Fifth Pseudo-Boolean Competition PB10

Vasco MANQUINHO and Olivier ROUSSEL

13th International Conference on
Theory and Applications of Satisfiability Testing,
SAT'10

July 14th, 2010

Outline

- Pseudo-Boolean constraints
- PBS, PBO, WBO
- Judges
- Benchmarks and Solvers
- Evaluation Environment
- Results

Linear Pseudo-Boolean Constraints

- A **linear** pseudo-Boolean (PB) constraint may be defined over Boolean variables by

$$\sum_i a_i.l_i \geq d \text{ with } a_i, d \in \mathbb{Z}, l_i \in \{x_i, \bar{x}_i\}, x_i \in \mathbb{B}$$

  Example: $3x_1 - 3x_2 + 2\bar{x}_3 + \bar{x}_4 + x_5 \geq 5$
- Extends both clauses and cardinality constraints
  - cardinalities: all $a_i = 1$ and $d > 1$
  - clauses: all $a_i = 1$ and $d = 1$
- PB constraints are more expressive than clauses (one PB constraint may replace an exponential number of clauses)
- A pseudo-Boolean instance is a conjunction of PB constraints

Non-Linear Pseudo-Boolean Constraints

- A **non-linear** pseudo-Boolean constraint may be defined over Boolean variables by

$$\sum_i a_i(\prod_j l_{i,j}) \geq d \text{ with } a_i, d \in \mathbb{Z}, l_{i,j} \in \{x_{i,j}, \bar{x}_{i,j}\}, x_{i,j} \in \mathbb{B}$$

  Example: $3x_1\bar{x}_2 - 3x_2x_4 + 2\bar{x}_3 + \bar{x}_4 + x_5x_6x_7 \geq 5$
- A product is a AND
- Compact encoding for several problems (e.g. factoring problem encoded by one constraint)
- Can be easily translated into linear pseudo-Boolean by introducing new variables and constraints such that

$$p \leftrightarrow x_0 \wedge x_1 \wedge \ldots \wedge x_n$$

  (requires 2 PB constraints or n+1 clauses)

Different problems: PBS, PBO,...

- **PBS (Pseudo Boolean Satisfaction)**
  decide of the satisfiability of a conjunction of PB constraints
- **PBO (Pseudo Boolean Optimization)**
  find a model of a conjunction of PB constraints which optimizes one objective function

$$\begin{cases} \text{minimize} & f = \sum_i c_i.x_i \text{ with } c_i \in \mathbb{Z}, x_i \in \mathbb{B} \\ \text{subject to} & \text{the conjunction of constraints} \end{cases}$$

Different problems: ... and WBO

WBO (Weighted Boolean Optimization)

- **new in the competition**
- generalization of maximum satisfiability for PB constraints
- hard constraints **must** be satisfied
- soft constraints may be violated, but this has a cost
- the cost of an interpretation is the sum of the costs of violated soft constraints
- as in WCSP, there is a top cost. Interpretations with a cost greater or equal to the top cost are non admissible.
- the goal is to find an admissible interpretation with the smallest cost
- **to avoid any intersection with the Max-SAT competition, at least one constraint must not be a clause.**

Judges

- 2 judges (the same as last year)
  - Heidi Dixon (pbChaff solver)
  - Peter Barth (opbdp solver)
- decided of the selection of instances
- suggested a comparison with CPLEX
- approved the results

Benchmark categories (1)

For PBS/PBO, 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.

For WBO, classification based on the existence of hard clauses

SOFT
: No hard clause at all.

PARTIAL
: At least one hard clause.

Benchmark categories (2)

Classification based on the size of coefficients

SMALLINT
: small integers: no constraint with a sum of coefficients greater than 2^{20} (20 bits): expected to be safe for solvers using 32 bits integers and simple techniques (be careful with learning), but strong limit to the encoding of concrete problems.

BIGINT
: big integers: at least one constraint with a sum of coefficients greater than 2^{20} (20 bits): requires arbitrary precision.

