

1 Automated grading of student constraint modelling 2 projects

3 Hendrik Bierlee¹ ✉ 

4 KU Leuven, Belgium

5 Tias Guns ✉ 

6 KU Leuven, Belgium

7 — Abstract —

8 Auto-grading student constraint models requires a checker that validates solutions against the
9 problem specification. In practice, for example in the Coursera Massive Open Online Course on
10 MiniZinc [5], the checker is developed separately from the teacher’s own constraint model [1]. This
11 duplicates effort and risks inconsistencies that produce confusing feedback and unfair grades.

12 We present a framework for CPMpy [3] in which the checker is automatically derived from the
13 teacher’s constraint model. By adding a feedback template to each constraint, granular feedback is
14 returned when a student solution violates a specified requirement constraint, domain, or objective.
15 Decision variables in the teacher’s model that are intentionally hidden from the student are assigned
16 consistently via a *Maximum Satisfiable Subset* (MSS) computation. The same model also provides a
17 performance baseline for grading, generates instances of various difficulty, and can be obfuscated and
18 distributed for local checking. We illustrate the approach on an example *Flexible Jobshop Scheduling*
19 (FJSS) project, demonstrating violation detection on a toy instance, and checker performance on
20 larger instances.

21 **2012 ACM Subject Classification** Theory of computation → Constraint and logic programming;
22 Applied computing → Interactive learning environments

23 **Keywords and phrases** constraint modelling, auto-grading, education

24 **Supplementary Material** *Software (Source Code)*: [link](#)

25 **Funding** This research is partly funded by the European Research Council (ERC) under the EU
26 Horizon 2020 research and innovation programme (Grant No 101002802, CHAT-Opt)

27 **1** Introduction

28 Automated grading is a common approach to assessing student programming projects. For
29 constraint modelling specifically, the prime example is the popular MiniZinc Coursera Massive
30 Open Online Course (MOOC)², where the students submit their models to a server. The
31 server solves the model under a time limit, then runs a *checker* on the generated candidate
32 solution to return feedback on the correctness and performance of the submission. The
33 feedback may include which constraints are violated and how. This feedback is a vital part
34 of graded assessments and ungraded learning exercises.

35 However, implementing a checker for a project is a time-consuming and error-prone
36 task. The main non-trivial aspect is to keep the checker consistent with the specification.
37 Otherwise, the checker may give incorrect feedback on a student model solution, which can
38 be costly. For the student, incorrect feedback can be confusing and demotivating. Worse, an
39 update to the checker may require a complete re-check, which entails running all submitted
40 models of all students over all instances for the entire time limit, in order to determine the

¹ Corresponding author

² <https://www.coursera.org/learn/solving-algorithms-discrete-optimization>

41 new best grade. If this happens during the project submission window, extensions may be
 42 needed to allow for the full time. The updated grade may be higher or lower for different
 43 students. A final concern is that if a checker is to be distributed to the student to allow local
 44 checking, then it should not reveal any knowledge outside the specification to the student.
 45 However, a procedural checker running asserts may look similar to a constraint model.

46 At ModRef'17, Coffrin *et al.* [1] described how solutions are checked for the MiniZinc
 47 MOOC. A constraint model solving the specified problem, which we will call a *solving*
 48 *model*, can be used to check a candidate solution by fixing the decision variables. An
 49 unsatisfiable solve result indicates a violation of the specification. This check is solver-free if
 50 all variables are fixed by the candidate solution. If there are *hidden variables* in the solving
 51 model, which are unspecified and thus not part of the candidate solution, a solver is used to
 52 assign them. However, this approach lacks granular feedback beyond correct and incorrect.
 53 Additionally, the solving model should not include *unspecified constraints* (e.g. symmetry
 54 breaking constraints to improve performance).

55 These obstacles motivate a separate *checking model*, in which constraints are annotated
 56 with feedback messages if violated, there are no unspecified constraints, and the non-hidden
 57 variables have become fixed parameters. The checking and solving models will look quite
 58 similar, to the point that checking models are obfuscated using e.g. a substitution cipher.
 59 However, the checking model can check but not solve the problem. When developing a
 60 project, a solving model is practically still required, given that the teacher needs to dry-run
 61 the specification, test the checking model, get performance baselines on instances, potentially
 62 distribute it as a sample model, and so on. Thus, two slightly different models need to be
 63 kept in sync with the specification at all times, which can cause the aforementioned issues.

64 In this paper, we build on this prior work and propose a CPMpy auto-grading framework
 65 which merges the functionality of both the checking and (efficient) solving model into a
 66 single *hidden model*. As needed, the hidden model can either be formulated as a solving
 67 model (for obtaining baseline objectives, generating instances), or as a checking model (for
 68 giving granular feedback, safely obfuscated for distribution). Furthermore, we propose a
 69 somewhat more robust (less reversible) method of checker model obfuscation, and handle
 70 another hidden variable issue using *Maximum Satisfiable Subset* (MSS) extraction.

