

# 1 **Mutational Fuzz-testing for Constraint Modeling**

## 2 **Systems**

3 **Anonymous author**

4 **Anonymous affiliation**

### 5 **Abstract**

---

6 Constraint programming (CP) modeling languages, like MiniZinc and CPMpy, play a crucial role in  
7 making CP technology accessible to non-experts. Both solver-independent modeling languages and  
8 solvers themselves are complex pieces of software that can contain bugs, which undermines their  
9 usefulness. Mutational fuzz-testing is a way to test complex systems by stochastically mutating  
10 input and verifying preserved properties of the mutated output. We investigate different mutations  
11 and verification methods that can be used in the context of constraint programming. This includes  
12 methods proposed in the context of SMT-solving, as well as new methods related to global constraints,  
13 optimisation and solution counting/preservation. Our results show that such a fuzz testing approach  
14 improves the overall code coverage of a modeling system compared to only unit testing, and is able  
15 to find bugs in the whole toolchain, from the modeling language transformations themselves to the  
16 underlying solvers.

17 **2012 ACM Subject Classification** Replace `ccsdesc` macro with valid one

18 **Keywords and phrases** Dummy keyword

19 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

20 **Acknowledgements** Anonymous acknowledgements

## 21 **1 Introduction**

22 Constraint solving is a declarative AI reasoning technique that is used in a variety of high-  
23 stakes applications ranging from scheduling production lines [17] to automated verification  
24 of computer programs [18] and aerospace applications [31]. All of these applications require  
25 constraint solvers to provide correct and reliable solutions to the constraint specifications.

26 To leverage the power of modern constraint solvers, it is common for users to write down  
27 the problem specification in a high level, declarative *constraint modeling language* such as  
28 MiniZinc [23], XCSP [30], Conjure [2] or CPMpy [15]. These modeling languages play a  
29 fundamental role in enabling the wider adoption of CP technology across various domains as  
30 they provide high-level, expressive, and intuitive methods for users to define complex problem  
31 constraints. They offer an abstraction from the details of encoding high-level constraints  
32 into the specific constraints supported by a solver, allowing users to focus on the problem at  
33 hand rather than the specifics of the solvers. Modeling systems then rewrite the high-level  
34 user-constraints into solver-specific expression such as clauses, linear constraints or unnested  
35 global constraints. Therefore, the code base of modeling systems typically contains multiple  
36 reformulation and encoding algorithms. They are also made more complex by optimizations  
37 used to reduce the number of generated low-level constraints such as Common Subexpression  
38 Elimination (CSE) [24, 25, 27]. In some cases, these transformations are mixed-and-matched  
39 in different ways for different solvers.

40 Like all complex software, modeling systems and constraint solvers can contain bugs. In  
41 the case of modeling systems, bugs can cause a range of undesired behavior: from experiencing  
42 crashes of the system itself to returning an invalid or non-optimal solution to the constraints  
43 stated by the user. Especially the latter can have a major impact on the user and the



© **Anonymous author(s);**

licensed under Creative Commons License CC-BY 4.0

42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:25

Leibniz International Proceedings in Informatics

 **LIPICS** Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

## 23:2 Mutational Fuzz-testing for Constraint Modeling Systems

44 application at hand. Moreover, it can also decrease the trust of users towards the underlying  
45 solving techniques.

46 To mitigate the number of bugs in computer programs, it is good practice to use some  
47 kind of *automated testing* during software development. *Unit testing* [12] is such a technique  
48 to test isolated parts of the code. While unit testing is very useful to verify the intended  
49 behavior of a program, it is time-consuming for developers to write as it necessitates testing  
50 for both expected and unexpected inputs. Therefore, tricky edge cases may be overlooked  
51 when designing the test suite. In constraint solving, this is especially the case for non-trivial  
52 combinations of constraints that share variables.

53 *Fuzz testing* is a family of techniques that automatically test computer programs on  
54 randomly constructed inputs. These techniques can either be *generation-based* or *mutation-based*: the former generates input from scratch, while the latter uses existing inputs and applies  
55 mutations to them in order to construct a valid new input. Fuzz-testing has proved to be  
56 extremely successful in finding bugs in a variety of computer programs: from testing Android  
57 apps [39], to crashes of Unix command-line utilities [22], and SMT solvers[20, 38]. Although  
58 fuzz testing has been used to test several solver-specific algorithms such as propagation  
59 routines [20, 38, 3, 9, 26], it has not yet been applied to solver-independent constraint  
60 modeling languages, despite their rapid development in recent years.

61 In this paper, we draw inspiration from systems such as STORM [20] and YinYang [38] tailored  
62 to test SMT solvers, and propose HURRICANE, a method to use *mutational fuzz testing*  
63 for generic constraint modeling systems. New opportunities for fuzz testing arise, because  
64 of the rich constraint specification that CP modeling languages allow. These include the  
65 use of global constraints and their decompositions [34], the use of n-ary aggregate functions,  
66 the possibility of arbitrarily nested expressions (even global constraints) that may require  
67 flattening, the use of objective functions, and the changing transformation flows that are  
68 used for different available backend solvers.

69 Our contributions are:

70 1. We propose a generic, mutation-based, automated testing framework, HURRICANE, for  
71 verifying the correctness of solver-independent CP modeling languages and their solvers;  
72 2. We investigate the use of 3 families of mutations; as well as 5 methods to verify the  
73 mutated models do not contain bugs; and  
74 3. We evaluate HURRICANE by testing it on the CPMpy constraint modeling system [15]  
75 and show that its effectiveness at finding bugs in the system itself as well as its underlying  
76 solvers.

## 77 2 Related work

78 Automated testing of computer programs finds its roots in *unit testing* [12]. A unit test  
79 consists of a small use case of a part of the software as envisioned by the developers. The  
80 technique was made popular by the JUnit testing framework in Java [32]. In recent years,  
81 researchers have studied ways to automatically synthesize unit tests in order to improve *code  
82 coverage* of the test suite [19]. Code coverage quantifies the number of lines of code in a  
83 program that is executed by a (set of) tests. While this is not a foolproof metric [36], it is a  
84 reasonable proxy to evaluate how thoroughly a system is tested.

85 Fuzz testing has been used in combinatorial solving before. An early form of testing  
86 SAT-solvers uses generation-based techniques [10], and more recently, several solvers who  
87 entered the 2022 edition of the Max-SAT competition were subjected to fuzz testing [26]. In  
88 the field of CP, *generation-based fuzz testing* has already been adopted as an automatic testing

90 technique for solvers. For example, the propagation algorithms present in the MINION  
 91 solver have been automatically fuzz-tested throughout its development [3]. The input used  
 92 for testing such propagation routines is a randomly generated set of constraints within the  
 93 relatively simple grammar supported by the solver. The output of the solver is verified using  
 94 simpler, but equivalent algorithms.

95 Compared to the API of a constraint solver, CP modeling languages allow for a much  
 96 richer set of expressions to be written down by a user. This makes stochastic generation of  
 97 inputs more complex, hence we turn our attention to *mutational fuzz testing* techniques that  
 98 were applied to satisfiability checking SMT solvers [20, 38, 8]. These techniques can generate  
 99 deeply nested expressions in the rich, nested language that SMT solvers natively accept as  
 100 input. While also applicable to high-level constraint modeling languages, we propose new  
 101 mutations and verification methods based on this richer input.

102 Finally, a very different kind of technique to detect bugs in combinatorial solvers is  
 103 through the use of *proof logging*. Proof logging requires a system to write down the result of  
 104 its algorithms as relatively simple mathematical reasoning steps. Such proofs are then *verified*  
 105 automatically by a third-party checker [13, 16, 14]. SAT solvers are required to output proof  
 106 logs (mathematical search certificates) in order to enter the yearly SAT competition<sup>1</sup>. In  
 107 recent years, proof logging has successfully found its way to other combinatorial search  
 108 algorithms such as those used in (Max-)SAT-, ASP-, SMT- and CP [35, 4, 5, 28, 7, 21].  
 109 However, proof logging for now remains a low-level technique that is not directly applicable  
 110 to algorithms that translate any high-level expressions into multiple equivalent low-level  
 111 solver constraints.

