

Towards Understanding Differences Between Modelling Pipelines: a Modelers Perspective

Csobán Balogh ✉

University of St Andrews, School of Computer Science, Scotland

Ruth Hoffmann¹ ✉ 🏠 

School of Computer Science, University of St Andrews, Scotland

Joan Espasa ✉ 🏠 

University of St Andrews, School of Computer Science, Scotland

Abstract

In this work we aim to investigate the capabilities of the `MiniZinc` and `Savile Row` constraint programming pipelines from the user’s perspective. We evaluate their modelling and reformulation capabilities on a selection of six diverse problem classes using the commonly supported `Chuffed` solver. Our preliminary findings show that both pipelines are very competitive in performance. However, they seem to cater to distinct user preferences. `MiniZinc` allows better modeler control, and provides a slightly more expressive language due to the facilities for code organization and reusability. Conversely, `Savile Row` provides a solid set of default settings and a more consistent performance profile.

2012 ACM Subject Classification Theory of computation → Constraint and logic programming; Computing methodologies → Modeling methodologies

Keywords and phrases Constraint Modelling, Constraint Pipelines, Constraint Satisfaction Problems

1 Introduction

As with imperative programming languages like C, the process from writing out the problem constraints to arriving at a solution involves a pipeline of translation from a high-level language to a low-level language that can be directly interpreted by the solver. Modelling problems to be efficiently solved is no trivial task, and as a result there are automated modelling assistants who tailor the models to the required input to specific solvers. More concretely, we consider tailoring as the process of translating the constraint model with given input parameters (an instance) to a form readable by a specific constraint solver.

In this work we are going to focus on two well-known pipelines: `MiniZinc` [12] and its homonymous language, and `Savile Row` [14] and its `Essence’` [13] language. Both `MiniZinc` and `Essence’` are high-level constraint modelling languages. The `Savile Row` modelling assistant was developed between the University of York and St Andrews. `MiniZinc` is an open-source language, developed at Monash University in collaboration with Data61 Decision Sciences and the University of Melbourne. Similarly, it has its own modelling assistant to tailor the input for the solvers. The following steps describe the usual workflow for a problem to be able to be solved using these two pipelines: **1.** Using a high-level constraint modelling language, a problem is modelled. **2.** The constraints are merged with the instance data and tailored to be fed as input to a specified constraint solver. **3.** The constraint solver reads the low-level description and searches for solutions. **4.** If a solution is found, the solution is then translated back to a user-readable form. We have decided to focus on `Chuffed` [5], as the solver is used by both pipelines, given its robustness, performance and for the maturity of its support by both pipelines.

¹ corresponding author

43 We selected a set of 6 diverse problems from CSPLib [7] and the Minizinc Challenges [17].
 44 While the benchmarks for these problems already existed in either MiniZinc or
 45 **Essence'**, no direct detailed comparison has been made between the two due to models
 46 not being available in both languages. In addition, to our knowledge our work is unique in
 47 considering the full pipeline: from modelling, translation and optimisation, to the effects
 48 upon the solver. Further, while two models for the same problem may arrive at the same
 49 solution, variations in a model, such as a different viewpoint, can result in great variations in
 50 solving time [9]. Our objective is to manually create models as similar as possible to try to
 51 isolate the pipeline effects as much as possible.

52 In summary, our contributions are as follows: (i) new models in both **Essence'** and
 53 **MiniZinc** of existing problems which allow to directly compare the considered pipelines, (ii)
 54 a discussion on how both languages compare when modelling the selected problems, (iii) an
 55 empirical evaluation of the pipelines performance over 3 different optimisation levels.

56 2 The Models

57 We now present the models for the categories and problems which we chose for this work:
 58 mathematical objects (Quasigroup Completion [15]), packing problems (Wordpress [4]),
 59 scheduling problems (Rotating Rostering [18] and Travelling Tournament Problem with
 60 Predefined Values [16]), planning problems (Multi-Skilled Project Scheduling Problem) and
 61 routing problems (Capacitated Vehicle Routing Problem with Time Windows). All models
 62 can be found at <https://github.com/stacs-cp/modref2024-pipeline-comparison>.

