

Fifth XCSP³ Competition
(2023 CSP and COP competition)
Second Call for Solvers and Benchmarks
(**Participation Deadline:** June 1, 2023)
<https://www.pycsp.org>

Changes with respect to First Call. The submission site <http://xcsp23.cril.fr/> is now open. The deadline is now June 1, 2023. When you upload a file (an archive containing instances or a solver), in the description field associated with your upload action, please write a line of the form:

- Type : XXX

where XXX is one value (or possibly several values) among:

- Benchmark
- CSP sequential
- COP fast (4 minutes)
- COP sequential
- COP parallel
- mini CSP
- mini COP

and indicate the line command to run your solver.

Of course, you can indicate any other relevant information about your submission.

The fifth international XCSP³ constraint solver competition is organized to improve our knowledge about components (e.g., filtering algorithms, heuristics, search strategies, encodings, reformulation techniques and learning procedures) that are behind the efficiency of solving systems (referred to as constraint solvers in this document) for combinatorial constrained problems. Two classical problems are considered for this competition:

- CSP (Constraint Satisfaction Problem)
- COP (Constrained Optimization Problem)

The intermediate¹ format XCSP³ is used as input format for the solvers. The effort required for entering the competition is limited because some tools (parsers) are available, and only a central set of popular (and important) constraints is considered.

This call for solvers and benchmarks presents the tracks that will be considered during the competition. In particular, we give important details about the format restrictions, the execution environment, and the rules that must be followed by the solvers.

Importantly, do note that:

- detailed results concerning the previous XCSP³ competitions can be found at <https://www.xcsp.org/competitions>.
- anybody (and in particular, contestants) can submit new benchmarks. You can use the modeling library called [PyCSP³](#) [LS22] for that.
- the proceedings of the 2022 competition are available [ALL22].

¹XCSP³ is neither a modeling language, nor a flat format. It is intermediate because it preserves the structure of problems through the concepts of variable arrays, constraint groups/blocks, and meta-constraints.

Contents

1	Main Points of the 2023 Competition	4
2	Timetable	4
3	Tracks	4
4	Modifications concerning XCSP³ with respect to the 2022 Competition	5
5	Format	6
5.1	Standard Tracks	6
5.1.1	Constraint precedence	7
5.1.2	Constraint knapsack	7
5.1.3	Constraint binPacking	7
5.2	Mini-solver Tracks	7
6	Resources: Benchmarks and Tools	8
7	Execution Environment	8
7.1	Command Line	8
7.2	Output Format	10
7.3	Special Considerations for Incomplete Solvers	13
7.3.1	Complete solvers	13
7.3.2	Incomplete solvers	14
7.4	Special Considerations for Parallel Solvers	14
8	Entering the Competition	15
9	Ranking	15
10	Organization	16

1 Main Points of the 2023 Competition

The organization of the 2023 competition will be rather close to the one conducted in 2022. There are however a few minor changes that are indicated below:

- The perimeter (syntax, constraints) has been slightly extended, mainly introducing the constraints `binPacking`, `knapsack` and `precedence` for the main tracks. This is discussed in Section 4.
- Two new rules will be used if necessary:
 - In case a team submits the same solver to both the main track and the mini-track of the same problem (CSP or COP), the solver will be ranked in the mini-track only if the solver is not one of the three best solvers in the main track.
 - In case several teams submit variations of the same solver to the same track, only the team who developed the solver and the best other team with that solver will be ranked (possibly, a second best other team, if the jury thinks that it is relevant).

2 Timetable

The deadlines of the competition are defined below:

Opening of the registration site at http://xcsp23.cril.fr/	April 2023
Pre-registration of contestants	May 1, 2023
Final registration (submission of solvers and benchmarks)	June 1, 2023
Test of solvers conformance	June 2023
Position paper (2 pages)	end-June 2023
Competition running	June-July 2023
Final results available	during CP 2023

Once submitted, solvers will be run on a limited number of benchmarks to make sure that they interact correctly with the evaluation environment. Potential problems will be reported to the authors so that they have some opportunities to fix bugs.