### 112 3 Preliminaries

113 A *Constraint Satisfaction Problem (CSP)* is a triple  $(\mathcal{X}, \mathcal{D}, \mathcal{C})$  [29] with

- 114   ■  $\mathcal{X}$  a set of *decision variables*;
- 115   ■  $\mathcal{D}$  a set of *domains of values* for each variable in  $\mathcal{X}$ ;
- 116   ■  $\mathcal{C}$  a set of *constraints*, each over some subset of  $\mathcal{X}$ .

117 An *assignment* maps variables in  $\mathcal{X}$  to a value in their domain. A *constraint* maps  
 118 assignments to true or false. An assignment *satisfies* a constraint if the constraint maps  
 119 it to true. We make no assumption on the structure of a constraint, that is, it can be a  
 120 nested expression as we will see below. A *solution* to a CSP is an assignment over all  $\mathcal{X}$   
 121 that satisfies all constraints in  $\mathcal{C}$ . The set of solutions of a set of constraints, projected  
 122 to a set of variables  $\mathcal{X}$  is written as  $sols_{\mathcal{X}}(\mathcal{C})$ . E.g., given the following set of constraints  
 123  $\mathcal{C} = \{p + q + z \leq 2, p < q\}$  and positive domains for  $p, q$  and  $z$ , we observe the following sets  
 124 of solutions:

$$125 \quad sols(\mathcal{C}) = \{\{p \mapsto 0, p \mapsto 1, z \mapsto 0\}, \{p \mapsto 0, p \mapsto 1, z \mapsto 1\}, \{p \mapsto 0, p \mapsto 2, z \mapsto 0\}\}$$

$$126 \quad sols_{\{p, q\}}(\mathcal{C}) = \{\{p \mapsto 0, p \mapsto 1\}, \{p \mapsto 0, p \mapsto 2\}\}$$

128 A CSP allowing no solutions is *unsatisfiable*. In CP it is common to use an *objective*  
 129 function to quantify the quality of a solution. A *Constraint Optimization Problem (COP)*  
 130 is a quadruple  $(\mathcal{X}, \mathcal{D}, \mathcal{C}, f)$  with  $f$  an function that maps to a numeric value. An *optimal*

Hélène: TODO:  
 update the ref  
 if the list of  
 accepted paper  
 at CPAIOR is  
 out before sub-  
 mission

<sup>1</sup> <https://satcompetition.github.io/>

## 23:4 Mutational Fuzz-testing for Constraint Modeling Systems

131 solution is a solution to the COP such that no solution exists with a lower/higher objective  
 132 value for minimization/maximization problems.

133 *Global constraints* are one of the essential features of constraint programming and capture  
 134 high-level relations between a (non-fixed) number of variables [34]. Well-known examples of  
 135 global constraints are the ALLDIFFERENT [33] constraint or the CUMULATIVE [1] constraint.  
 136 More examples can be found in the global constraint catalog [6].

137 Typically, constraints and objectives are represented by expressions in some formal syntax.  
 138 E.g., the constraint  $\neg\text{ALLDIFFERENT}(x_1, x_2 + x_3, \text{MAX}(x_4, 0))$  maps those assignments to  
 139 true where  $x_1, x_2 + x_3$ , and the maximum of  $x_4$  and zero do not all take different values.  
 140 Equivalently, constraints can be inductively defined as expression trees. Its leaves are variables  
 141 or values. Its non-leaf nodes are formed by applying *operators*, *global constraints*, *functions*,  
 142 and *comparisons* to other expressions. The expression tree representing the previously  
 143 mentioned complex expression is shown in Figure 1a.

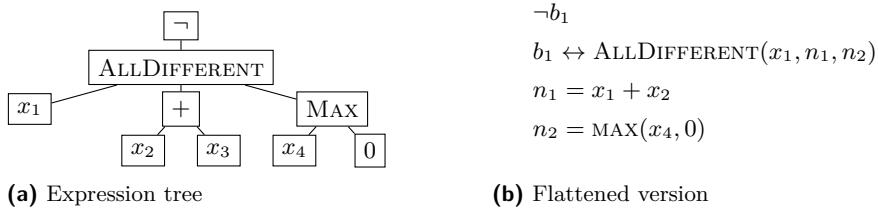


Figure 1 Expression tree and flattened version of  $\neg\text{ALLDIFFERENT}(x_1, x_2 + x_3, \text{MAX}(x_4, 0))$

### 144 3.1 Solvers and modeling systems

145 CSPs are solved by *constraint solvers*: highly optimized combinatorial search systems that  
 146 accept a set of constraints and return (optimal) solutions or report that none exist. Constraint  
 147 solvers do not accept arbitrary expression trees as constraints. Instead, they have a restricted  
 148 input and rarely a solver would accept a complex expression like the one given in Figure 1a  
 149 as an input constraint.

150 Instead of having to manually transform a problem to the format of each solver, a *model*  
 151 and *solve* approach is used, where a user specifies the constraints in an expressive, high-level  
 152 *modeling language*. Then, an underlying compiler translates these constraints to simpler,  
 153 low-level constraints that are passed to a solver. The translation involves multiple complex  
 154 *transformation* steps, with *flattening* (unnesting of nested expressions) and *global constraint*  
 155 *decomposition* (decomposition of unsupported global constraint) as notable examples [27].  
 156 Because different solvers can accept different inputs, distinct transformation paths are  
 157 necessary for different solvers. When using MIP solvers the constraints have to be linearised  
 158 into mixed integer linear inequalities, for SAT solvers only propositional clauses should be  
 159 left, or for CP solvers non-nested constraints over variables, where global constraints that  
 160 are not supported are decomposed.

161 ► **Example 1 (Flattening).** In Figure 1b, we show the flattened version of the expression  
 162  $\neg\text{ALLDIFFERENT}(x_1, x_2 + x_3, \text{MAX}(x_4, 0))$ . The flat output is constructed by traversing the  
 163 expression tree in Figure 1a and introducing auxiliary variable  $n_1, n_2$  and  $b_1$  for every non  
 164 leaf-node.  $n_1$  and  $n_2$  are numerical variables while  $b_1$  is Boolean.

165 **3.2 CPMpy**

166 As a concrete modeling system, we will use CPMpy [15], a constraint modeling library  
 167 embedded in the Python programming language. It translates high-level expressions written  
 168 by a user, to different constraint solvers using a sequence of generic *transformations*. Multiple  
 169 solvers are supported, including CP, SAT, MIP, SMT and Pseudo-Boolean solvers.

170 CPMpy's input language allows arithmetic operations ( $+, -, /, \times \dots$ ), comparisons ( $=$   
 171  $, \neq, <, >, \leq, \geq$ ), logical operations ( $\neg, \wedge, \vee, \rightarrow, \oplus$ ), functions (MAX, COUNT, ABS  $\dots$ ) and  
 172 global constraints (ALLDIFFERENT, CUMULATIVE  $\dots$ ). Expressions in CPMpy are either  
 173 of Boolean or integer type. With  $\mathcal{B}$  we denote the Boolean expressions, with  $\mathcal{N}$  the integer  
 174 ones. Any Boolean expression in CPMpy can also be used as an integer expression (with  
 175 true treated as 1 and false as 0). In other words,  $\mathcal{B} \subseteq \mathcal{N}$ .

