is welcome!
- **Next competition in 2027.**
- **We need new bencharks!**
- Write solvers, generate instances, be ready to submit!

Links

- PB26: https://www.cril.univ-artois.fr/PB26/
- VeriPB: https://gitlab.com/MIAOresearch/software/VeriPB
- CakePB: https://gitlab.com/MIAOresearch/software/cakepb