71 **2 Background**

72 **2.1 Constraint Programming (CP)**

73 We cover the basics of *Constraint Programming* (CP) following Rossi *et al.* [7]. We define
 74 a *Constraint Satisfaction Problem* (CSP) $P \equiv (\mathcal{X}, D, \mathcal{C})$ as a tuple of integer variables
 75 \mathcal{X} , with domains defined by D , and \mathcal{C} a set of constraints. A constraint c with scope
 76 $scope(c) = [x_1, \dots, x_n], x_i \in \mathcal{X}$ is defined (implicitly) as a set of solutions of the form
 77 $\{\hat{v}_1, \dots, \hat{v}_n\}$ for its variable sequence. A *solution* of P is an assignment θ over variables \mathcal{X} such
 78 that $\theta(x) \in D(x), x \in \mathcal{X}$ and $[\theta(x_1), \dots, \theta(x_n)] \in c$ for all $c \in \mathcal{C}$ where $scope(c) = [x_1, \dots, x_n]$.
 79 A *Constrained Optimisation Problem* (COP) extends a CSP with an objective function f over
 80 \mathcal{X} , yielding a tuple $\langle \mathcal{X}, D, \mathcal{C}, f \rangle$. A solution θ is *optimal* if there exists no other solution θ'
 81 achieving a better objective value, i.e. $f(\theta) \leq f(\theta')$ for minimization (or \geq for maximization).

82 **2.2 Maximum Satisfiable Subset (MSS)**

83 Given hard constraints \mathcal{C}_H and soft constraints \mathcal{C}_S where $\mathcal{C}_H \cup \mathcal{C}_S$ is unsatisfiable, a *Maximum*
 84 *Satisfiable Subset* (MSS) is a maximal subset $S \subseteq \mathcal{C}_S$ satisfiable together with \mathcal{C}_H [4].

3 Method

We use a *Flexible Jobshop Scheduling* (FJSS) modelling project as running example, with the following specification:

Given an instance of n jobs, m tasks per job, k machines, and an $n \times m$ duration matrix (e.g. $\{"n":3, "m":2, "k":2, "duration":[[3,2],[2,4],[1,3]]\}$), the model contains variables X where $X[i, j]$ denotes the start time for job i , task j . Constraints: tasks within a job must be executed in order; start times are non-negative; each task must run on exactly one machine; tasks on the same machine cannot overlap. Objective: minimize total flowtime (i.e. sum of end times).

3.1 Checking via hidden model

Formally, the hidden model constructed for a given instance is a constraint model, either a CSP or COP (if an objective is specified). An example of a hidden model constructor for FJSS is shown in Listing 1. We assume the hidden model correctly implements the specification, and that in a preparation phase it has been solved to assert feasibility, and possibly find a baseline objective value o (not necessarily optimal). Note that most constraints are added with a *description*: a template string that describes the requirement in natural language and may include placeholders for variable values. A subset of the variables, $\mathcal{X}' \subseteq \mathcal{X}$ is communicated to the student in the spec. For FJSS, these are the start times, X .

```

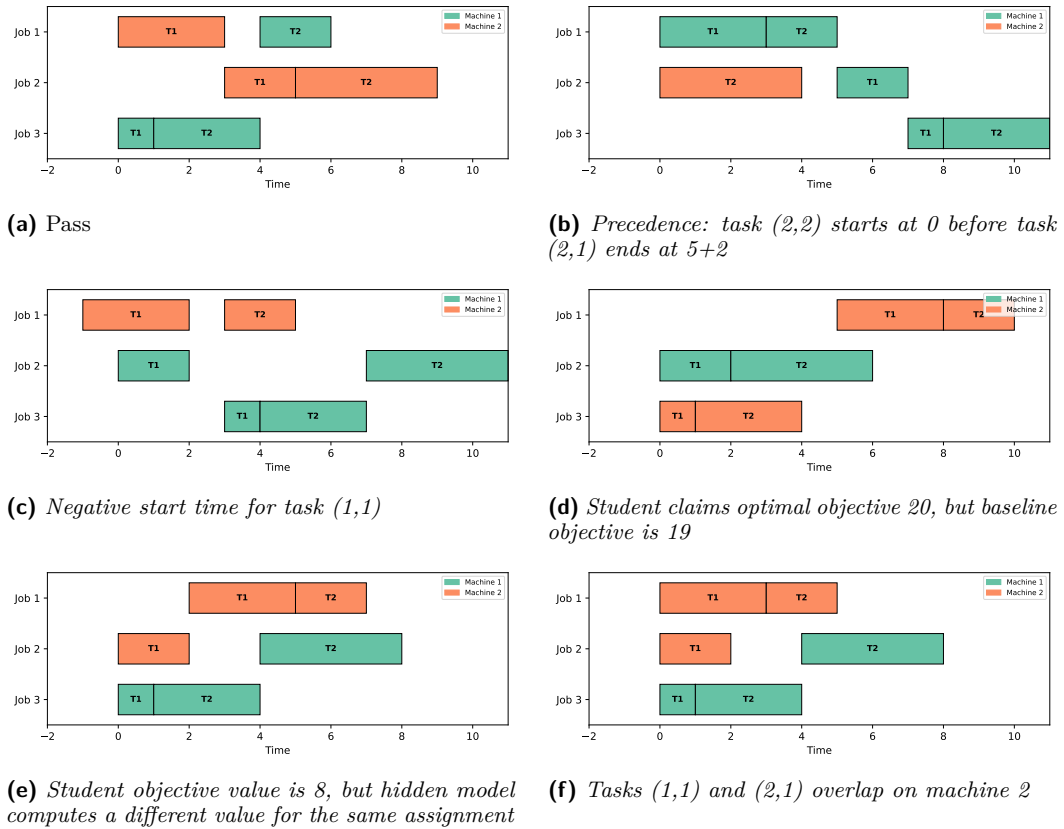
10 class FJSSModel(HiddenModel, FJSSPublic):
11     def __init__(self, instance, checking=True):
12         n, m, k = instance["n"], instance["m"], instance["k"]
13         dur = np.array(instance["duration"])
14
15         # Shared decision variables: start time for each task
16         X = cp.intvar(0, dur.sum(), shape=(n, m), name="X")
17         super().__init__(instance, shared_variables=X)
18
19         for i in range(n):
20             for j in range(m - 1):
21                 self.add(X[i, j] + dur[i, j] <= X[i, j + 1], descr=f"Precedence:
22                     ↪ task ({i + 1},{j + 2}) starts at {{X[i, j + 1].name}} before
23                     ↪ task ({i + 1},{j + 1}) ends at {{X[i, j].name}}+{dur[i,
24                     ↪ j]}")
25
26         # Non-negativity / domains
27         if checking:
28             for i in range(n):
29                 for j in range(m):
30                     self.add(X[i, j] >= 0, descr=f"Negative start time for task
31                     ↪ ({i + 1},{j + 1})")
32
33         self.minimize(sum(X[i, m - 1] + dur[i, m - 1] for i in range(n)))