176 CPMpy allows users to arbitrarily nest expressions. For example, a disjunction can be  
 177 used as a *constraint* or as an argument to an operator, a function or even a global constraint.  
 178 Similarly, global constraints can be arbitrarily nested and used as any Boolean expression.  
 179 E.g.,  $\text{MAX}(10 \cdot \text{CIRCUIT}(x_1, x_2, x_3), x_1/x_4) \neq 7$  is a valid CPMpy expression. Therefore,  
 180 we avoid the use of the word "constraint" to represent a Boolean expression, as such a  
 181 Boolean expression might be used as a subexpression instead. We use the concept of *top-level*  
 182 *expression* to denote that the expression was given to the solver as a constraint.

183 **4 Mutational testing**

184 We now introduce HURRICANE, a framework for *mutational fuzz testing* of constraint  
 185 modeling systems, inspired by the STORM [20] and YinYang [38] systems for testing SMT-  
 solvers. A high-level overview is shown in Algorithm 1.

**Algorithm 1** HURRICANE

---

**Input:** set of  $m$  CSP models  $\{(\mathcal{X}_j, \mathcal{D}_j, \mathcal{C}_j)\}$ , set of mutations  $\mathcal{M}$  and  $n$ , a number of mutations to apply to each instance

```

1 while true do
2    $(\mathcal{X}, \mathcal{D}, \mathcal{C}) \leftarrow$  pick an instance from the input set
3   for  $i = 1 \dots n$  do
4      $M \leftarrow$  pick a mutation from  $\mathcal{M}$ 
5      $\mathcal{C} \leftarrow \mathcal{C} \cup M(\mathcal{C})$ 
6   if  $\text{verify}(\mathcal{C})$  does not succeed then
7     yield bug with constraints  $\mathcal{C}$ 

```

---

186 Our method takes as input a set of  $m$  constraint satisfaction or optimization problems  
 187 that are known to be satisfiable. In each iteration of the algorithm, we randomly pick one of  
 188 the models and apply a number of *mutations* to its constraints. A mutation is a function  
 189  $M$  that takes as input as set of constraints and outputs a set of new constraints  $M(\mathcal{C})$ . We  
 190 investigate different mutations in Section 5. These newly generated constraints are then  
 191 *added to the model*. Notice this allows us to generate weaker constraints without altering  
 192 the set of solutions of the model. After applying these mutations, we *verify* whether the  
 193 resulting set of constraints satisfies certain properties, e.g., whether the mutated model is  
 194 still satisfiable. Whenever this check fails, the algorithm has found a bug and this is logged  
 195 to the user. In Section 5 we investigate different types of mutations to use and Section 6  
 196 discusses the methods that can be used in order to verify the mutated models.

## 23:6 Mutational Fuzz-testing for Constraint Modeling Systems

198 As our algorithm involves several random components, it is common to (re-)discover the  
199 same error or bug in multiple ways. In an attempt to minimize this to some extent, we  
200 exclude any mutation-model pairs which have already produced a bug, without showing this  
201 explicitely in the pseudocode.

### 202 Input models

203 To construct a varied dataset of feasible input models, we suggest to extract the constraint  
204 models used for the unit tests of the given modeling language. From a practical point of  
205 view, this is useful as unit test models are readily available and kept up-to-date. Many of  
206 models used in unit tests also tend to be small and hence fast to solve. Moreover, unit  
207 tests are highly diverse and it is reasonable to assume these models will contain all language  
208 constructs (such as global constraints and functions). Finally, additional test cases are often  
209 added to the unit tests as part of a bug-fix, hence a fix is tested more rigorously by applying  
210 fuzz-testing on the newly added test-model too.

211 Throughout this paper we use the following input model as a running example.

212 ► **Example 2** (Running example). Consider the following constraint satisfaction problem with  
integer variables  $x, y, z, p$  and  $q$  with domains  $[1..5]$  and a Boolean variable  $b$ .

$$\text{ALLDIFFERENT}(x, y, z), \quad y + \text{MIN}(p, q) > 3, \quad 2 \cdot (x + p) \leq 7$$

213  
214

## 215 5 Mutations

216 We consider three families of mutations. The first of which are based on the reformulation  
217 methods built into constraint modeling systems such as flattening or linearization of con-  
218 straints. Second, we focus on *top-level* mutations which combine existing *top-level expressions*  
219 to create a new expression, and lastly, we consider *sub-expression-level* mutations which can  
220 replace nodes at arbitrary depth in the expression tree. All of these mutations generate  
221 constraints which do not disallow any of the solutions of original constraints. Because  
222 we also leave the original constraints in the mutated model (see Algorithm 1), this means  
223 the set of solutions projected to the original variables should remain unchanged after any  
224 mutation. This property of our mutations is exploited in Section 6 in order to verify the  
225 output of the modeling system after mutating the constraint model.

### 226 5.1 Reformulation mutations

227 Constraint modeling systems implement *reformulation* methods in order to rewrite constraints  
228 into semantically equivalent ones. For example, when a modeling system interfaces a MIP  
229 solver, it implements some procedure to *linearize* constraints. That is, to rewrite any  
230 constraint into weighted sums and linear comparisons. Similarly, CP modeling systems  
231 *decompose* unsupported global constraints or *flatten* complex expression trees.

232 CPMpy provides this functionality as standalone *transformation functions* which take  
233 as input a set of constraints and output a set of (simpler) constraints that imply the input  
234 constraints.<sup>2</sup> As these transformations are supposed to create sets of constraints that leave

---

<sup>2</sup> <https://github.com/CPMpy/cmpy/tree/master/cmpy/transformations>

235 the solutions of the CSP unaltered, we can directly use them as mutations in the mutational  
 236 testing framework. By using these transformation functions, we are able to test these core  
 237 components of the modeling language on a wide range of expressions, even if the backend  
 238 solver does not require that specific transformation. The full list of the transformation  
 239 functions used and their description can be found in Appendix B.

240 **5.2 Top-level mutations**

241 The first set of mutations we use in our framework is based on logical operations with the  
 242 main idea being the following: given two Boolean expressions from the *top-level* of the  
 243 constraint model, combine them to create an implied expression. As both input expressions  
 244 will be enforced to be satisfied by the constraint solver, the newly generated expressions do  
 245 not alter the set of solutions when added to the model and can be considered *redundant*.

246 We compile a set of top-level mutations as summarized in Section 5.2. They are inspired  
 247 by the mutations described in [20] and derived from the truth table of the logical operation  
 248 relation whose name is shown as subscript in the function descriptions below. We repeat  
 249 that these operations are only done on top-level constraints, so they are all implied under  
 250 the condition of  $a \wedge b$  being enforced. Hence, all these constraints can be added to the model  
 251 without changing the set of solutions.

$$252 \quad M_{neg}(a) = \{a, \neg(\neg a)\} \quad (1a)$$

$$253 \quad M_{conj}(a, b) = \{(a \wedge b), \neg(a \wedge \neg b), \neg(\neg a \wedge b), \neg(\neg a \wedge \neg b)\} \quad (1b)$$

$$254 \quad M_{disj}(a, b) = \{(a \vee b), (a \vee \neg b), (\neg a \vee b), \neg(\neg a \vee \neg b)\} \quad (1c)$$

$$255 \quad M_{impli}(a, b) = \{(a \rightarrow b), (\neg a \rightarrow b), (b \rightarrow a), (\neg b \rightarrow a), \\ 256 \quad \quad \quad \neg(a \rightarrow \neg b), (\neg a \rightarrow \neg b), \neg(b \rightarrow \neg a), (\neg b \rightarrow \neg a)\} \quad (1d)$$

$$257 \quad M_{xor}(a, b) = \{(a \oplus \neg b), (\neg a \oplus b), \neg(a \oplus b), \neg(\neg a \oplus \neg b)\} \quad (1e)$$

259 Note that we add all these constraints as is, e.g. we do not simplify  $\neg(a \wedge \neg b)$  to  $(\neg a \vee b)$  but  
 260 leave this expression for future mutations to manipulate further, and for the transformations  
 261 and solvers to handle correctly.