63 **MiniZinc** has predicates and further rich functionality and syntax that **Essence'** is
 64 lacking. However, in all cases an equivalent expression can be created and run in **Essence'**
 65 to produce valid solutions. The main difference in both pipelines is the treatment of the
 66 search order. **MiniZinc** has an extensive number of search orders it may impose upon the
 67 variables, for example searching through an array of variables from the maximum of the
 68 domain to the minimum. **Essence'** on the other hand provides similar functionality for
 69 the target solver **Minion** with search heuristics, but these fall short in control compared to
 70 **MiniZinc**. We have tried to recreate the **MiniZinc** search order through a rough translation
 71 by specifying which variables to branch on first, but there is little control in how to perform
 72 the search over the variables and their respective domains.

73 2.1 Quasigroup Completion

74 In Quasigroup Completion CSPLib [15], an order n quasigroup can be viewed as a size n
 75 Latin square, where a Latin square is a $n \times n$ table of numbers ranging from 1 up to n
 76 inclusive. The table's rows and columns must be distinct in numbers. The Quasigroup
 77 Completion problem gives an incomplete quasigroup table and requires the remaining empty
 78 cells to be filled in with numbers.

79 The provided **MiniZinc** model on CSPLib uses an explicit representation, making use of
 80 the **allDifferent** constraint to ensure columns and rows are distinct. The **allDifferent**
 81 constraint takes in an array, and imposes that every element appears at most once in the
 82 array. Given the simplicity of the problem and its reliance upon **allDifferent** through
 83 the explicit model, an equivalent occurrence model was created (for both pipelines). The
 84 occurrence model represents numbers through an array of zeros and ones [8].

85 To impose **allDifferent** upon an occurrence representation, either **allDifferent_except**
 86 or **sum** could be used. Rather than compare slight variations of **allDifferent**, the interest

87 is in comparing the difference between using a simple operator such as `sum`, and a powerful
88 optimised propagator such as `allDifferent`.

89 There does not exist an *Essence*' model for Quasigroup Completion on `CSPLib`, so
90 one needed to be created. The explicit model follows the `MiniZinc` model exactly, except
91 *Essence*' has a nice spread operator for slicing columns and rows from matrices instead of
92 list comprehensions. In *Essence*' we represent the occurrence model, by finding a three
93 dimensional 0,1 matrix and using a `sum` constraint.

94 2.2 Wordpress Application Deployment in the Cloud

95 Wordpress Application Deployment in the Cloud [4] was used in the 2022 `MiniZinc` chal-
96 lenge [10] with minor adaptations. Wordpress is an application used to create websites and
97 needs to be deployed to the cloud with a set of hardware and virtual machines (VMs). The
98 goal is to deploy several Wordpress instances while minimising the cost of hardware, VM
99 configurations and providers to deploy the instances into the cloud. The Wordpress problem
100 is an example of a bin packing problem, a common problem in combinatorial optimization
101 ranging over many applications such as loading trucks or graphics card resource allocation.

102 The `MiniZinc` model used is taken from the `MiniZinc` challenge 2022 [10]. The original
103 `MiniZinc` model [6], which the `MiniZinc` challenge model is based on, included variables not
104 used within the input files given. More concretely, `FVariables` and `FVInstances` are Fixed
105 Variables and Fixed Variable Instances respectively and were used to determine the effects of
106 fixing certain variables otherwise left to be optimized and found. The paper also originally
107 left the number of VMs as a decision variable, however the `MiniZinc` challenge model has an
108 input parameter for the upper bound on the number of VMs to use. This restricts the search
109 space on the number of VMs and focusing on the objective of minimising the cost.

110 The *Essence*' model is a one-to-one translation of the `MiniZinc` model. Note the
111 `MiniZinc` model does not break symmetry upon the distribution of hardware on the VMs.
112 The hardware deployed on a VM can be swapped to any other VM for an equivalent
113 solution, resulting in the symmetry of the problem. The original model [6] does break
114 the symmetry of the problem, using multiple different methods such as pricing and the
115 lexicographical assignment of hardware on machines among others. This change from the
116 paper to the `MiniZinc` challenge model might have been done to test a solvers ability to
117 recognize symmetry and break it.