3 Tracks

Tracks remain the same as in 2022. First, let us recall that we consider two main problems: CSP (a decision problem) and COP (an optimization problem). For CSP, the goal is to exhibit one solution or to prove that none exists. For COP (mono-objective optimization), the goal is to exhibit a solution with the best possible objective value, ideally proving that it represents an optimum solution (note that some COP instances in the competition may be unsatisfiable, which must then be proved).

Anyone can submit a solver to any particular track. There are 4 main standard tracks, imposing absolutely no conditions on solvers. For example, they can be written in any language (provided that we can reasonably execute them in our environment; see section 7), and can

Problem	Goal	Exploration	Timeout
CSP	one solution	sequential	40 minutes
COP	best solution	sequential	4 minutes
COP	best solution	sequential	40 minutes
COP	best solution	parallel	40 minutes

Table 1: Standard Tracks.

be complete or incomplete solvers (e.g., based on local search). The tracks are described in Table 1.

Ranking for COP will be stated in two different manners: by considering and not considering possible proofs of optimality, permitting in the latter case to emphasize the quality of incomplete solvers.

There are also 2 Mini-Solver tracks. The goal of these tracks is to facilitate the submission of new systems as well as to allow the discovery of new simple ideas. Basically, solvers in these tracks will be evaluated on a restricted set of constraints and may be submitted as a patch of an existing, open-source solver. See Section 5.2 for details.

Problem	Goal	Exploration	Timeout
CSP	one solution	sequential	40 minutes
COP	best solution	sequential	40 minutes

Table 2: Mini-Solver Tracks.

4 Modifications concerning XCSP³ with respect to the 2022 Competition

Compared to 2022, here are some extensions with respect to the syntax:

1. three constraints are introduced: `binPacking`, `knapsack` and `precedence`. **Important:** these constraints can be easily decomposed.
2. Variables (instead of integers) can be used in the element `<coeffs>` of an objective. However, although authorized in XCSP³ Specifications 3.1, for the 2023 competition, general integer expressions will not be permitted in the element `<coeffs>` of a constraint `sum` or in the element `<coeffs>` of an objective.

Important: a set of instances is available in order to let the competitor to check all these changes/extensions:

<https://www.cril.univ-artois.fr/~lecoutre/compets/checkingInstancesXCSP23.zip>

5 Format

The complete description of the format (XCSP³) used to represent combinatorial constrained problems can be found in [BLAP22]. Do note that we refer to the **version 3.1** of the specifications. However, as in 2022, for the 2023 competition, we limit XCSP³ to its kernel, called XCSP³-core, whose description is given in [BLA22]. This means that the scope of XCSP³ is restricted to:

- integer variables,
- CSP and COP problems,
- a set of 24 popular (global) constraints for Standard tracks, and a small set of constraints for Mini-solver tracks.

For simplicity, as in 2022, we also impose the following restrictions:

- Integer variables have finite domains (and so, the special value *infinity* is forbidden).
- Variable arrays always start indexing at 0 (and so, the attribute `startIndex`, whose default value is 0, cannot be associated with `<array>`).
- The attribute `as` can only be associated with elements `<var>` and `<array>`; see Section 10.5 in [BLAP22].
- Undefined variables are not accepted but useless variables are (note that parsers/solvers can easily identify useless variables); see Section 2.10 in [BLAP22].
- Advanced forms of constraints (see Part III in [BLAP22]) are not accepted, except for the very specific cases explicitly described in the rest of this section.
- View extensions are accepted for some constraints: `allDifferent` and `sum`, as in 2019, but also `allEqual`, `count`, `nValues`, `minimum`, and `maximum`, as in 2022.
- The type of the objective (in case of a COP instance) cannot be "product" or "lex".
- Any integer value occurring in an XCSP³ file must belong to the interval $-2^{31}..2^{31} - 1$.

5.1 Standard Tracks

In the Standard tracks, we find twenty-four constraints. Do note that a large majority of the numerous instances that are currently available on our website <http://www.xcsp.org/instances/> only involve these 24 constraints.

Compared to 2022, there are no changes concerning the 21 constraints already involved in 2022, which are: `intension`, `extension`, `regular`, `mdd`, `allDifferent`, `allEqual`, `ordered`, `lex`, `sum`, `count`, `nValues`, `cardinality`, `maximum`, `minimum`, `element`, `channel`, `noOverlap`, `cumulative`, `instantiation`, `circuit`, `slide`. For more information, we refer the reader to Section 5.1 of the previous call :

<https://www.cril.fr/~lecoutre/compets/callXCSP22.pdf>

But three new constraints are introduced. They are described below.