262 Our proposed mutation will randomly pick one of the sets of implied constraints and add  
 263 all those.

264 ► **Example 3.** Given the constraint model shown in Example 2. Imagine HURRICANE  
 265 selects the constraints  $a := \text{ALLDIFFERENT}(x, y, z)$  and  $b := \text{MIN}(p, q) > 3$  and the top-level  
 266 mutation derived from the disjunction operator. Then the following set of constraints is  
 267 generated and added to the model, resulting in a CSP with seven constraints.

$$268 \quad \{(\text{ALLDIFF}(x, y, z)) \vee (2 \cdot (x + p) \leq 7), \quad \neg(\neg \text{ALLDIFF}(x, y, z) \vee \neg(2 \cdot (x + p) \leq 7)), \\ 269 \quad (\neg \text{ALLDIFF}(x, y, z)) \vee (2 \cdot (x + p) \leq 7), \quad (\text{ALLDIFF}(x, y, z)) \vee \neg(2 \cdot (x + p) \leq 7)\}$$

271 **5.3 Subexpression mutations**

272 The mutations described in the previous section operate on top-level Boolean expressions.  
 273 However, we can also modify the expression trees themselves by replacing any of the nodes  
 274 with equivalent ones. Such modified expression trees may trigger different code paths for  
 275 example during flattening of the expression tree before being posted to the solver.

## 23:8 Mutational Fuzz-testing for Constraint Modeling Systems

276 In order to find a set of subexpressions to use for the mutation, we recursively traverse  
 277 the expression tree of each of the constraints. Whenever we find a (sub)expression of the  
 278 required type - e.g., an arithmetic expression - we add the constraint and corresponding  
 279 subexpression to the set of candidates to sample from. Once this set of candidate expressions  
 280 is found, we sample the required amount of expressions to use in the mutation.

281 In the remainder of this section we discuss two types of subexpression mutations.

### 282 Semantic fusion

283 As a way to combine arithmetic sub-expressions, *semantic fusion* was introduced in the  
 284 context of testing SMT-solvers [38]. The key idea is to fuse two numeric expressions and  
 285 create an auxiliary variable for it, and then replace the original expressions by an equivalent  
 286 one involving that variable.

287 In general, semantic fusion requires a *fusion function*  $f(a, b)$  which takes as input two  
 288 arithmetic expressions; an auxiliary variable  $v$  and two *inversion functions*  $r_a(v, b)$  and  
 289  $r_b(v, a)$ . We can then mutate constraints in which  $a$  and  $b$  occur, by replacing the occurrences  
 290 of  $a$  and  $b$  by their now equivalent  $r_a(v, b)$  and  $r_b(v, a)$  expressions.

291 ▶ **Example 4.** We sample two arithmetic subexpressions from the CSP given in Example 2.  
 292 For example, we take  $a := \text{MIN}(p, q)$  and  $b := 2 \cdot (x + p)$ , which are sampled from the second  
 293 and third constraint in the CSP. Using the fusion function  $f(a, b) = a + b$ , we now define a  
 294 new auxiliary variable  $v$  to link the new fused expression as  $v = \text{MIN}(p, q) + 2 \cdot (x + p)$ . We  
 295 can now define a relation from  $a$  to  $b$  and vice versa involving the auxiliary variable. E.g., we  
 296 replace  $\text{MIN}(p, q)$  with  $v - 2 \cdot (x + p)$  and the occurrences of  $2 \cdot (x + p)$  with  $v - \text{MIN}(p, q)$ .

297 This yields the two constraints  $b \leftrightarrow (v - 2 \cdot (x + p)) > 3$  and  $v - \text{MIN}(p, q) \leq 7$  which are  
 298 then added to the model.

299 Multiple operations can be used for the fusion function, even Boolean operators (in which  
 300 case boolean sub-expressions should be selected) though an appropriate inverse function must  
 301 exists. For example  $f(x, y) = x \vee y$  and  $f(x, y) = x \wedge y$  do not allow constructing appropriate  
 inversion functions. In practice, we make use of the fusion functions shown in Table 1

Origin	Fusion Function	Inverse Functions
Sum	$f(a, b) = a + b$	$r_a(v, b) = v - b$ $r_b(v, a) = v - a$
Weighted sum	$f(a, b) = c_1 \cdot a + c_2 \cdot b + c_3$	$r_a(v, b) = (v - c_2 \cdot b - c_3) / c_1$ $r_b(v, a) = (v - c_1 \cdot a - c_3) / c_2$
Subtract	$f(a, b) = a - b$	$r_a(v, b) = v + b$ $r_b(v, a) = a - v$

302 □ **Table 1** Functions which can be used in semantic fusion of arithmetic expressions

302

### 303 Equivalent comparisons

304 The second type of subexpression mutators generates equivalent comparisons. This is done  
 305 by selecting a random comparison in the expression tree of the constraint model and applying  
 306 the same operation to both its sides. These operations can either *add* a constant, *subtract*  
 307 a constant or *apply multiplication* by a constant. The constant itself is picked at random.  
 308 Although this mutation is based on a straightforward idea, we did not find any mention of it  
 309 in literature.

310 ► **Example 5.** Imagine the algorithm picks the second constraint of the running Example 2:  
311  $y + \text{MIN}(p, q) > 3$  and the *multiply by a constant* mutator. If the constant used is “5”, then  
312 applying the mutation results in the expression  $5 \cdot (y + \text{MIN}(p, q)) > 5 \cdot 3$ .

313 The same could in principle be done with a fresh variable or an existing numeric subex-  
314 pression from another constraint, but in this case we just use an integer constant.

## 315 6 Verification methods

316 To detect whether a bug has occurred, we need to *verify* that certain properties hold for the  
317 mutated constraints. In fuzz-testing for SMT research [20, 38, 8], the authors check if after  
318 mutations, the model still admits a solution. However, more elaborate checks are possible as  
319 well. In particular, the mutations presented in Section 5 *should* not alter the set of solutions  
320 projected to the original variables. The verification methods as presented in the following  
321 sections are all methods in order to check whether indeed this set of solutions is preserved.  
322 Different trade-offs between efficiency, code coverage, and thoroughness of the verification  
323 present themselves. We compare and evaluate them experimentally in Section 9.

### 324 6.1 All-solutions

A first method to check the set of solutions is unchanged is to enumerate the solutions of the original model and those of the mutated model and checking for equivalence of solution sets. Some of the mutations presented in Section 5 can introduce auxiliary variables. E.g., semantic fusion introduces a *fusion variable* but also the built-in reformulations such as *flattening* can introduce new variables into the model. Therefore, in order to compare both sets of solutions, we need to project them to the original set of decision variables  $\mathcal{X}$ . I.e., this verification method checks whether the following equivalence holds:

$$sols_{\mathcal{X}}(\mathcal{C}) \equiv sols_{\mathcal{X}}(\mathcal{C} \cup M(\mathcal{C}))$$

325 Note that enumeration of all solutions is a costly operation -  $\#\mathcal{P}$ -complete in general [11]  
326 - but solvers oftentimes have built-in methods for doing so. CPMpy implements enumeration  
327 of all solutions using the `solveAll` function. This in turn calls the built-in enumeration  
328 method of the solver if available, otherwise it implements the enumeration using repeated  
329 solve calls and blocking clauses. Clearly, using this verification method does not only allow  
330 for a theoretically strong verification of the mutations, but can also trigger different code  
331 paths in either the modeling system or the solver itself.

### 332 6.2 Solution count

333 Instead of checking whether *projected* sets of solutions are equivalent, we also want to check  
334 whether new solutions with respect the auxiliary variables are introduced by the mutations.  
335 E.g., if a mutation introduces an unconstrained Boolean auxiliary variable, the total number  
336 of solutions will be doubled. While this behaviour is unwanted for any of the mutations  
337 presented in this paper, it is undetected by the **All-solutions** verification method as the  
338 sets of solutions are projected to the original variables.