```

Listing 1 Partial hidden model for the FJSS project

Now, the student submits a constraint model from which a candidate solution θ over \mathcal{X}' is generated with objective value o' , possibly with a claim of optimality. An optimal and

Automated grading of student constraint modelling projects



■ **Figure 1** Checker output for different candidate submissions on the example toy FJSS instance.

105 correct solution for the toy instance is shown in Figure 1a. First, θ is checked to ensure it
 106 assigns all shared variables. We first consider the simple case where all variables are shared,
 107 i.e. $\mathcal{X}' = \mathcal{X}$. Since θ is a total assignment, each constraint can be evaluated directly without
 108 a solver. If a constraint is violated, its description is rendered with the concrete values from
 109 θ , producing a human-readable feedback message. Figure 1b visualizes a student solution
 110 violating the precedence constraint, with the resulting feedback included in the caption.
 111 The domains in the hidden model are not explicitly checked, so that we can provide tight
 112 bounds when solving. For example, this is the case for the upper bounds of the start times,
 113 but the lower bounds are specified to be non-negative, so we explicitly add $x \geq 0$ for each
 114 $x \in \mathcal{X}$ in the hidden model (Figure 1c). Note that the upper bounds in the hidden model
 115 may be tighter than the student's to improve the performance of the solving model without
 116 restricting the student model.

117 We can verify that the student's objective value o' is correctly calculated by adding the
 118 constraint $f = o'$ (e.g. to catch the case where the student minimized makespan rather than
 119 total flowtime, see Figure 1e). If the solver reported optimality for (correct) objective value
 120 o' , we can also add the constraint $f(\theta) \leq o$ (if minimizing) to check against sub-optimality
 121 (Figure 1d). Super-optimal candidate solutions need no explicit checks as they should already
 122 violate an existing constraint.

123 When constructing the hidden model, a **checking** argument indicates whether we will
 124 solve the problem or check a candidate solution: essentially whether to generate a solving or
 125 a checking model. For instance, the lower bound constraints are redundant for solving, and

126 are added only if `checking` is true. More importantly, we omit unspecified constraints (e.g.
 127 symmetry breaking constraints) if `checking` is true. If `checking` is true, we can intentionally
 128 replace a constraint by its decomposition to give more granular feedback.

129 3.2 Hidden variables

130 As discussed by Coffrin *et al.* [1], hidden variables, i.e. $\mathcal{X} \setminus \mathcal{X}'$, can be the key challenge to
 131 a modelling project. The PHOTO problem is given as example, which has shared `pos` and
 132 hidden `who` variables, connected by an `Inverse(pos, who)` constraint. A MiniZinc solving
 133 and checking model are shown in Listings 4 and 5. A hidden model version for it is shown in
 134 Listing 2.

```

8 class PhotoModel(HiddenModel, PhotoPublic):
9     def __init__(self, instance, checking=True):
10         n, gender = instance["n"], cp.cpm_array(instance["gender"])
11
12         # Shared decision variables: position of each person (0-indexed)
13         pos = cp.intvar(0, n - 1, shape=n, name="pos")
14         super().__init__(instance, shared_variables=pos)
15
16         self.add(cp.AllDifferent(pos), descr="Positions are not unique")
17
18         if checking:
19             for i in range(n):
20                 self.add( # Position domain bounds are specified
21                     (pos[i] >= 0) & (pos[i] <= n - 1),
22                     descr=f"Person {i + 1} was not in range 0..{n - 1}",
23                 )
24
25         who = cp.intvar(0, n - 1, shape=n, name="who") # Hidden variables
26         self.add(cp.Inverse(pos, who)) # who[j] = person at position j
27
28         for j in range(n - 2):
29             self.add(
30                 (gender[who[j]] != gender[who[j + 1]])
31                 | (gender[who[j + 1]] != gender[who[j + 2]]),
32                 descr=f"3 consecutive same gender at ({j + 1},{j + 2},{j + 3})",
33             )
34
35         # Minimize total distance between consecutively numbered people
36         self.minimize(sum(abs(pos[i] - pos[i + 1]) for i in range(n - 1)))

```

■ **Listing 2** Photo Lineup (after [1]): arrange n people in a line minimizing total distance between consecutively numbered people, with no more than 2 consecutive people of the same gender. The `who` variables (inverse of `pos`) are hidden.