5.1.1 Constraint precedence

This constraint is described in Section 4.1.3.5 in [BLAP22]. Competition restrictions:

- the attribute `covered` will not be used (and so, never present in instances used for the 2023 competition)

5.1.2 Constraint knapsack

This constraint is described in Section 4.1.6.5 in [BLAP22]. Competition restrictions:

- none

This constraint can be very easily decomposed.

5.1.3 Constraint binPacking

This constraint is described in Section 4.1.6.4 in [BLAP22]. Competition restrictions:

- the general form will not be used (only the first three forms as described in Section 4.1.6.4)

This constraint can be easily decomposed.

5.2 Mini-solver Tracks

As previously stated, the goal of the mini-solver tracks is to allow to enter the competition without having to implement all the features required in the standard tracks.

The requirements to enter the mini-solver tracks are the following:

- solvers will be evaluated on a restricted set of constraints (see below),
- solvers will be only evaluated on sequential search,
- solvers may be submitted as a modification of an existing, open-source solver. The goal in this case is to save developers the effort of developing a complete solver and to focus on new techniques, heuristics and so on. In this case, the submission must contain both the source of the initial solver, and the source of the modified solver (or equivalently a patch file). The differences between the initial solver and the modified solver should be significant enough. As an example, Mini-CP (www.minicp.org) may be used as a starting point.
- obviously, solvers developed independently by their authors may be submitted as well.

Compared to 2022, there are no changes concerning the perimeter (constraints) of the mini-tracks. As in 2022, the constraints that are accepted for Mini-solver tracks are restricted to five types. For more information, we refer the reader to Section 5.2 of the previous call :

<https://www.cril.fr/~lecoutre/compets/callXCSP22.pdf>

6 Resources: Benchmarks and Tools

Many benchmarks can be found at <https://www.xcsp.org/instances>.

The organizers invite *anybody* to submit new benchmarks. The organizers are particularly interested in new problem instances originating from real-world applications. For generating new XCSP³ instances, one can use the Python library PyCSP³: see <https://pycsp.org/>.

Some tools are also provided. They can be found at www.xcsp.org/tools.

Currently, you can find:

- a C++ parser
- a Java parser
- a tool for checking solutions and costs
- a Java-based modeling API

7 Execution Environment

Solvers will run by the Slurm cluster management system on a cluster of computers executing the CentOS Stream 8.3 operating system. They will run under the control of another program (called runsolver) that will enforce some limits on both used memory and total CPU time.

Solvers can be run as either 32 bits or 64 bits applications. If you submit an executable, you are required to provide us with an ELF executable (preferably statically linked). Authors have to provide through the submission site the a document that enumerates the names and versions of the tools, libraries, ... that are needed to build and execute the project. The procedure to build and execute the solver must also be explained in this document.

Keep in mind that solvers will all be run on the same user account; softwares that alter the user environment (eg. Anaconda) must set it back to its initial state after execution, whatever the result of the execution (including a crash). It is advised to replace such tools by counterparts which do not alter the environment (eg. venv instead of Anaconda).

Two executions of a solver with the same parameters and system resources are expected to output the same result in approximately the same time (so that the experiments can be repeated).

7.1 Command Line

During the submission process, you will be asked to provide the organizers with a suggested command line that should be used to run your solver. In this command line, you will be asked to use the following placeholders, which will be replaced by the actual information by the evaluation environment.

- BENCHNAME will be replaced by the name of the file containing the XCSP³ instance to solve. Obviously, the solver must use this parameter or one of the following variants: BENCHNAMENOEXT (name of the file with path but without extension), BENCHNAMENOPATH (name of the file without path but with extension), BENCHNAMENOPATHNOEXT (name of the file without path nor extension).