Therefore, we propose to also check whether the total number of solutions of the mutated model is unchanged to the original number of solutions. I.e.. we check whether whether

$$|sols(\mathcal{C})| \equiv |sols(\mathcal{C} \cup M(\mathcal{C}))|$$

339     Similar to enumeration of all solutions, counting solutions is also a costly operation, but  
 340     may trigger new code paths in modeling systems or solvers. Note that solution counting and  
 341     checking equivalence of projected solutions sets are complementary to one another. While  
 342     solution counting discovers bugs related to auxiliary variables, **All-solutions** can discover  
 343     bugs related to assigned values of the decision variables.

### 344 6.3 1-solution

345     Instead of checking whether *all* solutions remain for the mutated constraints, we can check  
 346     whether a predefined solution is preserved by the mutations. In practice, we implement this by  
 347     adding the assignment of a pre-computed solution to the set of mutated constraints and check  
 348     if the resulting constraints are satisfiable. E.g., for the CSP from Example 2, we can test if after  
 349     mutation of the constraints, the assignment  $\{b \mapsto \text{false}, x \mapsto 2, y \mapsto 3, z \mapsto 1, p \mapsto 2, q \mapsto 1\}$   
 350     is still a solution of the CSP. Conceptually, we check for a given solution  $\theta$  whether

$$351 \quad \theta \in \text{sols}(\mathcal{C} \cup M(\mathcal{C}))$$

351     Notice that verifying if a solution satisfies a set of constraints is polynomial in time as  
 352     the solver does not require any search when all variables are fixed! Naturally, finding the  
 353     pre-computed solution for the original CSP requires invoking a solver nevertheless.

354     We expect this method to detect similar changes to the set of solutions such as the  
 355     **All-solutions**, while avoiding the enumeration of all solutions.

### 356 6.4 Satisfiability

357     Instead of checking whether a predefined assignment is a solution of the mutated model, we  
 358     can also check whether the mutated model admits a solution at all. This verification method  
 359     is similar to the work on fuzz-testing SMT-solvers [20, 38, 8]. Naturally, this check does not  
 360     detect subtle changes in the set of solutions of the mutated model, but rather checks if the  
 361     sets of solutions is non-empty.

### 362 6.5 Optimization

363     In constraint programming, it is common to use an objective function in order to quantify the  
 364     quality of a solution. E.g., when scheduling a set of tasks on a machine, it is common to find  
 365     a schedule which runs in the least amount of time or requires the smallest amount of energy.  
 366     When such an objective function is set in a constraint model, we can check whether solving  
 367     the mutated model to optimality yields the same objective value. That is, we check whether  
 368     at least one of the optimal solutions is still an optimal solution of the mutated model.

369     While this check is conceptually stronger compared to checking the satisfiability check  
 370     presented in Section 6.4, it has two disadvantages. Firstly, it requires the existence of  
 371     an objective function in the model and secondly, finding an optimal solution to a CSP is  
 372     conceptually harder, and hence will take more time, compared to finding any satisfying  
 373     solution to the constraints.

## 374 7 Dealing with bugs

375     Computer programs can exhibit several types of bugs. Similar to the authors of [20], we  
 376     define three classes for bugs to occur in constraint modeling languages. Section 7.1 discusses  
 377     errors in the logic of modeling systems and solvers, while Section 7.2 and Section 7.3 focus

378 on bugs which impact the runtime environment of modeling systems. Lastly, in Section 7.4,  
379 we discuss a practical method to find minimal examples of when a bug occurs.

### 380 7.1 Soundness bugs

381 The first type of bug are those where the modeling system are detected when the modeling  
382 system returns a wrong answer to a verification check from Section 6. Such bugs are critical  
383 as the user is given a wrong answer to the constraints, while the modeling system seems to  
384 run as normal. E.g., the solver returns a non-optimal solution to an optimization problem or  
385 declares a set of constraints to be unsatisfiable when in fact they admit a solution.

386 Soundness bugs can be caused by either the solver itself, or by the modeling system.  
387 In the case where the root-cause of the bug lies in the solver, an example can be when a  
388 propagation function for a (global) constraint removes values from a domain which allowed a  
389 solution. When the bugs is caused by the modeling system an example is flawed interface to  
390 the solver or an improper reformulation of the constraints.

391 Overall, soundness bugs are critical but difficult to detect in day-to-day use of a modeling  
392 language as their use rarely includes verifying the result in a later stage.

### 393 7.2 Crashes

394 During the execution of HURRICANE, it is possible the runtime of the modeling system  
395 crashes. We identify two main points of possible failure: applying a mutation and verifying  
396 the mutated model.

397 We noticed crashes or errors occurring during the mutation of set of constraints are often  
398 triggered when a reformulation mutation is chosen. For example, during linearization of a set  
399 of constraints, an assertion error was thrown because certain edge cases were not covered.

400 When a crash occurs during verification of the set of mutated constraints, this can  
401 be caused by either the backend solver or the modeling system. For example, during the  
402 development of our tool, a crash in a solver was caused by an integer overflow error - causing  
403 the solver to return an error message. An example when CPMpy was identified to be the  
404 cause of a crash happened when one of the interfaces to a solver did not implement all  
405 primitive constraints properly.

406 Most crashes are easy to detect in the day-to-day use of modeling systems as a user  
407 always receives an error message. Still, the severity of a crash can vary widely as it mostly  
408 depends on how the system is used. E.g., when the modeling system crashes when used in  
409 an integrated system of a manufacturing plant, the crash has likely far greater implications  
410 compared when it is used in an interactive session.

### 411 7.3 Performance issues

412 The last type of bugs we identified are related to the performance and efficiency of the library.  
413 For example, when we verify whether the mutated model satisfies at least one solution,  
414 the time it takes for the modeling system to receive an answer from the solver may be  
415 significantly higher compared to the original model. This can again have several reasons  
416 caused by either the modeling system or the solver. For example, the mutated model may  
417 contain global constraints which get decomposed in a particularly inefficient way when nested  
418 by HURRICANE. Sometimes, either the solver or modeling system may even get stuck in an  
419 infinite loop! In practice we overcome this by setting a hard time-limit on the call to the  
420 verification method. Naturally, this may trigger false-positives as the mutated model may

## 23:12 Mutational Fuzz-testing for Constraint Modeling Systems

421 simply be harder to solve due to the surplus in variables and constraints. Still, we log these  
422 bugs as it may uncover interesting inefficiencies in the code.

### 423 7.4 Minimizing buggy models

424 The mutations defined in this paper can result in very large and deeply nested constraint  
425 models. However, often only a (small) subset of the constraints are the root cause of the  
426 bug. In our work, we utilize a simple deletion-based method that iteratively removes a single  
427 constraint from the model as long as the remaining model exhibits the bug. This method  
428 is similar to delta-debugging and is often used in combination with fuzz-testing [40]. It  
429 should be noted that a crash of the system often gives some sort of message pointing to  
430 the expressions that caused the crash. Therefore, we deem delta debugging to be especially  
431 useful when dealing with a soundness bug.

432 Another way to simplify the debugging process is by automatically detecting bugs that  
433 are already identified. HURRICANE will keep logging a bug until it is fixed, so the same bug  
434 will be logged many times over. A first way to find out which of the bugged models are cause  
435 by Bug X, is to fix Bug X and then simply check which buggy models do no longer exhibit a  
436 bug. It's ofcourse not always possible to quickly fix a bug, even after it is identified. We  
437 then resort to matching the error messages and location of the error in the code, as well as  
438 the input model or transformation that lead to the bug. For soundness bugs we can compare  
439 the results of multiple solvers to see if they match. This is enough information to confidently  
440 categorise most bugs in a semi-automated process.