135 The presence of hidden variables makes checking harder in two ways. First, to assign
 136 the hidden variables, we need to invoke a solver. Note that this effectively splits up the
 137 checker model into *soft* constraints (which can be violated, returning messages), and *hard*
 138 constraints, which connect hidden variables to shared variables and avoid trivial assignments.
 139 This leads to the second problem, namely that the hard constraints may make the checking
 140 model unsatisfiable for certain incorrect candidate solutions (e.g. one which has `pos` not

141 all different), which is an uninformative result. In the MiniZinc PHOTO checking model
 142 (Listing 5), this is resolved by adding *guards* to the hard constraints so that they remain
 143 satisfiable for any candidate solution.

144 In the hidden model version of the PHOTO problem (see Listing 2), hard constraints
 145 are added by omitting the description. To add a guard, we use constraints (e.g. by
 146 `(AllDifferent(pos)).implies(Inverse(pos, who))`), which are resolved by the solver,
 147 rather than using the output phase as is done in the MiniZinc checking model. To fix
 148 the hidden variables to their candidate solution, we also use explicit hard constraints, i.e.
 149 $C_\theta \equiv \forall_{x \in \mathcal{X}'} x = \theta(x)$.

150 However, by employing a *sequential checking strategy*, we can avoid the guard altogether
 151 in a natural way. Note that in Listing 2, the `AllDifferent(pos)` is naturally placed before
 152 the `Inverse`. If `pos` is not all different, the infeasibility will signal a natural end to the
 153 check. In general, we alternate between incrementally posting and solving hard constraints
 154 (initially C_θ), checking soft constraints, then incrementally solving the next hard constraint
 155 (to assign additional hidden variables, e.g. `Inverse(pos, who)`), and so on. If the hidden
 156 model is carefully ordered, soft constraints will always have all their variables assigned, and
 157 we will not get unsatisfiable without violating at least one soft constraint. If no ordering
 158 works, guards can be used.

159 3.3 Improving hidden variable assignments

160 Yet, a third problem, which cannot be prevented with a guard, is posed by the hidden
 161 variables in the FJSS model, namely the task-machine assignment variables, `Y` (Listing 3).
 162 Hard constraints (linear equalities) enforce a task-machine assignment to avoid a trivial
 163 assignment setting all `Y` to false. No guard is needed as clearly no shared variables are
 164 in scope, so there is no risk of infeasibility. The shared and hidden variables connect via
 165 the `NonOverlapOptional` constraints, which should be soft, as we want to relay how the
 166 candidate solution violates them. In fact, a decomposition is added when `checking` is enabled,
 167 in order to show on which machine which tasks overlap. However, solving the task-machine
 168 constraint yields an arbitrary task-machine assignment (e.g. all tasks assigned to the same
 169 machine) without regard for how many task-machine overlap soft constraints are violated.

170 Instead of an arbitrary assignment of hidden variables, we get a favourable one by
 171 extracting a MSS from the hidden model. The complement of the MSS are the soft constraints
 172 violated by the associated assignment of the shared **and** hidden variables. In Figure 1f,
 173 three tasks are scheduled at time 0, which cannot be covered by the two machines, leading
 174 to an unavoidable violation. However, everywhere else, the particular task-machine MSS
 175 assignment carefully avoids the overlap. While a MSS can be costly to compute or even
 176 construct (because every soft constraint requires reification), note that we only execute it on
 177 a subset of the soft and hard constraints due to the sequential checking strategy.

178 3.4 Obfuscation

179 As mentioned, due to the similarity between a checking model and a solving model, obfuscation
 180 is required if we are to provide local checking for students of graded projects (e.g. outside
 181 the context of self-directed learning). A cipher method will suffice in most cases, however,
 182 in open-source software (or in interpreted languages such as Python), even encryption is
 183 perfectly reversible in principle, since both key and method are exposed. A less thorough, but
 184 also less reversible method of obfuscation is to transform the constraints using equivalence-
 185 preserving mutations, similar to those used in fuzz testing of constraint modelling systems [8].

```

31     # Hidden variables: which machine does which job-task
32     Y = cp.boolvar(shape=(n, m, k), name="Y")
33
34     # Hard constraint: each task runs on exactly one machine
35     for i in range(n):
36         for j in range(m):
37             self.add(sum(Y[i, j, :] == 1)
38
39     if not checking:
40         for r in range(k):
41             self.add( # No two tasks on the same machine may overlap
42                 cp.NoOverlapOptional(X.flatten(), dur.flatten(),
43                                     ↪ is_present=Y[:, :, r].flatten())
44
45     else:
46         all_tasks = [(i, j) for i in range(n) for j in range(m)]
47         for r in range(k):
48             for (i1, j1), (i2, j2) in combinations(all_tasks, 2):
49                 self.add(
50                     (Y[i1, j1, r] & Y[i2, j2, r]).implies(
51                         (X[i1, j1] + dur[i1, j1] <= X[i2, j2])
52                         | (X[i2, j2] + dur[i2, j2] <= X[i1, j1])
53                     ),
54                     descr=f"Tasks ({i1 + 1},{j1 + 1}) and ({i2 + 1},{j2 + 1})
55                         ↪ overlap on machine {r + 1}",

```

■ **Listing 3** Hidden model for FJSS (continued): hidden variables and overlap decomposition

186 Concretely, we randomly shuffle constraint order (within soft/hard blocks, to preserve the
187 sequential checking strategy) and commutative arguments; flip comparison sides ($a \leq b$ to
188 $b \geq a$); swap strict and non-strict comparators ($a \leq b$ to $a < b+1$); and shift both sides by
189 a random constant ($a \leq b$ to $a+k \leq b+k$). Note that we have more freedom to obfuscate
190 if the feedback is less granular, which is more usually the case for graded assignments.
191 The description templates attached to each constraint cannot be obfuscated further, which
192 does limit the effect. We can go further by *transforming* the hidden model into a low-level
193 representation such as *Conjunctive Normal Form* (CNF). We can even apply obfuscation
194 before and after this process. However, this introduces new hidden variables with a mix of
195 soft and hard constraints, so this idea is currently not implemented. An obfuscated hidden
196 model for FJSS for the toy instance is shown in Listing 6.