118 While the non-symmetric version of the problem is of interest following the `MiniZinc`
119 challenge to compare to, the symmetric model is more attuned to how a modeller would
120 write this given problem. We created a version of the problem in *Essence*' that breaks the
121 aforementioned symmetries. To remain fair, an equivalent symmetry breaking version of the
122 `MiniZinc` model was created to compare the symmetry breaking version of the *Essence*'
123 model. Both the symmetry-breaking and non symmetry-breaking models will be compared.

124 2.3 Rotating Rostering Problem

125 The Rotating Rostering Problem [18] generalises many real life rostering problems, such as
126 nurse rostering. The goal is to find a satisfiable assignment of shifts for each day to fulfill
127 the shifts requirements and avoid conflicts. A shift type may either be a day off, a morning
128 shift, a late shift, or a night shift. The shifts are ordered such that the following constraints
129 are satisfied: **1.** The number of staff required for a day is satisfied as needed. **2.** There is a
130 min and max number of consecutive shift assignments for the same shift type. **3.** The shift
131 ordering is forward rotating; a shift must be preceded by a larger or equal shift type, or a

rest day. 4. Weekends (Saturday Sunday) have the same shift type. 5. At least two days must be rest days every 14 days.

When converting the MiniZinc model over to *Essence'*, we identified what we believe is an overlooked edge case. Incorrect solutions rarely appeared, but surfaced when running generated instance solutions through a solution checker. The issue related to the minimum number of shift assignments mandated by the constraints. As part of the original constraints, if the shift type of one day is different to the next day, the following days must meet the minimum number of consecutive shifts. This works well, until the edge case of the very first shift forgoing the minimum number of consecutive shifts, as the constraints ensure the following days meet the minimum number of consecutive shifts as seen in Listing 1. To fix this, we simply an extra constraint from the very first shift as seen in Listing 2.

■ Listing 1 MiniZinc minimum number of consecutive shifts constraint

```
143
144 constraint forall(day in 1..numberOfDays - s_min) (
145     plan1d[day] != plan1d[day+1] -> all_equal(plan1d[day+1..day+s_min]));
146
```

■ Listing 2 MiniZinc minimum number of consecutive shifts additional constraint

```
147
148 constraint(all_equal(plan1d[1..s_min]));
149
```

2.4 Traveling Tournament Problem with Predefined Venues

The Traveling Tournament Problem with Predefined Venues (TTPPV) [16], is a specialisation of the Traveling Tournament Problem. A set of teams playing in a tournament is organized as a simple round robin schedule, with each game playing at different venues. The objective is to minimize the distance travelled by the teams between different venues.

The difference to the Traveling Tournament Problem is that in TTPPV the venues of each game are predefined. This means if team *a* plays against team *b*, the venue is predefined as being at either *a*'s home or *b*'s home. With the predefined venues, the problem lies in the scheduling of the games and minimizing the sum of the traveling distances of the teams. The problem is further specialised by using circular distances between the venues for simplicity.

The `regular` predicate is used within the MiniZinc model to assert there are at most 3 consecutive home games and at most 3 consecutive away games. More generally, the `regular` predicate asserts that a sequence of variables take a value from a finite automaton, where the automaton in the MiniZinc model asserts at most two consecutive away or home games. When translating the `regular` predicate to *Essence'*, we use a `forall` statement, checking that there are not four consecutive assignments.

2.5 Capacitated Vehicle Routing problem with Time Windows, Service Times and Pickup and Deliveries

The Capacitated Vehicle Routing problem with Time Windows, Service Times and Pickup and Deliveries (CVRPTW) is an example of a routing problem which specialises the Capacitated Vehicle Routing Problem [19] and was sourced from the 2022 MiniZinc challenge [17]. It is defined as follows: there are several vehicles with a given capacity for goods, and there are a number of pickup and drop-off locations of customers to deliver to. These pickup and delivery locations have an associated demand for the goods the vehicles need to pick up or deliver. As an added constraint, there are specified time windows for the deliveries to each

175 customer such that the delivery truck must arrive and leave within this time window from
 176 the delivery location. The route chosen may not take any sub-tours for any route it takes.

177 As part of the MiniZinc model, the `circuit` predicate is used to ensure the vehicle
 178 delivery routes do not take sub-tours in their route and visits each location uniquely for
 179 optimisation. A circuit is such that the cell value of an array points to the index of the next
 180 number, and this forms a circuit that continues around. For the translation to `Essence'`,
 181 we used the decomposition of the `circuit` predicate from the MiniZinc library ².