- RANDOMSEED will be replaced by a random seed which is a number between 0 and 4294967295. This parameter MUST be used to initialize the random number generator when the solver uses random numbers. It is recorded by the evaluation environment and will allow to run the program on a given instance under the same conditions if necessary.
- TIMELIMIT (or TIMEOUT) represents the total CPU time (in seconds) that the solver may use before being killed. May be used to adapt the solver strategy.
- MEMLIMIT represents the total amount of memory (in MiB) that the solver may use before being killed. May be used to adapt the solver strategy.
- NBCORE will be replaced by the number of processing units that have been allocated to the solver. Note that, depending on the available hardware, a processing unit may be either a processor, a core of a processor or a “logical processor” (in hyper-threading).
- TMPDIR is the name of the only directory where the solver is allowed to read/write temporary files
- DIR is the name of the directory where the solver files will be stored

Examples of command lines:

```
DIR/mysolver BENCHMARK RANDOMSEED
DIR/mysolver --mem-limit=MEMLIMIT --time-limit=TIMELIMIT --tmpdir=TMPDIR BENCHMARK
java -jar DIR/mysolver.jar -c DIR/mysolver.conf BENCHMARK
```

As an example, these command lines could be expanded by the evaluation environment as:

```
/solver10/mysolver /tmp/zebra.xml 1720968
/solver10/mysolver --mem-limit=900 --time-limit=1200 --tmpdir=/tmp/job12345 /tmp/zebra.xml
java -jar /solver10/mysolver.jar -c /solver10/mysolver.conf /tmp/zebra.xml
```

The command line provided by the submitter is only a suggested command line. Organizers may have to modify this command line (e.g., memory limits of the Java Virtual Machine (JVM) may have to be modified to cope with the actual memory limits).

The solver may also (optionally) use the values of the following environment variables:

- TIMELIMIT (or TIMEOUT) (the number of seconds it will be allowed to run)
- MEMLIMIT (the amount of RAM in MiB available to the solver)
- TMPDIR (the absolute pathname of the only directory where the solver is allowed to create temporary files)

After TIMEOUT seconds have elapsed, the solver will first receive a SIGTERM to give it a chance to output the best solution it found so far (in the case of an optimization problem). One second later, the program will receive a SIGKILL signal from the controlling program to terminate the solver.

The solver cannot write to any file except standard output, standard error and files in the TMPDIR directory. A solver is not allowed to open any network connection or launch unexpected external commands. Solvers may use several processes or threads. Children of a solver process are allowed to communicate through any convenient means (Pipes, Unix or Internet sockets, IPC, ...). Any other communication is strictly forbidden. Solvers are not allowed to perform actions that are not directly related to the resolution of the problem.

7.2 Output Format

To communicate their answers, solvers must print messages to the standard output and those messages will be used to check the results. The first two characters of a line allow us to classify it into different categories, which indicate the meaning of the line. With the exception of "o " lines, there is no specific order imposed on the lines output by solvers.

- **status line**

This line starts by the two characters: lower case s followed by a space (ASCII code 32). Only one such line is allowed, and it is mandatory. This line gives the answer of the solver. It must be one of the following answers:

- **s UNSUPPORTED**

This line should be printed by the solver when it discovers that the XCSP³ instance contains a non-supported feature. As an example, a solver that cannot deal with a global constraint should print this line when such a constraint is present.

- **s SATISFIABLE**

This line indicates that the solver has found a solution, and in such a case, a "v " line (see below) is mandatory. For CSP, the solver answers SATISFIABLE when it has found a solution. For COP, the solver answers SATISFIABLE when it has found a solution that it couldn't prove to be optimal.

- **s OPTIMUM FOUND**

This line must be printed when the solver has found an optimal solution for a COP instance, and in such a case, a "v " line (see below) is mandatory. This answer implies that the solver has proved that no other solution can give a better value of the objective function. This answer must not be used for CSP instances.

- **s UNSATISFIABLE**

This line must be output when the solver can prove that the instance has no solution.

- **s UNKNOWN**

This line may be output in any other case, i.e. when the solver is not able to tell anything about the instance.

It is of uttermost importance to respect the exact spelling of these answers. Any mistake in the writing of these lines will cause the answer to be disregarded.

Solvers are not required to provide any specific exit code corresponding to their answer.

If the solver does not output a status line, or if the status line is misspelled, then UNKNOWN will be assumed.

- **values line**

This line starts by the two characters: lower case v followed by a space (ASCII code 32). It is mandatory when the instance is satisfiable. More than one "v " line is allowed but the evaluation environment will act as if their content was merged.