### 441 8 Summary of found bugs

442 We coded up HURRICANE in Python 3.11 for CPMpy using the mutations and verification  
443 methods described previously. During development, which covers a period of about 1.5  
444 years, we discovered **52 unique bugs** in total. This includes 19 bugs found in CPMpy  
445 during a master thesis that preceded this work<sup>3</sup>. Out of all bugs discovered, **13 bugs where**  
446 **soundness bugs**, 5 of which had their origin in backend solvers. In particular, we found 2  
447 soundness bugs in the OR-tools solver and three in the MiniZinc system. The vast majority  
448 (29) bugs were crashes of the CPMpy runtime environment. One of these crashes was traced  
449 back to a backend solver crashing. Lastly, we found three performance issues, one of which  
450 was again found in a backend solver.

451 Out of these 52 bugs, 14 remained at the time of the experiments described in the next  
452 section. 6 bugs in backend solvers and 8 in CPMpy. We shortly discuss these bugs in  
453 Appendix A. Full experimental data is also shown there.

### 454 9 Experimental evaluation

455 In this section, we investigate each of the components of our fuzz-testing framework. In  
456 particular, we aim to answer the following experimental questions:

457 **EQ1.** What are the tradeoffs between increasing the number of mutations on each model  
458 and increasing the number of models being tested?

459 **EQ2.** How effective are the different verification methods for finding bugs in constraint  
460 modeling systems?

---

<sup>3</sup> Reference temporarily omitted for anonymity

461 **EQ3.** To what extend improves fuzz-testing the overall coverage of tested code compared to  
 462 CPMpy's builtin suite of unit tests?

463 We configure HURRICANE to use different numbers of mutations and different types of  
 464 verification methods. We test each of the five verification methods described in Sections 6.4 -  
 465 6.5 separately. For each of the verification methods, we employ four numbers of mutations  
 466 applied to the input model before verification:  $n = \{1, 2, 5, 10\}$ . As backend solvers, we test  
 467 the OR-Tools CP-SAT solver v.9.9 and MiniZinc v.2.8.3 with Gecode version 6.3.0. This  
 468 combination of settings results in a total of 40 configurations, each of which was ran for 10  
 469 hours on an Ubuntu 20.04.6 LTS machine with an Intel Core i7-2600 CPU@3.40Ghz and  
 470 16GB of RAM. During these experiments, we keep track of which lines in the code-base are  
 471 executed using the `coverage` utility in Python.

472 We used 1240 constraint models as input, 7 of which are optimization problems. As  
 473 discussed in Section 4, the models were extracted from the unit tests of CPMpy. All code  
 474 and experimental data will be made available upon acceptance of this paper. In the following  
 475 sections, we aggregate the results of the above evaluation in order to answer the experimental  
 476 questions.

## 477 **9.1 EQ1: effect of number of mutations**

478 In this first experiment, we investigate the influence of the number of mutations ( $n$ ) used  
 479 in Algorithm 1 before verifying the mutated models. The more mutations used, the more  
 480 diverse the output can be, and the more likely it is for a bug to be found. This can clearly  
 481 be seen from the `#unique` column in Table 2 where we notice a steady increase in number of  
 482 unique bugs found, with respect to the number of applied mutations. Notice this number  
 483 of unique bugs is not in direct correlation with the number of errors reported. E.g., when  
 484 testing OR-Tools and using two mutations before verification, many errors with the same  
 485 root-cause (bug) are found by HURRICANE.

486 Mutations can increase the size of a model hyperlinearly: when applying a transformation  
 487 such as flattening or decomposing global constraints, a single constraint can easily become a  
 488 large set of constraints. Hence, it is likely the subsequent mutations will be slower as they  
 489 have to run on bigger input, as does the verification check. From the `#models` column in  
 490 Table 2, we can indeed conclude more mutations will result less models tested for the given  
 491 time-frame of ten hours.

492 The optimal value for  $n$  will of course depend on the time HURRICANE is ran for, since  
 493 for smaller  $n$  we can find bugs more quickly, but for big  $n$  we expect to find those bugs  
 494 *eventually*. We therefore propose that the best way of using HURRICANE would be to  
 495 increase  $n$  over time, causing the easily detected bugs to get found quickly while making it  
 496 possible to find the more obscure bugs later on.

■ **Table 2** Number of mutations for each iteration compared to the number of bugs found and  
 number of models handled. (Aggregated over the different verification methods)

#mutations	OR-Tools			MiniZinc			Total #unique
	#models	#errors	#unique	#models	#errors	#unique	
1	9166418	5747	1	218377	289	3	3
2	6672588	11002	3	216527	723	6	6
5	2270441	8975	5	128884	1495	8	11
10	344710	2783	7	57191	423	9	13

497 **9.2 EQ2: effect of verification methods**

498 The next dimension of our algorithm we investigate is the different types of verification  
 499 methods. We aggregate the results for this experiments for all number of mutations. I.e., the  
 500 results as reported in Table 3 result from testing the algorithm with all settings of  $n$ .

501 First of all we notice a big difference in the amount of models that the different methods can  
 502 verify. The results for the optimization verification method should be interpreted cautiously,  
 503 because they run on a smaller subset of input models that have an objective function. These  
 504 models happen to be small, explaining why the optimization verification solves more models  
 505 than we would expect it to. More interesting is the difference in the number of models checked  
 506 for the satisfiability and 1-solution verifications compared to counting and equivalence. This  
 507 however does not translate to a large advantage in discovered bugs, indicating the usefulness  
 508 of the computationally more expensive counting and equivalence verifications.

509 The 1-Solution verification performs best, regarding the number of unique bugs. This can be  
 510 understood because it is a stronger check than the satisfiability check, but seems even faster.  
 511 This is due to the fact that we send the instantiated solution to the solver when verifying  
 512 the mutated model, leading to faster propagation.

513 Interestingly we observe that the solution counting, 1-solution and optimization methods  
 514 all found at least 1 bug that was not detected by any of the other methods. This was not  
 515 the case for All-solutions or satisfiability checking, and we could consider those redundant in  
 516 the context of our experiments. Although verifying All-solutions is theoretically a stronger  
 517 check than solution counting, and they can test models at a similar speed, both methods  
 518 found bugs that the other did not. For example in an earlier experiment a bug was found in  
 519 the `solveAll` routine of CPMpy, only detected using solution count. This highlights the  
 520 advantage of using different verification methods to cover all aspects of the toolchain.

521 **Table 3** Number of verification steps and errors found for different verification methods in 40  
 522 hours. (Aggregated over the different values of  $n$ )

verification	OR-Tools			MiniZinc			Total #unique
	#models	#errors	#unique	#models	#errors	#unique	
All sol	13441	460	4	11167	312	7	8
Counting	14551	539	5	11623	325	6	8
One sol	4095185	25695	5	194495	1983	8	10
Sat	3679400	180	4	186119	116	5	8
Opt	10651580	1633	2	217575	194	3	4

521 **9.3 EQ3: effect on code coverage**

522 As mentioned in Section 2, code coverage is a common proxy to measure the efficacy of  
 523 a test suite. In this experiment, we compare the code coverage of running all unit test  
 524 *models* (**unit-models**), running HURRICANE for 400 hours with these unit test models  
 525 (200 hours for each backend solver) (**HURRICANE**), running all unit tests (not just the  
 526 models that appear in them) (**unit-tests**), and the combined code coverage (**combined**) of  
 527 **HURRICANE** and **unit-tests**.

528 The results are presented in Table 4. The rows in this table are split on the different  
 529 solvers, with each subrow representing a part of the code base. *expressions* contains the  
 530 construction and evaluation code for all expressions (operators, functions, global constraints,  
 531 etc.), *transformations* the internal transformation routines, and *ortools.py* and *minizinc.py*  
 532 contain the solver-specific interfacing code

533 The results show that HURRICANE improves code coverage over just solving the unit  
 534 models, but not over running all unit tests. Still, HURRICANE does cover new parts of  
 535 the code, as the combined coverage is higher than just unit tests on its own. Because  
 536 HURRICANE uses the internal transformations as mutations, we see a high code coverage  
 537 on *transformations* too, even when using a solver like MiniZinc that requires only few of the  
 538 transformations in CPMpy.