182 To create the equivalent expression of circuit in `Essence'`, a new variable is introduced to
 183 determine and constrain the ordering of values to form the circuit. The `Essence'` equivalent
 184 of circuit for the decision variable `successor` in the model is as follows in Listing 3.

■ Listing 3 `Essence'` circuit predicate equivalent

```

185 allDiff(successor),
186 forAll i : NODES . successor[i] != i,
187 allDiff(successorOrder),
188 successorOrder[1] = 1,
189 forAll i : NODES .
190     (successorOrder[i] = maxNodes -> successorOrder[successor[i]] = 1) /\
191     (successorOrder[i] != maxNodes -> successorOrder[successor[i]] =
192         successorOrder[i] + 1)
193

```

195 2.6 Multi-Skilled Project Scheduling Problem

196 The Multi-Skilled Project Scheduling Problem (MSPSP) is a variation on the basic resource-
 197 constraint project scheduling problem [1] used in the 2012 Minizinc Challenge. In this
 198 problem there are a series of workers, with each worker having a specific skill set. There are
 199 several activities with an associated skill requiring completion to finish the project, and the
 200 overall goal is to minimize the project time.

201 The MiniZinc formulation uses set variables, where `Essence'` (unlike `Essence` [3]) lacks
 202 modelling support for them. To overcome this limitation, the sets from MiniZinc were
 203 translated into the occurrence representation of the numbers. This allowed for each matrix to
 204 be equivalent in size to satisfy the `Essence'` language limitations. A disadvantage of using
 205 this method is that by using the occurrence representation the parameter files become larger.

206 The MiniZinc model also makes use of `letting` to create variables within constraints,
 207 but `Essence'` cannot do the same. As a result, an equivalent expression is created. On Line 4
 208 of Listing 4, a new Boolean variable is introduced into the scope of the constraint. This
 209 variable acts like a normal decision variable, with the goal of assigning a satisfiable value. In
 210 Line 5 and Line 6 the Boolean variable `before` in combination with an implication ensures
 211 at least one of the expressions following the implication is true. This can be compactly
 212 expressed in `Essence'` by using an `or`, as shown in Listing 5 on Line 5.

■ Listing 4 Usage of MiniZinc `letting` in MSPSP

```

213 constraint
214     forall (i, j in Tasks where i < j /\ not(j in suc[i]) /\ not(i in suc[j]))(
215         if exists( k in Skills )( rr[k,i] + rr[k,j] > rc[k] ) then
216             let { var bool: before } in (
217

```

² fzn_circuit.mzn in [11]

6 Towards Understanding Differences Between Modelling Pipelines

```
218         (before -> s[i] + d[i] <= s[j])
219         /\ (not(before) -> s[j] + d[j] <= s[i]) )
220     else true endif);
221
```

Listing 5 Essence' letting equivalent to Listing 4

```
222
223 forall i : Tasks . forall j : Tasks .
224     (i < j /\ suc[i,j] = 0 /\ suc[j,i] = 0) ->
225         ((exists k : Skills .
226             rr[k,i] + rr[k,j] > rc[k]) ->
227             ((s[i] + d[i] <= s[j]) /\ (s[j] + d[j] <= s[i]))),
228
```

229 The MiniZinc model (Listing 6) has a series of further lettings: `WTasks` and `TWorkers`.
230 `WTasks` and `TWorkers` are sets, where `WTasks` is the set of tasks where a skill for that task
231 exists, and `TWorkers` is the set of workers who have an existing skill required. To create an
232 equivalent in `Essence'`, a single variable is introduced, `TWorkers`, encompassing `WTasks`
233 and `TWorkers` together in a 2d matrix. This is expressed in Listing 7.

Listing 6 Additional MiniZinc lettings in MSPSP

```
234
235 let { set of int: WTasks =
236     { i | i in Tasks where exists(k in has_skills[j])(rr[k, i] > 0) }
237 } in...
238 let { set of int: TWorkers =
239     { j | j in Workers where exists(k in has_skills[j])(rr[k, i] > 0) }
240 } in...
241
```

Listing 7 Essence' equivalent to Listing 6

```
242
243 forall i : Tasks . forall j : Workers .
244     TWorkers[j, i] = 1 <->
245     exists k : Skills . has_skills[j, k] = 1 /\ rr[k,i] > 0,
246
```

247 `TWorkers` is then leveraged in all further constraints equivalent to the lettings of
248 `TWorkers` and `WTasks` as seen in Listing 8. Listing 8 constraints the number of workers with
249 a set skill working upon a task to satisfy the requirements. Using the `TWorkers` variable
250 created in Listing 7, the equivalent of the MiniZinc Listing 8 is created for `Essence'` in
251 Listing 9.