Classification based on the linearity of constraints

LIN
: All constraints are linear

NLC
: At least one constraint is non linear (contains products of literals)

Categories

- **DEC-SMALLINT-LIN (452 instances)**
- **DEC-SMALLINT-NLC (100 instances)**
- DEC-BIGINT-LIN
- DEC-BIGINT-NLC
- **OPT-SMALLINT-LIN (699 instances)**
- **OPT-SMALLINT-NLC (409 instances)**
- **OPT-BIGINT-LIN (532 instances)**
- OPT-BIGINT-NLC
- **PARTIAL-SMALLINT-LIN (536 instances)**
- **PARTIAL-BIGINT-LIN (263 instances)**
- **SOFT-SMALLINT-LIN (201 instances)**
- SOFT-BIGINT-LIN (46 instances)

Submitted solvers: (1)

7 teams, 8 solvers, 30 solver versions

Solvers with only PBS/PBO support

borg-pb Bryan Silverthorn
a portfolio solver (In Python. Uses clasp, SAT4J and the PB10 versions of bsolo/wbo)

bsolo Vasco Manquinho and José Santos
a SAT-like solver with lower bound estimation techniques

PBPASSolver Amir Aavani
written in Pascal

PB-wave Cédric Piette
a local search solver

Submitted solvers (2)

Solvers with both PBS/PBO and WBO support

PB/CT Anders Franzen, Roberto Bruttomesso
based on OpenSMT

SAT4J Pseudo Daniel Le Berre and Anne Parrain
3 versions: learn clauses, learn PB constraints, run both in // and exchange intermediate values of the objective function

SCIP Stefan Heinz, Marc E. Pfetsch, and Michael Winkler
3 versions: with SoPlex as LP solver, with Clp as LP solver, without any LP solver

wbo Vasco Manquinho, Jordi Planes and João Marques-Silva
an unsatisfiability-based solver; iterates over the identification of unsatisfiable subformulas;

An extra solver: pb_cplex

- a direct interface to CPLEX 12.1, a state of the art linear programming solver
- support for PBS/PBO as well as WBO
- written by Vasco Manquinho after a suggestion of the judges

PBS/PBO Instances submitted this year

- resource-constrained project scheduling problem (A. Oliveras)
  converted from the PSPLib
  6216 submitted instances, (4080 DEC-SMALLINT-LIN, 2040 OPT-SMALLINT-LIN)
  80 instances randomly selected in each category
- dependency of packages in a Linux distribution (D. Le Berre)
  converted from the Mancoosi project
  1 DEC-SMALLINT-LIN, 65 OPT-SMALLINT-LIN, 327 OPT-BIGINT-LIN
  at most 80 instances randomly selected in each category
- Tolerant Algebraic Side-Channel Attack (TASCA) on the Keeloq cipher (Y. Oren)
  4 OPT-SMALLINT-NLC

WBO Instances

- no submission at all !!
- generation of WBO from unsatisfiable PBS/PBO instances by adding a random cost between 1 and 100 to
  - 100% of the constraints (only soft constraints)
  - 66% of the constraints (majority of soft constraints)
  - 33% of the constraints (majority of hard constraints)

  No top cost imposed in these instances.
- Conversion of WCSP instances
  - 1 hard equality constraint to encode each variable
  - 1 soft constraint to encode the cost of a tuple

Evaluation environment

kindly provided by the CRIL, University of Artois, France
For PBS/PBO: same hardware as last competitions

- Cluster of bi-Xeon 3 GHz, 2MB cache, 2GB RAM
- Each solver was given a time limit of 30 minutes (1800s) and a memory limit of 1800 MB (to avoid swapping).
- 280 days of CPU time used

For WBO: new hardware

- Cluster of bi-Xeon quad-core 2.66 GHz, 8 MB cache, 32 GB RAM
- Each solver was given a time limit of 30 minutes (1800s) and a memory limit of 3800 MB (to avoid swapping).
- 2 solvers per node (limited interactions because of the 2 CPU and the memory limit)
- 90 days of CPU time used

Verification of results

- The environment performs the following, efficient checks:
  - for SATISFIABLE answers, solvers must output a complete instantiation and the system checks that it satisfies all constraints
  - for UNSATISFIABLE answers, the system only checks that no other solver proved satisfiability
  - for OPTIMUM FOUND answers, solvers must output a complete instantiation; the system checks if all constraints are satisfied and that no other solver found a better solution
- UNSATISFIABLE and OPTIMUM FOUND answers cannot be completely checked efficiently and therefore should be taken with caution.
- Solvers giving a wrong answer in a category are disqualified in that category.