197 3.5 Instance generation

198 The hidden model also serves as the basis for an instance generator. Each project defines a
199 list of *instance configurations* (e.g. for FJSS: number of jobs, tasks, machines, and duration
200 range). Instances are generated from each configuration to obtain the desired performance
201 profile with regards to the hidden model within the set timeout. A more systematic approach
202 to generating instances from a constraint model, e.g. AutoIG [2], can be incorporated.

203 Note that, unlike the MiniZinc checking model, a separate hidden model is created for
204 each instance, as a CPMpy model is not generic over instances (there is no split such as
205 between .mzn and .dzn files). This does mean that large instances can create large checkers,

206 especially if they are obfuscated as CNF. It also means that students cannot make and
 207 check new instances, which could or could not be a helpful limitation depending on the
 208 project context. A large instance may also become hard to check (especially if the MSS is
 209 used). The flow objective was intentionally chosen because even small instances are hard to
 210 solve, whereas a makespan objective (without further changes, such as machine-dependent
 211 durations) requires large and slow checker binaries.

212 4 Evaluation

213 Table 1 shows the checker performance on five FJSS instances of increasing size, using Google’s
 214 CP-SAT solver [6] via CPMpy on an Apple M1 (8 cores, 16 GB RAM), solved single-threaded
 215 with a 60-second timeout. For each instance, we check three candidate solutions: the correct
 216 (optimal or best-found) solution, a random assignment, and a perturbed solution in which
 217 variables from the correct solution are changed one at a time (by $\pm 1-3$) until a violation is
 218 detected. The solve time shows that the basic instance generation is a useful tool, as the
 219 third instance can be re-used with multiple seeds to find a range of instances. All checks
 220 complete in under 0.5 seconds, fast enough for interactive use.

■ **Table 1** Checker performance on FJSS instances. Config: $(n, m, k, (dur_{\min}, dur_{\max}))$. Solve and check times in seconds.

Config	Solve	Check time		
		Correct	Random	Perturbed
(3, 2, 2, (1, 5))	0.007	0.03	0.04	0.04
(4, 3, 3, (2, 6))	0.061	0.05	0.11	0.11
(5, 3, 3, (3, 8))	32	0.07	0.17	0.16
(5, 4, 3, (3, 8))	TO	0.09	0.29	0.33
(6, 4, 3, (3, 8))	TO	0.12	0.44	0.38

221 We also report qualitative feedback from MSc. students enrolled at a constraint solving
 222 course which included two constraint modelling projects using initial prototypes of the checker.
 223 When asked which aspects of the course as a whole were good (not limited to the projects),
 224 out of 29 responses, the project was mentioned 20 times (‘practical’, ‘challenging’, ‘realistic’,
 225 ‘fun’), and the auto-grader was mentioned specifically as well. When asked which aspects in
 226 the course could be improved, the project was mentioned 9 times out of 28 responses. The
 227 main issue was that the project was too difficult and time-consuming. This is primarily a
 228 matter of tuning the project workload to the number of course credits, but it also shows that
 229 an auto-grader with unlimited attempts sets a rigid bar.

230 5 Conclusion

231 We presented a CPMpy framework for automated grading of constraint modelling projects in
 232 which the checker is derived directly from a single hidden model. By annotating constraints
 233 with feedback templates, the framework provides granular violation messages; hidden variables
 234 are handled via a sequential checking strategy and, if needed, MSS computation; and the same
 235 model generates instances, produces baselines, and can be obfuscated for local distribution.
 236 The approach eliminates the need for a separate checker, reducing development effort and
 237 the risk of inconsistencies.