Listing 8 MiniZinc letting over TWorkers

```
252
253 constraint forall ( i in Tasks ) (
254     let {
255         set of int: TWorkers =
256             { j | j in Workers where exists(k in has_skills[j])(rr[k, i] > 0) }
257     } in (
258         forall ( k in Skills where rr[k, i] > 0 )
259             (sum(j in TWorkers where k in has_skills[j])(
260                 bool2int(w[j, i])) >= rr[k, i])
261         /\ forall ( j in Workers where not(j in TWorkers) )
262             (w[j, i] = false));
263
```

Listing 9 Essence' equivalent to Listing 8

```

264 forAll i : Tasks . forAll k : Skills .
265     rr[k, i] > 0 ->
266         sum([w[j,i] /\ TWorkers[j,i] = 1 /\ has_skills[j,k] = 1 | j : Workers])
267             >= rr[k, i],
268 forAll i : Tasks . forAll j : Workers .
269     TWorkers[j, i] = 0 -> w[j,i] = false,
270
271

```

272 The `cumulative` predicate is used to determine if the cumulative resource usage is within
 273 bounds. That is, a set of tasks with start times, durations, and resource requirements, never
 274 exceed the global resource bound at any time. The `cumulative` predicate is a common
 275 predicate used in scheduling problems and is therefore optimized in most solvers such as
 276 `Chuffed`. In the translation to `Essence'` we used the default `MiniZinc` decomposition of
 277 `cumulative`³. The `cumulative` predicate is used twice in the MSPSP `MiniZinc` model, both
 278 with `letting` statements. The first `cumulative` imposes that at least one worker fulfills the
 279 task assignment while respecting the duration and timings. The second `cumulative` ensures
 280 the resources requirements is exceeded or equaled by workers while respecting durations and
 281 orderings. The equivalent `Essence'` is created in Listing 10.

■ Listing 10 `Essence'` equivalent of `cumulative` in MSPSP

```

282 forAll work : Workers .
283     sum(TWorkers[work,..]) > 1 ->
284         (forAll j : Tasks .
285             1 >= sum([(TWorkers[work, j] = 1) /\ (TWorkers[work, i] = 1) /\
286                 (s[i] <= s[j]) /\ (s[j] < (s[i] + d[i]))
287                 /\ w[work,i] | i : Tasks])),
288 forAll k : Skills .
289     (sum([rr[k,i] > 0 | i : Tasks]) > 1 /\
290         sum([rr[k,i] | i : Tasks, rr[k,i] > 0]) > rc[k]) ->
291         forAll j : Tasks .
292             rr[k,j] > 0 ->
293                 rc[k] >= sum([(s[i] <= s[j] /\ s[j] < (s[i] + d[i]))*rr[k,i] | i :
294                     Tasks, rr[k,i] > 0]),
295
296

```

3 Experiments

298 Our experiments aim to identify performance differences between `Savile Row` 1.9.1 and
 299 `MiniZinc` 2.7.5. The experiments were run using 3 different optimisation levels that the
 300 respective developers offer for each pipeline. That is, no optimisation (O0S0 in `Savile Row`,
 301 O0 in `MiniZinc`), intermediate optimisation (O2S1 in `Savile Row`, O1 in `MiniZinc`) and full
 302 optimisations (O3S2 in `Savile Row`, O5 in `MiniZinc`). Note that these will differ between
 303 pipelines, in particular due to `Savile Row` allowing control over the amount of symmetry
 304 breaking constraints introduced during tailoring, something that `MiniZinc` does not enable.

305 Certain problems were lacking in the number of instances or in their variety. To com-
 306 pensate, additional instances were generated through a combination of Python scripts or
 307 parameterised generators as constraint models, similarly to previous approaches [2].