Ranking of solvers and Virtual Best Solver (VBS)

Ranking based on two criteria:

1. the number of solved instances
2. ties are broken by considering the cumulated time on solved instances

The Virtual Best Solver (VBS)

- is the virtual solver obtained by combining the best results of all submitted solvers.
- could 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

Results for DEC-SMALLINT-LIN

Rank	Solver	#solved	Detail	%inst.	%VBS
Total number of instances: 452					
Virtual Best Solver (VBS)		434	180 S, 254 U	96%	100%
1	borg-pb	415	179 S, 236 U	92%	96%
2	SAT4J Res//CP	382	173 S, 209 U	85%	88%
3	bsolo *3.2 Card*	380	172 S, 208 U	84%	88%
4	wbo *1.4a*	378	171 S, 207 U	84%	87%
5	PB/CT *bugfix*	369	164 S, 205 U	82%	85%
6	SAT4J Res.	367	174 S, 193 U	81%	85%
7	bsolo *3.2 Cl*	355	170 S, 185 U	79%	82%
8	SCIPspx *bugfix*	351	139 S, 212 U	78%	81%
9	SCIPspx	351	141 S, 210 U	78%	81%
10	SCIPclp	344	144 S, 200 U	76%	79%
11	pb_cplex	337	155 S, 182 U	75%	78%
12	SCIPnone	288	154 S, 134 U	64%	66%
13	SAT4J CP	228	106 S, 122 U	50%	53%
14	PB-wave	66	66 S	15%	15%

DEC-SMALLINT-LIN

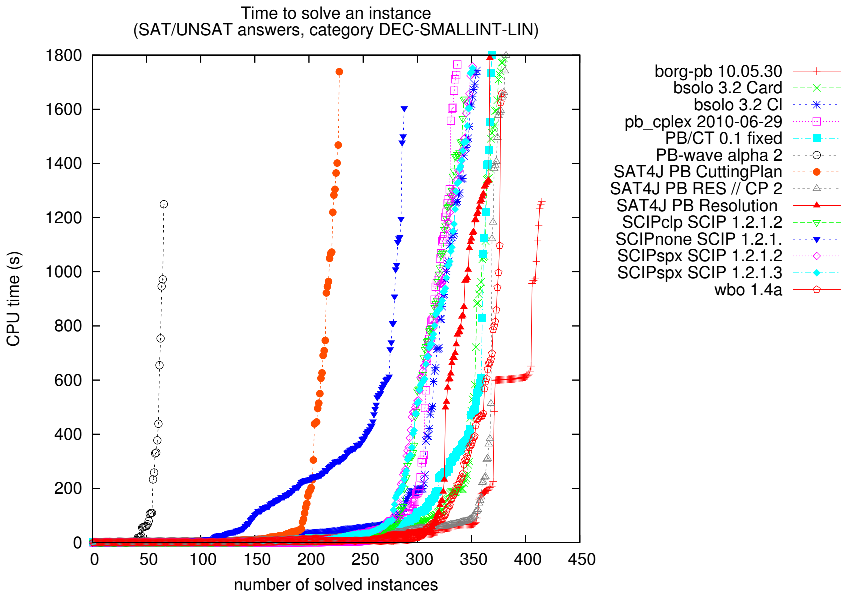

Results for DEC-SMALLINT-NLC

Rank	Solver	#solved	Detail	%inst.	%VBS
Total number of instances: 100					
Virtual Best Solver (VBS)		70	50 S, 20 U	70%	100%
1	pb_cplex	70	50 S, 20 U	70%	100%
2	SCIPspx *bugfix*	70	50 S, 20 U	70%	100%
3	SCIPclp	69	50 S, 19 U	69%	99%
4	SCIPspx	69	50 S, 19 U	69%	99%
5	SAT4J Res//CP	65	50 S, 15 U	65%	93%
6	SAT4J CP	65	50 S, 15 U	65%	93%
7	PB/CT	65	50 S, 15 U	65%	93%
8	PB/CT *bugfix*	63	50 S, 13 U	63%	90%
9	bsolo *3.2 Card*	61	46 S, 15 U	61%	87%
10	wbo *1.4a*	57	42 S, 15 U	57%	81%
11	SCIPnone	49	39 S, 10 U	49%	70%
12	borg-pb	27	17 S, 10 U	27%	39%
13	bsolo *3.2 Cl*	26	16 S, 10 U	26%	37%
14	PB-wave	25	25 S	25%	36%
15	SAT4J Res.	25	10 S, 15 U	25%	36%