238 — **References** —

- 239 **1** Peter J Stuckey Carleton Coffrin, Siqi Liu, and Guido Tack. Solution checking with minizinc. In *ModeRef2017, The Sixteenth International Workshop on Constraint Modelling and Reformulation*, 2017.
- 240
- 241
- 242 **2** Nguyen Dang, Özgür Akgün, Joan Espasa, Ian Miguel, and Peter Nightingale. A framework for generating informative benchmark instances. In *International Conference on Principles and Practice of Constraint Programming*, 2022.
- 243
- 244
- 245 **3** Tias Guns. Increasing modeling language convenience with a universal n-dimensional array, cppy as python-embedded example. In *Proceedings of the 18th workshop on Constraint Modelling and Reformulation at CP (Modref 2019)*, volume 19, 2019.
- 246
- 247
- 248 **4** Mark H. Liffiton, Alessandro Previti, Ammar Malik, and João Marques-Silva. Fast, flexible MUS enumeration. *Constraints An Int. J.*, 21(2):223–250, 2016. URL: <https://doi.org/10.1007/s10601-015-9183-0>, doi:10.1007/S10601-015-9183-0.
- 249
- 250
- 251 **5** Nicholas Nethercote, Peter J. Stuckey, Ralph Becket, Sebastian Brand, Gregory J. Duck, and Guido Tack. Minizinc: Towards a standard CP modelling language. In Christian Bessiere, editor, *Principles and Practice of Constraint Programming - CP 2007, 13th International Conference, CP 2007, Providence, RI, USA, September 23-27, 2007, Proceedings*, Lecture Notes in Computer Science, pages 529–543. Springer, 2007. doi:10.1007/978-3-540-74970-7_38.
- 252
- 253
- 254
- 255 **6** Laurent Perron and Frédéric Didier. Cp-sat. URL: https://developers.google.com/optimization/cp/cp_solver/.
- 256
- 257
- 258 **7** Francesca Rossi, Peter van Beek, and Toby Walsh, editors. *Handbook of Constraint Programming*, volume 2 of *Foundations of Artificial Intelligence*. Elsevier, 2006. URL: <https://www.sciencedirect.com/science/bookseries/15746526/2>.
- 259
- 260
- 261 **8** Wout Vanroose, Ignace Bleukx, Jo Devriendt, Dimos Tsouros, Hélène Verhaeghe, and Tias Guns. Mutational fuzz testing for constraint modeling systems. In Paul Shaw, editor, *30th International Conference on Principles and Practice of Constraint Programming, CP 2024, Girona, Spain, September 2-6, 2024*, LIPICs, pages 29:1–29:25. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. URL: <https://doi.org/10.4230/LIPICs.CP.2024.29>, doi:10.4230/LIPICs.CP.2024.29.
- 262
- 263
- 264
- 265
- 266

267 **A MiniZinc Photo solving and checking model**

```

1  int: n;                                % number of people
2  set of int: PERSON = 1..n;             % set of people
3  enum GENDER = { M, F, O };           % set of genders
4  array[PERSON] of GENDER: g;          % the gender of each person
5  set of int: POSN = 1..n;              % set of positions
6
7  array[PERSON] of var POSN: pos;       % decisions: a position for each person
8
9  array[POSN] of var PERSON: who;       % view: a person for each position
10 include "inverse.mzn";
11 constraint inverse(pos,who);          % channel from decisions to view
12 constraint forall(i in 1..n-2)
13     (g[who[i]] != g[who[i+1]] \ / g[who[i+1]] != g[who[i+2]]);
14
15 solve minimize sum(i in 1..n-1)(abs(pos[i] - pos[i+1]));

```

4 MiniZinc solving model for the PHOTO problem.

```

1  int: n;
2  set of int: PERSON = 1..n;

```

```

3  enum GENDER = { M, F, O };
4  array[PERSON] of GENDER: g;
5  set of int: POSN = 1..n;
6  array[int] of int: pos;
7
8  array[POSN] of var PERSON: who;
9  int: _objective;
10
11 test alldifferent(array[int] of int: x) =
12     forall(i,j in index_set(x) where i < j)(x[i] != x[j]);
13
14 include "inverse.mzn";
15 constraint if alldifferent(pos) then inverse(pos,who)
16     else forall(i in 1..n)(who[i] = 1) endif;
17
18 output [if
19     check_array(pos, n, POSN, "pos") /\
20     check_alldifferent(pos,"pos") /\
21     check_array(fix(pos), n, PERSON, "who") /\
22     forall(i in 1..n-2)
23         (check(g[fix(who[i])] != g[fix(who[i+1])] /\
24             g[fix(who[i+1])] != g[fix(who[i+2])]),
25             "three people of the same gender \(g[fix(who[i])]) in positions
26             ↪ \(\i..\(i+2)\n")) /\
27     let { int: obj = sum(i in 1..n-1)(abs(pos[i] - pos[i+1])); } in
28     check(obj = _objective, "calculated objective \(obj) does not agree with
29     ↪ objective from the model \(\_objective)\n")
30 then
31     "CORRECT: All constraints hold"
32 else
33     "INCORRECT"
34 endif];
35
36 test check(bool: b,string: s) =
37     if b then true else trace_stdout("ERROR: "+s+"\n",false) endif;
38
39 test check_alldifferent(array[int] of int: x, string: name) =
40     forall(i, j in index_set(x) where i < j)
41         (check(x[i] != x[j], name ++ "[\(\i)] = \(x[i]) = " ++
42             name ++ "[\(\j)] " ++
43             "when they should be different\n"));
44
45 test check_int(int: x, set of int: vals, string: name) =
46     check(x in vals, "integer \(name) is not in values \(vals)\n");
47
48 function bool: check_array(array[int] of int: x, int: length, set of int: vals,
49     ↪ string: name) =
50     check(length(x) = length, "array \(name) is not of length \(length)\n") /\
51     forall(i in index_set(x))
52         (check(x[i] in vals, "element \(\i) of array \(name), \(x[i]) is not in
53             ↪ values \(vals)\n"));

```

- 5 MiniZinc checking model for the PHOTO problem (after [1]). Note the guard on `inverse`.