308 Timing information is presented as the quotient of `Essence'` over `MiniZinc`. That is, a

³ `fzn_cumulative_task` of `fzn_cumulative.mzn` in [11]

number > 1 suggests **MiniZinc** was faster, while a number < 1 suggests **Essence'** was faster. These timings use the geometric mean, which uses the product rather than the sum, giving a better indicator of the central tendency of runs. In Table 1 we can see that both pipelines produce very similar and consistent results with respect to solving a given problem. Both pipelines solve all instances for the Rostering and Scheduling problems, and neither finds any solutions to the vehicle routing problems (within the set timeout). When the problems get harder (or more varied), such as the Quasigroup problems, modelling in **Savile Row** seems to be more consistent between the model representations and different optimisations levels. In the other problem groups both pipelines can be considered equal.

From a timing perspective, **MiniZinc** clearly outperforms **Savile Row** in MSPSP and Rostering (the problem classes where both solvers find all solutions). Meanwhile, in the two Quasigroup models **Savile Row** performs faster, and is solving more of the instances. We see this as an indicator that, in the given dataset, **Savile Row** performs better on harder instances. We believe that **MiniZinc** outperforms **Savile Row** on easy instances, as **MiniZinc** has a very low initialisation time when compared to **Savile Row**, as it is written in C++. In problems where both pipelines solved some of the problems, the results are inconclusive. Although the Wordpress problem without the explicit symmetry breaking constraints performs better in the highest optimisation level in **MiniZinc**, with the explicit symmetry breaking constraints **Savile Row** performs better the higher the optimisation levels.

Problem	#	Essence'			MiniZinc			Timing Ratio		
		O0S0	O2S1	O3S2	O0	O1	O5	$\frac{O0S0}{O0}$	$\frac{O2S1}{O1}$	$\frac{O3S2}{O5}$
Quasigroup	43	41	42	41	40	39	40	0.08	0.5	0.6
Quasigroup Occ.	43	41	41	42	32	37	38	0.12	0.08	0.38
Wordpress	9	6	6	6	6	6	6	1.54	1.83	5.29
Wordpress Symm.	9	4	4	6	4	4	4	1.47	1.36	0.49
TTPPV	20	3	3	3	3	3	3	0.99	1.35	1.98
MSPSP	6	6	6	6	6	6	6	138.48	88.57	612.61
CVRPTW	5	0	0	0	0	0	0	1.0	1.0	1.0
Rostering	7	7	7	7	7	7	7	29.5	15.15	77.7

Table 1 Columns **Essence'** and **MiniZinc** show the number of solved instances per problem, split between the 3 considered optimisation levels. Timing ratios show the ratio between **Essence'** and **MiniZinc** options, where > 1 denotes **MiniZinc** was faster and < 1 otherwise.

4 Conclusions and Further Work

These initial findings seem to suggest that **Minizinc** might be better suited for scenarios where an expert modeler can leverage the capabilities of the pipeline and where code maintainability are crucial. Conversely, **Savilerow** strong reformulation capabilities and good default settings would be the preferred choice for tackling complex problems where consistent performance is paramount. To solidify these findings, a wider selection of both solvers and problems have to be considered.

References

- 1 Rina Agarwal, Manoj K Tiwari, and Sanat K Mukherjee. Artificial immune system based approach for solving resource constraint project scheduling problem. *The International Journal*