DEC-SMALLINT-NLC

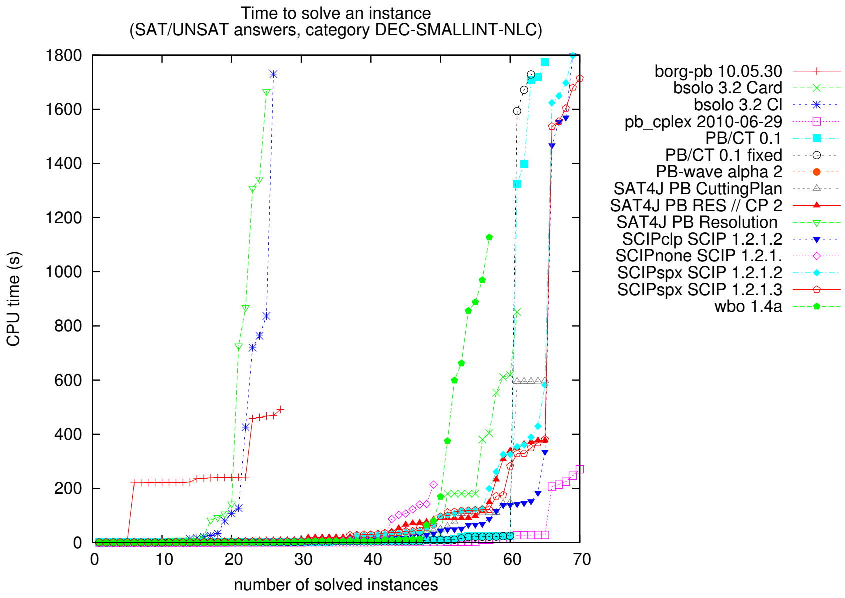

Results for OPT-SMALLINT-LIN

Rank	Solver	#solved	Detail	%inst.	%VBS
Total number of instances: 699					
Virtual Best Solver (VBS)		481	446 O, 35 U	69%	100%
1	pb_cplex	417	384 O, 33 U	60%	87%
2	SCIPspx *bugfix*	354	321 O, 33 U	51%	74%
3	bsolo *3.2 Card*	333	300 O, 33 U	48%	69%
4	bsolo *3.2 Cl*	328	295 O, 33 U	47%	68%
5	SCIPclp	319	286 O, 33 U	46%	66%
6	SCIPspx	317	284 O, 33 U	45%	66%
7	SAT4J Res//CP	315	282 O, 33 U	45%	65%
8	SAT4J Res.	303	270 O, 33 U	43%	63%
9	PB/CT *bugfix*	283	251 O, 32 U	40%	59%
10	SAT4J CP	255	226 O, 29 U	36%	53%
11	SCIPnone	187	158 O, 29 U	27%	39%