If the solver finds a solution (i.e., if the solver outputs "s SATISFIABLE" or "s OPTIMUM FOUND"), it must provide a solution. For CSP or COP, this solution is an instantiation that satisfies every constraint. For COP, this instantiation must be such

that the value of the objective function corresponds to the best one that the solver was able to find.

Solutions must respect the format described in Section 2.11 of [BLAP22]. However, it is important to note that the attributes `type` and `cost` that can be associated with the element `<instantiation>` are **not required** in the context of the competition. These attributes, if present, will simply be ignored.

Importantly, the solution can be output on several successive "v " lines, provided that each "v " line must be terminated by a Line Feed character (the usual Unix line terminator `'\n'`). A "v " line that does not end with that terminator will be ignored because it will be considered that the solver was interrupted before it could print a complete solution.

As an illustration, the following output is valid for the COP instance (Example 4) given in Chapter 1 of [BLAP22]:

```
v <instantiation type="optimum" cost="1700">
v   <list> b c </list>
v   <values> 2 2 </values>
v </instantiation>
```

and the following output is valid for the CSP instance (Example 25) given in Section 2.11 of [BLAP22]:

```
v <instantiation type="solution">
v   <list> x[] </list>
v   <values> 1 1 2 * </values>
v </instantiation>
```

As the attributes `type` and `cost` are not required (and simply ignored by our environment), we could have written:

```
v <instantiation>
v   <list> b c </list>
v   <values> 2 2 </values>
v </instantiation>
```

and

```
v <instantiation>
v   <list> x[] </list>
v   <values> 1 1 2 * </values>
v </instantiation>
```

- **objective line**

These lines start by the two characters: lower case o followed by a space (ASCII code 32). **These lines are mandatory for incomplete solvers.** As far as complete solvers are concerned, they are not strictly mandatory but solvers are strongly invited to print them. These lines are only relevant for COP instances.

Whenever the solver finds a solution with a better value of the objective function, it is asked to print an "o " line with the current value of the objective function. Therefore, an "o " line must contain the lower case o followed by a space and then by an integer

that represents the better value of the objective function. "o " lines should be output as soon as the solver finds a better solution and be ended by a standard Unix end of line character ('\n'). Programmers are advised to flush immediately the output stream.

As an example; let us consider Example 2 in Chapter 1 of [BLAP22]. Let us assume that the solver finds first this solution:

```
<instantiation id='sol1' type='solution' cost='450'>
  <list> b c </list>
  <values> 0 1 </values>
</instantiation>
```

and later:

```
<instantiation id='sol2' type='solution' cost='1700'>
  <list> b c </list>
  <values> 2 2 </values>
</instantiation>
```

which is finally proved to be optimal by the solver. The output by the solver can be (using this time only one "v " line):

```
o 450
o 1700
s OPTIMUM FOUND
v <instantiation> <list> b c </list> <values> 2 2 </values> </instantiation>
```

The evaluation environment will automatically timestamp each of these lines so that it is possible to know when the solver has found a better solution and how good the solution was. The goal is to be able to analyze the way solvers progress toward the best solution. As an illustration, here is a sample of the output of a solver, with each line timestamped (first column, expressed in seconds of wall clock time since the start of the program):

```
0.00      c Time Limit set via TIMEOUT to 1800
0.51      c Initial problem consists of 6774 variables and 100 constraints.
0.55      c preprocess terminated. Elapsed time: 0.45
0.55      c Initial Lower Bound: 0
0.63      o 235947
0.63      o 226466
0.63      o 217758
0.75      o 186498
1.16      o 178319
2.42      o 168389
3.13      c Restart #1 #Var: 6774 LB: 0 @ 3.03
4.89      c Restart #2 #Var: 6774 LB: 0 @ 4.79
5.73      o 160358
6.44      o 159206
7.52      o 150077
9.09      o 149533
12.14     o 140853
17.74     o 140264
19.61     o 131636
29.81     o 15450
34.00     o 7066
41.66     o 5000
84.01     o 3905
```

```
84.01    c NEW SOLUTION FOUND: 3905 @ 83.873
84.61    s OPTIMUM FOUND
84.61    v ... // solution not shown here for space reasons
84.61    c Total time: 84.478 s
```

- **diagnostic line**

These lines are optional and start with the two following characters: lower case `d` followed by a space (ASCII code 32). Then, a keyword followed by a value must be given on this line.