■ **Table 4** Segmented code coverage for different components of CPMpy

Solver	files	unit-models	HURRICANE	unit-tests	combined
OR-Tools	<i>expressions</i>	54.6%	<b>64.6%</b>	87.3%	<b>88.6%</b>
	<i>transformations</i>	59.3%	<b>83.6%</b>	86.4%	<b>88.2%</b>
	<i>ortools.py</i>	64.1%	<b>81.5%</b>	90.4%	<b>91.5%</b>
MiniZinc	<i>expressions</i>	51.1%	<b>64.0%</b>	87.3%	<b>88.6%</b>
	<i>transformations</i>	22.1%	<b>82.6%</b>	86.4%	<b>88.2%</b>
	<i>minizinc.py</i>	70.6%	<b>84.3%</b>	83.0%	<b>89.2%</b>

## 539 10 Discussion and future work

540 We presented a method to automatically test constraint modeling languages given a set  
 541 of input CSPs and COPs. We show that a sufficiently diverse set of input models can be  
 542 obtained from the unit tests of the modeling language. Based on recent work in SMT-testing,  
 543 we proposed a set of mutations to use over these models, in order to generate new and more  
 544 complex inputs to the modeling language.

545 As shown in Section 9, our method is able to find a significant number of bugs for the  
 546 CPMpy framework and its solvers, ranging from crashes to soundness bugs and finding  
 547 downstream bugs in MiniZinc and OR-Tools. Moreover, using our framework improves  
 548 the code coverage compared to the unit testing implemented in the library. Our proposed  
 549 fuzz testing techniques also neatly allow *continuous integration* with modeling language  
 550 development: when new features and bug fixes are added to a modeling language, the fuzz  
 551 testing framework can just continue with the latest version on some remote server, testing  
 552 the codebase 24/7.

553 While our methods are highly effective in finding bugs, one of the major difficulties  
 554 remains how to avoid re-finding similar bugs, and producing minimal bug instances. We  
 555 leave this topic for future investigation. Compared to testing SMT-solvers, CP offers several  
 556 interesting dimensions on which we only briefly touched in this paper. These features include  
 557 optimization, which can be tested more thoroughly in the future by also mutating objective  
 558 functions. Another key feature of CP is the notation of global constraints. Based on [8], we  
 559 would like to include mutations which can introduce *new* global constraints into the models  
 560 as currently we rely on the global constraints already being present in the input.

561 Recent work in SMT-solving showcases the power of using voting between multiple solvers  
 562 to verify the answer any of the solvers produce [37]. Crucially, solver voting allows to use  
 563 mutations where the result of the solver does not have to be known upfront, i.e., one does  
 564 not have to know what properties the mutations have. Using multiple solvers perfectly suits  
 565 the testing of constraints modeling languages, as their core function is to translate constraint  
 566 specifications to multiple solvers and solving paradigms. We are optimistic that this work  
 567 will remain useful in the future, by applying it to more solvers, adding more mutations, and  
 568 encouraging more developers to make use of it.

---

569 **References**

570 1 Abderrahmane Aggoun and Nicolas Beldiceanu. Extending CHIP in order to solve complex  
571 scheduling and placement problems. In Jean-Paul Delahaye, Philippe Devienne, Philippe  
572 Mathieu, and Pascal Yim, editors, *JFPL'92, 1ères Journées Francophones de Programmation  
573 Logique, 25-27 Mai 1992, Lille, France*, page 51, 1992.

574 2 Özgür Akgün, Alan M. Frisch, Ian P. Gent, Christopher Jefferson, Ian Miguel, and Peter  
575 Nightingale. Conjure: Automatic generation of constraint models from problem specifications.  
576 *Artif. Intell.*, 310:103751, 2022. doi:10.1016/j.artint.2022.103751.

577 3 Özgür Akgün, Ian P. Gent, Christopher Jefferson, Ian Miguel, and Peter Nightingale. Meta-  
578 morphic testing of constraint solvers. In John N. Hooker, editor, *Principles and Practice of  
579 Constraint Programming - 24th International Conference, CP 2018, Lille, France, August  
580 27-31, 2018, Proceedings*, volume 11008 of *Lecture Notes in Computer Science*, pages 727–736.  
581 Springer, 2018. doi:10.1007/978-3-319-98334-9\\_46.

582 4 Mario Alviano, Carmine Dodaro, Johannes Klaus Fichte, Markus Hecher, Tobias Philipp, and  
583 Jakob Rath. Inconsistency proofs for ASP: the ASP - DRUPE format. *Theory Pract. Log.  
584 Program.*, 19(5-6):891–907, 2019. doi:10.1017/S1471068419000255.

585 5 Haniel Barbosa, Andrew Reynolds, Gereon Kremer, Hanna Lachnitt, Aina Niemetz, Andres  
586 Nötzli, Alex Ozdemir, Mathias Preiner, Arjun Viswanathan, Scott Viteri, Yoni Zohar, Cesare  
587 Tinelli, and Clark W. Barrett. Flexible proof production in an industrial-strength SMT solver.  
588 In Jasmin Blanchette, Laura Kovács, and Dirk Pattinson, editors, *Automated Reasoning - 11th  
589 International Joint Conference, IJCAR 2022, Haifa, Israel, August 8-10, 2022, Proceedings*,  
590 volume 13385 of *Lecture Notes in Computer Science*, pages 15–35. Springer, 2022. doi:  
591 10.1007/978-3-031-10769-6\\_3.

592 6 Nicolas Beldiceanu, Mats Carlsson, and Jean-Xavier Rampon. Global constraint catalog,  
593 (revision a), 2012.

594 7 Bart Bogaerts, Stephan Gocht, Ciaran McCreesh, and Jakob Nordström. Certified symmetry  
595 and dominance breaking for combinatorial optimisation. In *Thirty-Sixth AAAI Conference  
596 on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications  
597 of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in  
598 Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 3698–3707.  
599 AAAI Press, 2022. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/20283>.

600 8 Mauro Bringolf. Fuzz-testing smt solvers with formula weakening and strengthening. Master's  
601 thesis, ETH Zurich, 2021.

602 9 Robert Brummayer and Armin Biere. Fuzzing and delta-debugging smt solvers. In *Proceedings  
603 of the 7th International Workshop on Satisfiability Modulo Theories*, pages 1–5, 2009.

604 10 Robert Brummayer, Florian Lonsing, and Armin Biere. Automated testing and debugging of  
605 SAT and QBF solvers. In Ofer Strichman and Stefan Szeider, editors, *Theory and Applications  
606 of Satisfiability Testing - SAT 2010, 13th International Conference, SAT 2010, Edinburgh,  
607 UK, July 11-14, 2010. Proceedings*, volume 6175 of *Lecture Notes in Computer Science*, pages  
608 44–57. Springer, 2010. doi:10.1007/978-3-642-14186-7\\_6.

609 11 Nadia Creignou and Miki Hermann. Complexity of generalized satisfiability counting problems.  
610 *Inf. Comput.*, 125(1):1–12, 1996. doi:10.1006/inco.1996.0016.

611 12 Ermira Daka and Gordon Fraser. A survey on unit testing practices and problems. In *25th  
612 IEEE International Symposium on Software Reliability Engineering, ISSRE 2014, Naples,  
613 Italy, November 3-6, 2014*, pages 201–211. IEEE Computer Society, 2014. doi:10.1109/ISSRE.  
614 2014.11.

615 13 Stephan Gocht, Ciaran McCreesh, and Jakob Nordström. Veripb: The easy way to  
616 make your combinatorial search algorithm trustworthy. In *workshop From Constraint  
617 Programming to Trustworthy AI at the 26th International Conference on Principles and  
618 Practice of Constraint Programming (CP'20)*. Paper available at [http://www.cs.ucc.ie/bg6/cptai/2020/papers/CPTAI\\_2020\\_paper\\_2.pdf](http://www.cs.ucc.ie/bg6/cptai/2020/papers/CPTAI_2020_paper_2.pdf), 2020.