OPT-SMALLINT-LIN

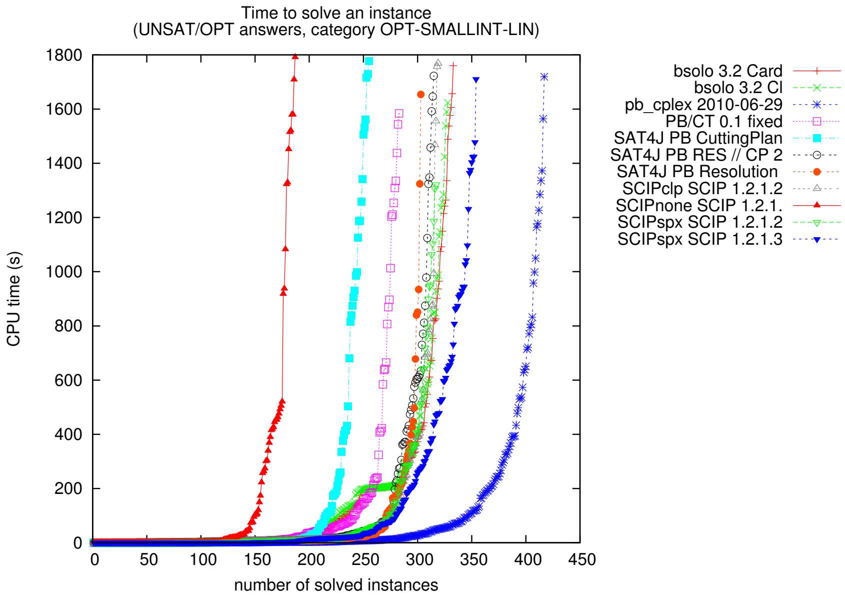

Results for OPT-SMALLINT-NLC

Rank	Solver	#solved	Detail	%inst.	%VBS
Total number of instances: 409					
Virtual Best Solver (VBS)		289	289 O	71%	100%
1	SCIPspx *bugfix*	288	288 O	70%	100%
2	SAT4J Res	271	271 O	66%	94%
3	SCIPnone	260	260 O	64%	90%
4	SAT4J Res//CP	250	250 O	61%	87%
5	bsolo *3.2 Cl*	230	230 O	56%	80%
6	bsolo *3.2 Card*	217	217 O	53%	75%
7	PB/CT	194	194 O	47%	67%
8	PB/CT *bugfix*	186	186 O	45%	64%
9	SAT4J CP	117	117 O	29%	40%

OPT-SMALLINT-NLC

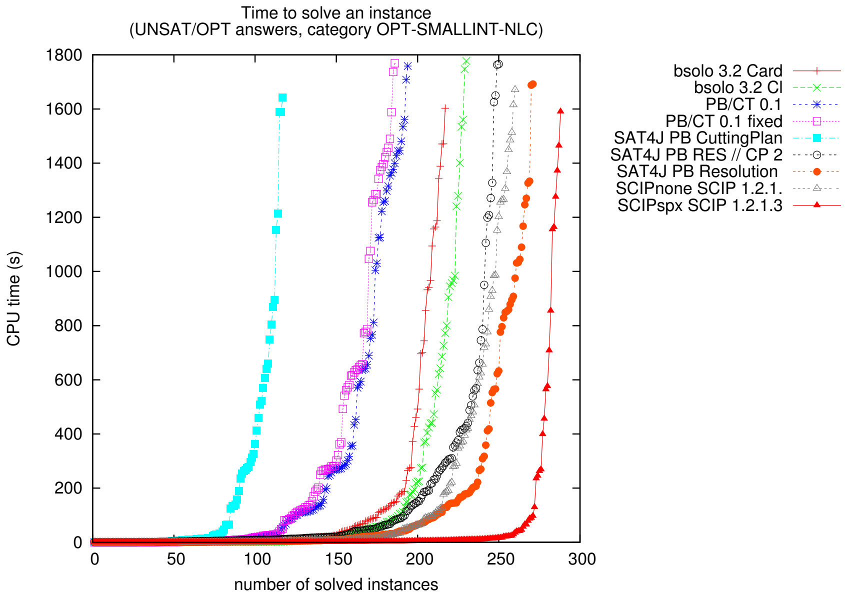

Results for OPT-BIGINT-LIN

Rank	Solver	#solved	Detail	%inst.	%VBS
Total number of instances: 532					
Virtual Best Solver (VBS)		211	152 O, 59 U	40%	100%
1	SAT4J Res//CP	205	146 O, 59 U	39%	97%
2	SAT4J Res	198	141 O, 57 U	37%	94%
3	SAT4J CP	168	110 O, 58 U	32%	80%
4	PB/CT *bugfix*	87	53 O, 34 U	16%	41%