More precisely, a *diagnostic* is a (name,value) pair that gives an information about the work carried out by the solver. As indicated above, each diagnostic is a line of the form `'d NAME value'`, where `NAME` is a sequence of letters describing the diagnostic, and `value` is a sequence of characters defining its value. The following diagnostic is predefined:

WRONG DECISIONS: The total number of wrong decisions which have been carried out (as defined in [BZF04]).

Contestants wishing to record other diagnostics than the one listed before above should inform the organizers.

- **comments line**

A line which is not one the special lines defined above, or which explicitly starts with the two characters: lower case `c` followed by a space (ASCII code 32) is a comment line, and is ignored. These lines are thus optional and may appear anywhere in the solver output.

They contain any information that authors want to emphasize, such as `#backtracks`, `#flips`,... or internal CPU time. They are recorded by the evaluation environment for later viewing but are otherwise ignored. At most one megabyte of solver output will be recorded. So, if a solver is very verbose, some comments may be lost.

Submitters are advised to avoid printing comment lines which may be useful in an interactive environment but otherwise useless in a batch environment. For example, printing comment lines with the number of constraints read so far only increases the size of the logs with no benefit.

If a solver is really too verbose, the organizers will ask the submitter to remove some comment lines.

7.3 Special Considerations for Incomplete Solvers

Complete solvers are solvers which can always decide the satisfiability of a CSP instance and the optimality of a COP instance, provided that enough time and memory are given. Incomplete solvers may loop endlessly in a number of cases; local search algorithms are examples of incomplete solvers. Both kinds of solvers are welcome in this competition. Submitters will have to indicate if their solver is complete or incomplete on the submission form.

7.3.1 Complete solvers

There is no special requirement about complete solvers. See the input and output format that all solvers must respect for details.

7.3.2 Incomplete solvers

Incomplete solvers are definitely welcome in the competition.

For CSP, an incomplete solver will stop as soon as it finds a solution and will time out if it can't find one. The only difference with a complete solver is that it will time out systematically on unsatisfiable instances.

For COP, an incomplete solver will systematically time out because it will be unable to prove that it has found the optimum solution. Yet, it may have found the optimum value well before the time out. In order to get relevant information in these categories, an incomplete solver must fulfill two requirements:

1. it must intercept the SIGTERM signal sent to the solver on timeout and output either "s UNKNOWN" or "s SATISFIABLE" with the "v " line(s) corresponding to the best solution it has found
2. it MUST output an "o " line whenever it finds a better solution so that, even if the solver always timeouts, the timestamp of the last "o " line indicates when the best solution was found. Keep in mind that it is the evaluation environment which is in charge of timestamping "o " lines.

7.4 Special Considerations for Parallel Solvers

The execution environment will bind the solvers to a subset of all available processing units. The environment variable NBCORE will indicate how many processing units have been granted to the solver. The solver will not have access to more processing units than NBCORE. This implies that if the solver uses x threads or processes (with $x > \text{NBCORE}$), $x - \text{NBCORE}$ threads or processes will necessarily sleep at one time. Note that only scenarios $\text{NBCORE}=8$ and $\text{NBCORE}=4$ will be possible this year; we will not share processors between two runs to prevent them from stepping on each other too much (processor cache, access memory, ...).

As an example, if the competition is run on hosts with 2 quad-core processors (8 cores in total), several scenarios are possible:

- one single solver is run on the host, it is allowed to use all 8 cores ($\text{NBCORE}=8$).
- two solvers are run simultaneously, each one being assigned to a given processor (which means that a solver is assigned 4 cores, hence $\text{NBCORE}=4$).
- 4 solvers are run simultaneously, each one being assigned to a fixed set of 2 cores (belonging to the same CPU), hence $\text{NBCORE}=2$.
- more generally, a single solver may be assigned any number x of cores (from 1 to 8 in this example) to simulate the availability of x processing units.

The solver might use the NBCORE environment variable to adapt itself to the number of available processing units.

A solver must not modify its processor affinity (calls to `sched_setaffinity(2)` or `taskset(1)`) to get access to a processing unit that was not initially allocated to the solver. It may however modify its processor affinity to use a subset of the initially allocated processing units.

