The VeRoLog Solver Challenge 2016-2017

The Set-up of the challenge

Overview

The VeRoLog Solver Challenge 2016-2017 will consist of two parts that run parallel in time:

1. An all-time-best challenge. The organizers will disclose 25 instances (dataset ORTEC-early) in October 2016. The participants are invited to submit their solution to an instance if it is better than the best solution submitted so far for this instance. Progress on the instances will be shown on `https://verolog.ortec.com`. The all-time-best challenge will run till May 1, 2017. In this part you are allowed to use any means you want.

2. A restricted resources challenge, this is a challenge in the more "traditional" form. Here the resources are restricted, especially time. To participate you have to run your algorithm on a set of instances (dataset ORTEC-late) and send in your results and solver binaries on May 1, 2017. Based on these results, around 5 finalists will be selected. The solvers of the finalists are then applied to a set of hidden instances. Per instance the participants are ranked and the participant with the lowest mean rank is the winner of the challenge.

Initially all instances and solutions will be published on `https://verolog.ortec.com` and later in the VRP repository.

In the next section a more detailed description of the challenge is given, formalized by rules.

Rules

Please pay attention to the following rules.

General rules

Rule 1.1

This challenge seeks to encourage research into vehicle routing methods, and to offer prizes to the most successful methods. It is the spirit of these rules that is important, not the letter. With any set of rules for any challenge it is possible to work within the letter of the rules but outside the spirit. The organizers ask that you please don't do this.

Rule 1.2

The organizers reserve the right to remove a participant from the challenge if the participant is determined by the organizers to have worked outside the spirit of the challenge rules. The organizers' decision is final in any matter. Decisions will be made by the organizers with the chair of VeRoLog having the casting vote in the case of doubt.

Rule 1.3

The rules explained below might be updated in due course. Any change of rules will be accompanied by a general email to all registered participants.

Rule 1.4

The challenge consists of two parts, both with deadline around May 1, 2017; the precise date and time will be published on `https://verolog.ortec.com` . The first part is called the *all-time-best challenge* and the second part the *restricted resources challenge*. Participants can take part in one part or both parts. The participants of the restricted resources challenge also have to submit their binaries, that should work under Windows.

Rule 1.5

The goal is to produce feasible routes minimizing the cost.

The All-time-best challenge

Rule 2.1

The first part of the challenge consists of constructing all-time-best solutions to the 25 instances in *ORTEC-early*. For this part of the challenge there are no time or technology restrictions. Submissions of (new best) solutions should be done via `https://verolog.ortec.com` .

Rule 2.2

Per instance in ORTEC-early the best submitted solution at the deadline is rewarded. In case of ties the time that the solution was received is decisive.

The restricted resources challenge

Rule 3.1

The restricted resources challenge will work with the 25 instances in *ORTEC-late* and with hidden instances in *ORTEC-hidden*. The hidden instances will be 'similar' to the late instances. The ORTEC-late instances will be made available around February 2017, and the ORTEC-hidden instances after the VeRoLog-2017 conference.

Rule 3.2

In the restricted resources challenge the participants have to implement a single threaded algorithm to tackle the problem; they can use any programming language. Third party software can be used if this is freely available for all, even for commercial use. This *excludes* the use of, for example, CPLEX, ILOG, Gurobi and SCIP.

Rule 3.3

The algorithm can be either deterministic or stochastic. The participants that use a stochastic algorithm should code their program in such a way that the exact run can be reproduced by specifying a random seed.

Rule 3.4

In the restricted resources challenge the time limit for an instance is 10 seconds plus 2 seconds per (delivery) request on an imaginary default machine. Participants have to benchmark their machines with a benchmark program that will be provided by the organizers in order to know how much time they have available to run their programs on their machines. For example, the solving time for an instance with 1000 requests is 2010 seconds on the default machine. The benchmark program returned the factor 0.8740 for your machine, hence your program is allowed to run 1756.74 seconds on your machine. This is rounded to the nearest integer, 1757 seconds. (The variability in the benchmark program does not justify to work in milliseconds.)

Rule 3.5

The algorithm should take as input an instance file in text or XML format, and produce as output a solution in the same format as well, all within the allowed time.

The solver should run under Windows with command-line arguments: input file name, output file name, the calculation time in seconds (integer), and for stochastic solvers, the random seed. For example (for a stochastic solver):

```
>> my_solver.exe instance1.txt solution1.txt 1757 1542955064
```

Rule 3.6

When submitting solutions to the late instances, the same version of the algorithm must be used for all instances. That is, the algorithm should not "know" which of the 25 instances it is solving - while your particular algorithm might analyze the problem instance and set parameters accordingly, it should not "recognize" the particular instance. See also Rule 1.1.

Rule 3.7

The organizers will provide a set of 9 random seeds two weeks before the deadline. Random seed number $n = 1, 2, \ldots, 9$ will lie between $100,000,000n$ and $100,000,000(n+1)$. Participants should submit for each instance the costs and solutions found by their algorithm in the time set according to Rule 3.5, using the 9 random seeds. Even in case of a deterministic solver 9 runs could be executed to avoid the effects of breaking off the execution abruptly.

Rule 3.8

We calculate a rank score per instance for each solver. First, per instance, we remove the two best solutions and the two worst solutions found by the solver. Hence, we are left with the middle 5 solutions on which we base the score of the solver. If these five solutions are all feasible, we take their average as score of the solver, and rank all solvers accordingly. Solvers that produce infeasible solutions among the middle 5 solutions are first ranked with respect to the *number* of feasible solutions, and secondary by the average cost of the feasible ones. The organizers will select around 5 participants as finalists for the second part of the challenge, based on the mean of the ranks of the scores of the solver. For this the results of potential finalists will be verified on the instances in ORTEC-late with the given random seeds, and with other random seeds, if necessary.

Rule 3.9

The finalists' solvers will be run by the organizers on the instances in ORTEC-hidden again with 9 random seeds. The same executables which were delivered for validation will be used in the final. With the same rules as above the score per instance will be calculated as well as the rank of the participant based on these scores. The mean ranks will produce the final place list.

Ranking

Here we explain how the ordering works. Consider the following example with the following scores (for ease assumed integer values) of 7 participants on 6 instances.

Instance	1	2	3	4	5	6
Solver A:	34	35	142	132	10	12
Solver B:	332	124	144	133	13	15
Solver C:	33	36	230	512	111	17
Solver D:	36	32	146	132	12	13
Solver E:	37	30	143	129	9	4
Solver F:	268	29	141	55	10	5
Solver G:	36	30	243	58	10	4

The ranks are the following:

Instance	1	2	3	4	5	6
Solver A:	2	5	2	4.5	3	4
Solver B:	7	7	4	6	6	6
Solver C:	1	6	6	7	7	7
Solver D:	3.5	4	5	4.5	5	5
Solver E:	5	2.5	3	3	1	1.5
Solver F:	6	1	1	1	3	3
Solver G:	3.5	2.5	7	2	3	1.5

We define for each solver the mean of the ranks. The winner of the challenge will be the solver with the lowest mean ranks. In the example, the mean ranks are:

Mean rank
Solver A: 3.42
Solver B: 6
Solver C: 5.66
Solver D: 4.5
Solver E: 2.66
Solver F: 2.5
Solver G: 3.25

Hence solver F is the winner.

Awards

The awards are supplied by ORTEC. They are:

1. In the all-time-best challenge €40 per instance, so €1000 in total.
2. In the restricted resources challenge the award for the winner is €2017. The runner-up is rewarded €500 and for the third place the award is €250.