- 339 of *Advanced Manufacturing Technology*, 34(5):584–593, 2007.
- 340 2 Özgür Akgün, Nguyen Dang, Ian Miguel, András Z. Salamon, and Christopher Stone. Instance
341 generation via generator instances. In Thomas Schiex and Simon de Givry, editors, *Principles*
342 *and Practice of Constraint Programming*, pages 3–19, Cham, 2019. Springer International
343 Publishing.
- 344 3 Özgür Akgün, Alan M. Frisch, Ian P. Gent, Christopher Jefferson, Ian Miguel, and Peter
345 Nightingale. Conjure: Automatic Generation of Constraint Models from Problem Specifications.
346 *Artif. Intell.*, 310:103751, 2022. doi:10.1016/J.ARTINT.2022.103751.
- 347 4 David Bogdan-Nicolae. CSPLib problem 090: Wordpress application deployment in the cloud.
348 <http://www.csplib.org/Problems/prob090>.
- 349 5 Geoffrey Chu, Peter J Stuckey, Andreas Schutt, Thorsten Ehlers, Graeme Gange, and Kathryn
350 Francis. Chuffed, a lazy clause generation solver, 2018. URL: [https://github.com/chuffed/](https://github.com/chuffed/chuffed)
351 [chuffed](https://github.com/chuffed/chuffed).
- 352 6 Mădălina Eraşcu, Flavia Micota, and Daniela Zaharie. Scalable optimal deployment in the cloud
353 of component-based applications using optimization modulo theory, mathematical programming
354 and symmetry breaking. *Journal of Logical and Algebraic Methods in Programming*, 121:100664,
355 2021.
- 356 7 Ian P Gent and Toby Walsh. Csplib: a benchmark library for constraints. In *International*
357 *Conference on Principles and Practice of Constraint Programming*, pages 480–481. Springer,
358 1999.
- 359 8 Christopher Jefferson and Alan M Frisch. On the effectiveness of set and multiset repres-
360 entations in constraint programming. *Modelling and Reformulating Constraint Satisfaction*
361 *Problems*, page 125, 2004.
- 362 9 Vipin Kumar. Algorithms for constraint-satisfaction problems: A survey. *AI magazine*,
363 13(1):32–32, 1992.
- 364 10 MiniZinc. Minizinc challenge 2022 results, 2022. URL: [https://www.minizinc.org/](https://www.minizinc.org/challenge2022/results2022.html)
365 [challenge2022/results2022.html](https://www.minizinc.org/challenge2022/results2022.html).
- 366 11 MiniZinc Developers. MiniZinc Standard Library. [https://github.com/MiniZinc/](https://github.com/MiniZinc/libminizinc/tree/master/share/minizinc/std)
367 [libminizinc/tree/master/share/minizinc/std](https://github.com/MiniZinc/libminizinc/tree/master/share/minizinc/std).
- 368 12 Nicholas Nethercote, Peter J. Stuckey, Ralph Becket, Sebastian Brand, Gregory J. Duck, and
369 Guido Tack. Minizinc: Towards a standard cp modelling language. In Christian Bessière,
370 editor, *Principles and Practice of Constraint Programming – CP 2007*, pages 529–543, Berlin,
371 Heidelberg, 2007. Springer Berlin Heidelberg.
- 372 13 Peter Nightingale. Savile row manual. *CoRR*, abs/2201.03472, 2022. URL: [https://arxiv.](https://arxiv.org/abs/2201.03472)
373 [org/abs/2201.03472](https://arxiv.org/abs/2201.03472), [arXiv:2201.03472](https://arxiv.org/abs/2201.03472).
- 374 14 Peter Nightingale, Özgür Akgün, Ian P. Gent, Christopher Jefferson, Ian Miguel, and
375 Patrick Spracklen. Automatically improving constraint models in savile row. *Artificial*
376 *Intelligence*, 251:35–61, 2017. URL: [https://www.sciencedirect.com/science/article/](https://www.sciencedirect.com/science/article/pii/S0004370217300747)
377 [pii/S0004370217300747](https://www.sciencedirect.com/science/article/pii/S0004370217300747), doi:<https://doi.org/10.1016/j.artint.2017.07.001>.
- 378 15 Gilles Pesant. CSPLib problem 067: Quasigroup completion. [http://www.csplib.org/](http://www.csplib.org/Problems/prob067)
379 [Problems/prob067](http://www.csplib.org/Problems/prob067).
- 380 16 Gilles Pesant. CSPLib problem 068: Traveling tournament problem with predefined venues
381 (tttpv). <http://www.csplib.org/Problems/prob068>.
- 382 17 Peter J Stuckey, Thibaut Feydy, Andreas Schutt, Guido Tack, and Julien Fischer. The minizinc
383 challenge 2008–2013. *AI Magazine*, 35(2):55–60, 2014.
- 384 18 Petra Hofstedt Sven Löffler, Ilja Becker. CSPLib problem 087: Rotating rostering problem.
385 <http://www.csplib.org/Problems/prob087>.
- 386 19 Özgür Akgün. CSPLib problem 086: Capacitated vehicle routing problem. [http://www.](http://www.csplib.org/Problems/prob086)
387 [csplib.org/Problems/prob086](http://www.csplib.org/Problems/prob086).