8 Entering the Competition

Contestants can enter the competition with one or two solvers per track. Contestants are expected to submit their solver(s) and contribute some instances (as many instances as wished). Submitted instances will be made available on the evaluation web site shortly after the actual beginning of the competition. We cannot accept benchmarks which cannot (for various reasons) be publicly available (because anyone must be able to reproduce the experiments of the competition). In a second stage, they will also have to submit a position paper (1 page, or preferably 2 pages, or even more) indicating the main components of the submitted solver(s).

Of course, we expect that contestants propose solvers that recognize XCSP³ (either natively or by embedding a conversion procedure).

The deadline for submitting both benchmarks and solvers is June 1, 2023. Submission of solvers and benchmarks will be possible online in May 2023 at <http://xcsp23.cril.fr/>.

9 Ranking

About Scoring. The number of points won by a solver S is decided as follows:

- for CSP, this is the number of times S is able to solve an instance, i.e., to decide the satisfiability of an instance (either exhibiting a solution, or indicating that the instance is unsatisfiable)
- for COP, this is, roughly speaking, the number of times S gives the best known result, compared to its competitors. More specifically, for each instance I :
 - if I is unsatisfiable, 1 point is won by S if S indicates that the instance I is unsatisfiable, 0 otherwise,
 - if S provides a solution whose bound is less good than another one (found by another competing solver), 0 point is won by S ,
 - if S provides an optimal solution, while indicating that it is indeed the optimality, 1 point is won by S ,
 - if S provides (a solution with) the best found bound among all competitors, this being possibly shared by some other solver(s), while indicating no information about optimality: 1 point is won by S if no other solver proved that this bound was optimal, 0.5 otherwise.

Other ranking schemes may be introduced to help identify remarkable features.

Wrong Answers. Note that a solver is declared to give a wrong answer in the following cases:

- It outputs UNSATISFIABLE for an instance which can be proved to be satisfiable.
- For CSP and COP, it outputs SATISFIABLE or OPTIMUM FOUND, but provides an instantiation that does not satisfy every constraint. The only exception is when the solver outputs an incomplete "v " line (which does not end by '\n') in which case it is assumed that the solver was interrupted before it could output the complete model and the answer will be considered as UNKNOWN.

- It outputs OPTIMUM FOUND but there exists an instantiation with a better value of the objective function/cost than the one corresponding to the printed solution.

When a solver provides a wrong answer in a given track, the solver's results in that track will be excluded from the final evaluation results because they cannot be trusted.

A solver that ends without giving any solution, or just crashes for some reason (internal bugs...), is simply considered as giving an UNKNOWN result.

10 Organization

Gilles Audemard, Christophe Lecoutre and Emmanuel Lonca from CRIL.

They can be reached at lecoutre@cril.fr and lonca@cril.fr.

References

- [ALL22] G. Audemard, C. Lecoutre, and E. Lonca. *Proceedings of the 2022 XCSP3 Competition*. Technical Report. on CoRR, [arXiv:2209.00917](https://arxiv.org/abs/2209.00917), 2022. 97 pages.
- [BLA22] F. Boussemart, C. Lecoutre, and G. Audemard. *XCSP3-core: A Format for Representing Constraint Satisfaction/Optimization Problems*. Technical Report. v3.1 on CoRR, [arXiv:2009.00514](https://arxiv.org/abs/2009.00514), 2020–2022. 105 pages.
- [BLAP22] F. Boussemart, C. Lecoutre, G. Audemard, and C. Piette. *XCSP3: An Integrated Format for Benchmarking Combinatorial Constrained Problems*. Technical Report. v3.1 on CoRR, [arXiv:1611.03398](https://arxiv.org/abs/1611.03398), 2016–2022. 241 pages.
- [BZF04] C. Bessiere, B. Zanuttini, and C. Fernandez. Measuring search trees. In *Proceedings of ECAI'04 workshop on Modelling and Solving Problems with Constraints*, pages 31–40, 2004.
- [LS22] C. Lecoutre and N. Szczepanski. *PyCSP3: Modeling Combinatorial Constrained Problems in Python*. Technical Report. v2.1 on CoRR, [arXiv:2009.00326](https://arxiv.org/abs/2009.00326), 2020–2022. 155 pages.