268

B Obfuscated FJSS hidden model

```

1 60 <= 60 + (X[0,0])
2   description: Negative start time for task (1,1)
3 0 < (X[2,0]) + 1
4   description: Negative start time for task (3,1)
5 -89 <= (X[1,0]) + -89
6   description: Negative start time for task (2,1)
7 X[2,1] > -1
8   description: Negative start time for task (3,2)
9 ((X[2,0]) + 1) < (1 + (X[2,1]))
10  description: Precedence: task (3,2) starts at {X[2,1]} before task (3,1) ends at {X[2,0]}+1
11 (sum(X[0,0], np.int64(3), -58) < (sum(X[0,1], 1, -58))
12  description: Precedence: task (1,2) starts at {X[0,1]} before task (1,1) ends at {X[0,0]}+3
13 53 < sum(1, 53, X[0,1])
14  description: Negative start time for task (1,2)
15 82 < sum(X[1,1], 82, 1)
16  description: Negative start time for task (2,2)
17 ((X[1,1]) + 68) > (sum(68, X[1,0], np.int64(2), -1))
18  description: Precedence: task (2,2) starts at {X[1,1]} before task (2,1) ends at {X[1,0]}+2
19 58 == sum(57, Y[2,0,0], Y[2,0,1])
20 1 == (Y[1,1,1]) + (Y[1,1,0])
21 (Y[0,0,0]) + (Y[0,0,1]) == 1
22 1 == (Y[1,0,0]) + (Y[1,0,1])
23 1 == (Y[0,1,1]) + (Y[0,1,0])
24 1 == (Y[2,1,0]) + (Y[2,1,1])
25 ((Y[0,1,1]) and (Y[0,0,1])) -> (((X[0,0]) + 3) <= (X[0,1])) or (((X[0,1]) + 2) < (1 + (X[0,0])))
26  description: Tasks (1,1) and (1,2) overlap on machine 2
27 ((Y[1,1,1]) and (Y[0,1,1])) -> (((sum(np.int64(2), X[0,1], -71)) <= (-71 + (X[1,1]))) or (((X[1,1]) + 4)
<= (X[0,1])))
28  description: Tasks (1,2) and (2,2) overlap on machine 2
29 ((Y[0,1,0]) and (Y[0,0,0])) -> (((X[0,1]) + 2) <= (X[0,0])) or ((sum(X[0,0], np.int64(3), -12)) <
<= (sum(X[0,1], -12, 1)))
30  description: Tasks (1,1) and (1,2) overlap on machine 1
31 ((Y[0,1,1]) and (Y[1,0,1])) -> (((X[0,1]) + 2) < ((X[1,0]) + 1)) or (((X[0,1]) + -58) >
<= (sum(np.int64(2), -58, -1, X[1,0])))
32  description: Tasks (1,2) and (2,1) overlap on machine 2
33 ((Y[2,1,0]) and (Y[2,0,0])) -> (((3 + (X[2,1])) <= (X[2,0])) or (((X[2,1]) + -62) > (sum(X[2,0], -1, -62,
np.int64(1))))
34  description: Tasks (3,1) and (3,2) overlap on machine 1
35 ((Y[2,0,1]) and (Y[0,0,1])) -> (((-63 + (X[0,0])) > (sum(-1, -63, X[2,0], np.int64(1)))) or ((X[2,0]) >
<= (sum(X[0,0], -1, np.int64(3))))
36  description: Tasks (1,1) and (3,1) overlap on machine 2
37 ((Y[1,0,1]) and (Y[2,1,1])) -> (((sum(31, np.int64(2), X[1,0])) < (sum(1, 31, X[2,1]))) or (((X[2,1]) +
<= 3) <= (X[1,0])))
38  description: Tasks (2,1) and (3,2) overlap on machine 2
39 ((Y[1,0,0]) and (Y[0,1,0])) -> (((X[0,1]) + 39) >= (sum(np.int64(2), 39, X[1,0]))) or ((X[1,0]) >=
<= ((X[0,1]) + 2)))
40  description: Tasks (1,2) and (2,1) overlap on machine 1
41 ((Y[1,1,0]) and (Y[1,0,0])) -> (((X[1,1]) + 4) <= (X[1,0])) or ((X[1,1]) > (sum(-1, np.int64(2),
np.int64(1))))
42  description: Tasks (2,1) and (2,2) overlap on machine 1
43 ((Y[1,1,0]) and (Y[0,0,0])) -> (((X[0,0]) + -34) > (sum(-34, -1, np.int64(4), X[1,1]))) or ((sum(33,
np.int64(3)) < (sum(X[1,1], 33, 1)))
44  description: Tasks (1,1) and (2,2) overlap on machine 1
45 ((Y[0,0,1]) and (Y[1,0,1])) -> (((X[1,0]) >= (3 + (X[0,0]))) or ((X[0,0]) >= (2 + (X[1,0])))
46  description: Tasks (1,1) and (2,1) overlap on machine 2
47 ((Y[2,1,1]) and (Y[1,1,1])) -> (((sum(X[2,1], 77, np.int64(3)) <= (77 + (X[1,1]))) or (((X[1,1]) + 4) <
<= ((X[2,1]) + 1)))
48  description: Tasks (2,2) and (3,2) overlap on machine 2
49 ((Y[0,0,1]) and (Y[1,1,1])) -> (((sum(np.int64(4), 8, X[1,1])) <= (8 + (X[0,0]))) or ((sum(-51, X[0,0],
np.int64(3)) < (sum(-51, 1, X[1,1])))
50  description: Tasks (1,1) and (2,2) overlap on machine 2
51 ((Y[0,1,1]) and (Y[2,0,1])) -> (((X[0,1]) >= (1 + (X[2,0]))) or ((X[2,0]) >= ((X[0,1]) + 2)))
52  description: Tasks (1,2) and (3,1) overlap on machine 2
53 ((Y[1,0,0]) and (Y[2,1,0])) -> (((sum(29, np.int64(2), X[1,0])) <= (29 + (X[2,1]))) or ((44 + (X[1,0])) >
<= (sum(-1, X[2,1], 44, np.int64(3))))
54  description: Tasks (2,1) and (3,2) overlap on machine 1
55 ((Y[2,1,0]) and (Y[0,0,0])) -> (((X[0,0]) + 3) <= (X[2,1])) or (((X[2,1]) + 3) < ((X[0,0]) + 1)))
56  description: Tasks (1,1) and (3,2) overlap on machine 1
57 ((Y[0,1,0]) and (Y[2,0,0])) -> (((X[0,1]) + 2) < (1 + (X[2,0]))) or ((85 + (X[0,1])) > (sum(85, X[2,0],
np.int64(1), -1)))
58  description: Tasks (1,2) and (3,1) overlap on machine 1
59 ((Y[1,1,1]) and (Y[1,0,1])) -> (((-31 + (X[1,0])) > (sum(np.int64(4), -1, X[1,1], -31))) or
<= ((sum(np.int64(2), -5, X[1,0])) < (sum(X[1,1], 1, -5))))
60  description: Tasks (2,1) and (2,2) overlap on machine 2
61 ((Y[1,0,0]) and (Y[0,0,0])) -> (((7 + (X[0,0])) >= (sum(7, np.int64(2), X[1,0]))) or ((sum(85,
np.int64(3), X[0,0])) < (sum(X[1,0], 85, 1)))
62  description: Tasks (1,1) and (2,1) overlap on machine 1