OPT-BIGINT-LIN

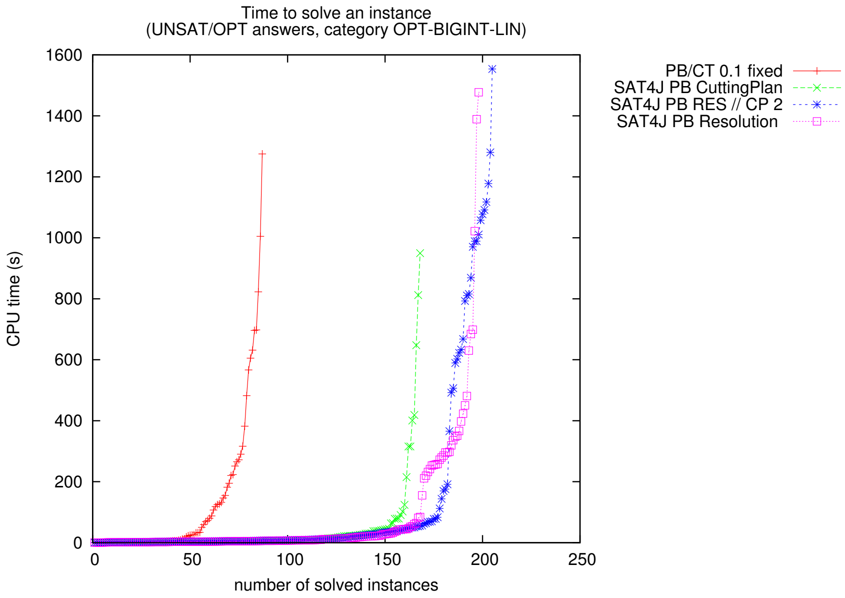

Results for PARTIAL-SMALLINT-LIN

Rank	Solver	#solved	Detail	%inst.	%VBS
Total number of instances: 536					
Virtual Best Solver (VBS)		532	531 O, 1 U	99%	100%
1	SAT4J Res. *bugfix*	446	445 O, 1 U	83%	84%
2	pb_cplex	428	428 O	80%	80%
3	PB/CT *bugfix*	375	374 O, 1 U	70%	70%
4	wbo *1.4a*	373	372 O, 1 U	70%	70%
5	SCIPclp	297	296 O, 1 U	55%	56%
6	SCIPclp	282	281 O, 1 U	53%	53%
7	SCIPspx	269	268 O, 1 U	50%	51%
8	SCIPnone	146	145 O, 1 U	27%	27%

PARTIAL-SMALLINT-LIN

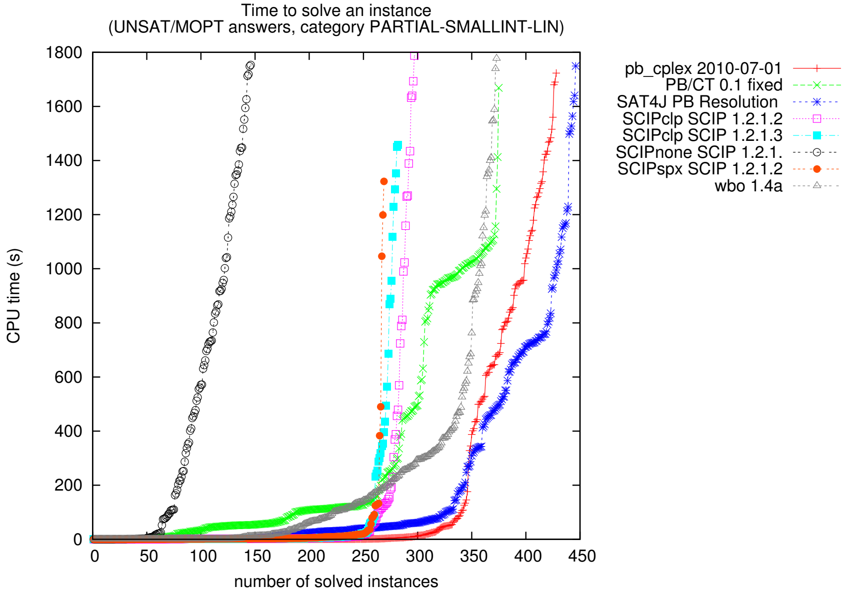

Results for SOFT-SMALLINT-LIN

Rank	Solver	#solved	Detail	%inst.	%VBS
Total number of instances: 201					
Virtual Best Solver (VBS)		201	201 O	100%	100%
1	wbo *1.4a*	161	161 O	80%	80%
2	pb_cplex	160	160 O	80%	80%
3	SAT4J Res. *bugfix*	160	160 O	80%	80%
4	PB/CT *bugfix*	138	138 O	69%	69%
5	SCIPclp	113	113 O	56%	56%
6	SCIPspx	113	113 O	56%	56%
7	SCIPclp	113	113 O	56%	56%
8	SCIPnone	22	22 O	11%	11%