```

```

63 ((Y[0,1,1] and Y[2,1,1]) -> (((sum(-44, np.int64(3), X[2,1]) <= ((X[0,1] + -44)) or ((18 + (X[2,1]))
↳ > (sum(18, X[0,1], np.int64(2), -1))))
64     description: Tasks (1,2) and (3,2) overlap on machine 2
65 ((Y[0,0,1] and Y[2,1,1]) -> (((X[0,0] + 94) >= (sum(np.int64(3), 94, X[2,1]))) or ((X[0,0] + 3) <
↳ ((X[2,1] + 1)))
66     description: Tasks (1,1) and (3,2) overlap on machine 2
67 ((Y[2,0,1] and Y[2,1,1]) -> (((sum(85, X[2,1], np.int64(3)) <= (85 + (X[2,0])) or ((40 + (X[2,1])) >
↳ (sum(40, X[2,0], -1, np.int64(1))))
68     description: Tasks (3,1) and (3,2) overlap on machine 2
69 ((Y[0,1,0] and Y[1,1,0]) -> (((X[0,1] >= (4 + (X[1,1]))) or ((sum(-66, X[0,1], np.int64(2))) <=
↳ ((X[1,1] + -66)))
70     description: Tasks (1,2) and (2,2) overlap on machine 1
71 ((Y[2,0,1] and Y[1,0,1]) -> (((X[1,0] > (sum(X[2,0], -1, np.int64(1)))) or ((71 + (X[2,0])) >
↳ (sum(np.int64(2), X[1,0], 71, -1))))
72     description: Tasks (2,1) and (3,1) overlap on machine 2
73 ((Y[0,0,0] and Y[2,0,0]) -> (((X[0,0] > (sum(np.int64(1), -1, X[2,0]))) or ((3 + (X[0,0])) < (1 +
↳ (X[2,0])))
74     description: Tasks (1,1) and (3,1) overlap on machine 1
75 ((Y[2,0,1] and Y[1,1,1]) -> (((X[2,0] + 23) > (sum(X[1,1], 23, -1, np.int64(4)))) or ((67 +
↳ (X[1,1])) >= (sum(67, X[2,0], np.int64(1))))
76     description: Tasks (2,2) and (3,1) overlap on machine 2
77 ((Y[2,0,0] and Y[1,0,0]) -> (((sum(-87, X[1,0], np.int64(2)) < (sum(-87, 1, X[2,0])) or (((X[1,0] +
↳ -56) > (sum(-1, X[2,0], -56, np.int64(1))))
78     description: Tasks (2,1) and (3,1) overlap on machine 1
79 ((Y[2,0,0] and Y[1,1,0]) -> (((X[2,0] + 10) >= (sum(X[1,1], 10, np.int64(4)))) or ((sum(46,
↳ np.int64(1), X[2,0]) <= (46 + (X[1,1])))
80     description: Tasks (2,2) and (3,1) overlap on machine 1
81 ((Y[0,1,0] and Y[2,1,0]) -> (((X[2,1] + -92) > (sum(-1, np.int64(2), -92, X[0,1])) or ((17 +
↳ (X[0,1])) > (sum(17, X[2,1], -1, np.int64(3))))
82     description: Tasks (1,2) and (3,2) overlap on machine 1
83 ((Y[1,1,0] and Y[2,1,0]) -> (((sum(X[1,1], 33, np.int64(4)) < (sum(X[2,1], 1, 33)) or ((X[2,1] +
↳ 3) <= (X[1,1])))
84     description: Tasks (2,2) and (3,2) overlap on machine 1
85
86 jective: min sum(X[0,1], np.int64(2), (X[1,1] + 4, (X[2,1] + 3)
87 seline: 19 (optimal: True)

```

■ 6 Obfuscated hidden model for the toy FJSS instance.