SOFT-SMALLINT-LIN

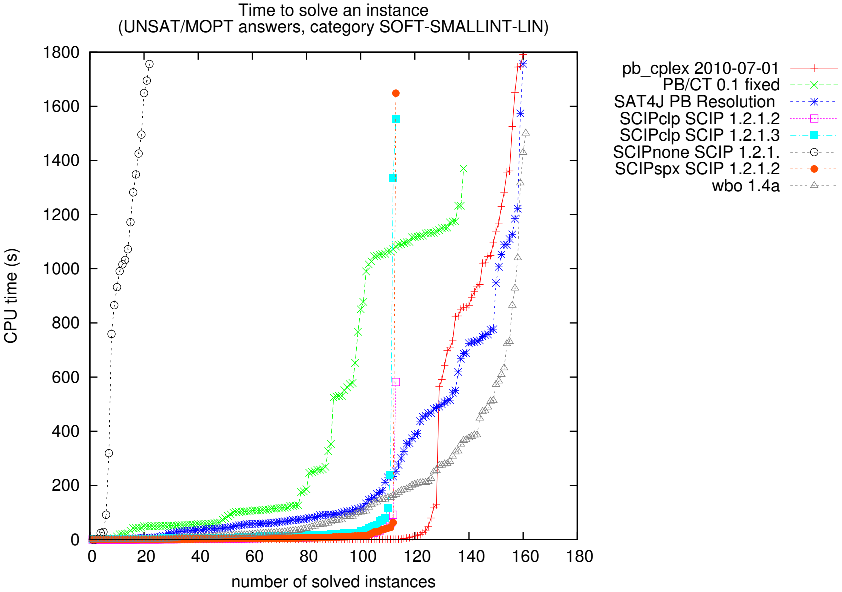

Results for PARTIAL-BIGINT-LIN

Rank	Solver	#solved	Detail	%inst.	%VBS
Total number of instances: 263					
Virtual Best Solver (VBS)		135	117 O, 18 U	51%	100%
1	SAT4J Res. *bugfix*	113	113 O	43%	84%
2	PB/CT *bugfix*	78	78 O	30%	58%

PARTIAL-BIGINT-LIN

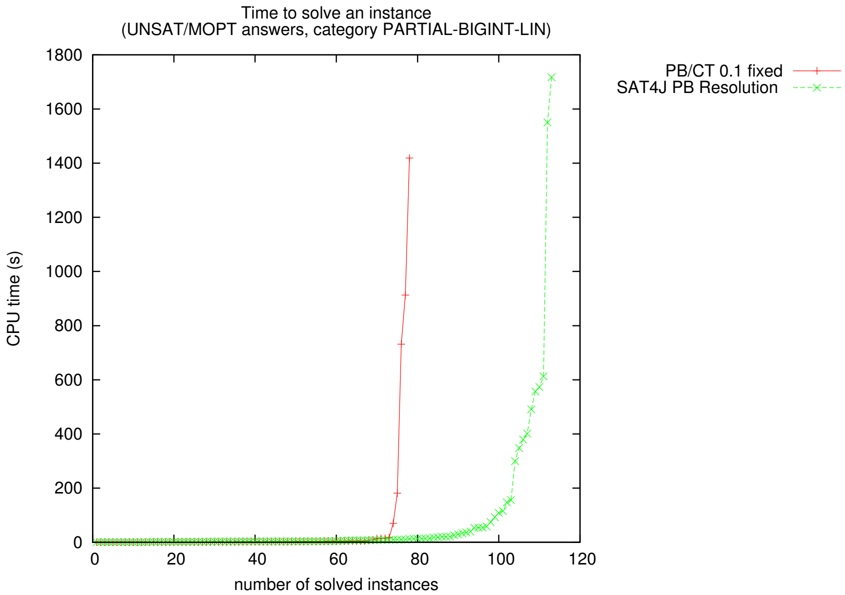

A word of warning...

- Keep in mind that the competition only takes a snapshot from a given angle.
- The rankings represent a user point of view, on a specific set of instances. This is only one small part of the picture.
- There are more points of view which are also relevant: innovation, robustness...

Some lessons

- A portfolio approach is valuable
- CPLEX outperforms all other solvers in OPT-SMALLINT-LIN and DEC-SMALLINT-NLC but is not so strong in other categories
- Linear programming techniques can help
- Learning PB constraints has a cost. Alternative approaches are valuable.

More information

- All details are on the web site http://www.cril.univ-artois.fr/PB10
- Get your solvers ready for PB11!
- Thanks to all participants!