620 14 Stephan Gocht, Ciaran McCreesh, and Jakob Nordström. An auditable constraint programming  
621 solver. In Christine Solnon, editor, *28th International Conference on Principles and Practice*  
622 of *Constraint Programming, CP 2022, July 31 to August 8, 2022, Haifa, Israel*, volume 235 of *LIPICS*, pages 25:1–25:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.  
623 [doi:10.4230/LIPIcs.CP.2022.25](https://doi.org/10.4230/LIPIcs.CP.2022.25).

625 15 Tias Guns. Increasing modeling language convenience with a universal n-dimensional array,  
626 cppy as python-embedded example. In *Proceedings of the 18th workshop on Constraint*  
627 *Modelling and Reformulation at CP (Modref 2019)*, volume 19, 2019.

628 16 Marijn J. H. Heule. Proofs of unsatisfiability. In Armin Biere, Marijn Heule, Hans van  
629 Maaren, and Toby Walsh, editors, *Handbook of Satisfiability - Second Edition*, volume 336  
630 of *Frontiers in Artificial Intelligence and Applications*, pages 635–668. IOS Press, 2021.  
631 [doi:10.3233/FAIA200998](https://doi.org/10.3233/FAIA200998).

632 17 Ahmet B. Keha, Ketan Khowala, and John W. Fowler. Mixed integer programming formulations  
633 for single machine scheduling problems. *Comput. Ind. Eng.*, 56(1):357–367, 2009. [doi:10.1016/j.cie.2008.06.008](https://doi.org/10.1016/j.cie.2008.06.008).

635 18 Shuvendu K. Lahiri and Shaz Qadeer. Back to the future: revisiting precise program verification  
636 using SMT solvers. In George C. Necula and Philip Wadler, editors, *Proceedings of the*  
637 *35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL*  
638 *2008, San Francisco, California, USA, January 7-12, 2008*, pages 171–182. ACM, 2008.  
639 [doi:10.1145/1328438.1328461](https://doi.org/10.1145/1328438.1328461).

640 19 Stephan Lukasczyk, Florian Kroiß, and Gordon Fraser. Automated unit test generation for  
641 python. *CoRR*, abs/2007.14049, 2020. URL: <https://arxiv.org/abs/2007.14049>, arXiv:  
642 2007.14049.

643 20 Muhammad Numair Mansur, Maria Christakis, Valentin Wüstholtz, and Fuyuan Zhang.  
644 Detecting critical bugs in SMT solvers using blackbox mutational fuzzing. In Prem Devanbu,  
645 Myra B. Cohen, and Thomas Zimmermann, editors, *ESEC/FSE '20: 28th ACM Joint*  
646 *European Software Engineering Conference and Symposium on the Foundations of Software*  
647 *Engineering, Virtual Event, USA, November 8-13, 2020*, pages 701–712. ACM, 2020. [doi:10.1145/3368089.3409763](https://doi.org/10.1145/3368089.3409763).

649 21 Matthew J. McIlree and Ciaran McCreesh. Proof logging for smart extensional constraints.  
650 In Roland H. C. Yap, editor, *29th International Conference on Principles and Practice of*  
651 *Constraint Programming, CP 2023, August 27-31, 2023, Toronto, Canada*, volume 280 of  
652 *LIPICS*, pages 26:1–26:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. URL:  
653 <https://doi.org/10.4230/LIPIcs.CP.2023.26>, [doi:10.4230/LIPICS.CP.2023.26](https://doi.org/10.4230/LIPICS.CP.2023.26).

654 22 Barton P. Miller, Lars Fredriksen, and Bryan So. An empirical study of the reliability of UNIX  
655 utilities. *Commun. ACM*, 33(12):32–44, 1990. [doi:10.1145/96267.96279](https://doi.org/10.1145/96267.96279).

656 23 Nicholas Nethercote, Peter J. Stuckey, Ralph Becket, Sebastian Brand, Gregory J. Duck, and  
657 Guido Tack. Minizinc: Towards a standard CP modelling language. In Christian Bessiere,  
658 editor, *Principles and Practice of Constraint Programming - CP 2007, 13th International*  
659 *Conference, CP 2007, Providence, RI, USA, September 23-27, 2007, Proceedings*, volume  
660 4741 of *Lecture Notes in Computer Science*, pages 529–543. Springer, 2007. [doi:10.1007/978-3-540-74970-7\\_38](https://doi.org/10.1007/978-3-540-74970-7_38).

662 24 Peter Nightingale, Özgür Akgün, Ian P. Gent, Christopher Jefferson, and Ian Miguel. Auto-  
663 matically improving constraint models in savile row through associative-commutative common  
664 subexpression elimination. In Barry O’Sullivan, editor, *Principles and Practice of Constraint*  
665 *Programming - 20th International Conference, CP 2014, Lyon, France, September 8-12, 2014.*  
666 *Proceedings*, volume 8656 of *Lecture Notes in Computer Science*, pages 590–605. Springer,  
667 2014. [doi:10.1007/978-3-319-10428-7\\_43](https://doi.org/10.1007/978-3-319-10428-7_43).

668 25 Peter Nightingale, Patrick Spracklen, and Ian Miguel. Automatically improving SAT en-  
669 coding of constraint problems through common subexpression elimination in savile row.  
670 In Gilles Pesant, editor, *Principles and Practice of Constraint Programming - 21st Inter-*  
671 *national Conference, CP 2015, Cork, Ireland, August 31 - September 4, 2015, Proceed-*

672        *ings*, volume 9255 of *Lecture Notes in Computer Science*, pages 330–340. Springer, 2015.  
 673        [doi:10.1007/978-3-319-23219-5\\\_\\\_23](https://doi.org/10.1007/978-3-319-23219-5_23).

674        **26** Tobias Paxian and Armin Biere. Uncovering and classifying bugs in maxsat solvers through  
 675        fuzzing and delta debugging. Update reference when published, 2022. URL: [http://www.pragmaticosfsat.org/2023/live/POS23\\_paper\\_4.pdf](http://www.pragmaticosfsat.org/2023/live/POS23_paper_4.pdf).

676        **27** Andrea Rendl. *Effective compilation of constraint models*. PhD thesis, University of St  
 677        Andrews, UK, 2010. URL: <https://hdl.handle.net/10023/973>.

678        **28** Robert Robere, Antonina Kolokolova, and Vijay Ganesh. The proof complexity of SMT solvers.  
 679        In Hana Chockler and Georg Weissenbacher, editors, *Computer Aided Verification - 30th  
 680        International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC  
 681        2018, Oxford, UK, July 14-17, 2018, Proceedings, Part II*, volume 10982 of *Lecture Notes in  
 682        Computer Science*, pages 275–293. Springer, 2018. [doi:10.1007/978-3-319-96142-2\\\_\\\_18](https://doi.org/10.1007/978-3-319-96142-2_\_18).

683        **29** Francesca Rossi, Peter van Beek, and Toby Walsh, editors. *Handbook of Constraint  
 684        Programming*, volume 2 of *Foundations of Artificial Intelligence*. Elsevier, 2006. URL:  
 685        <https://www.sciencedirect.com/science/bookseries/15746526/2>.

686        **30** Olivier Roussel and Christophe Lecoutre. XML representation of constraint networks: Format  
 687        XCSP 2.1. *CoRR*, abs/0902.2362, 2009. URL: <http://arxiv.org/abs/0902.2362>, arXiv:  
 688        0902.2362.

689        **31** Gilles Simonin, Christian Artigues, Emmanuel Hebrard, and Pierre Lopez. Scheduling scientific  
 690        experiments for comet exploration. *Constraints An Int. J.*, 20(1):77–99, 2015. URL: <https://doi.org/10.1007/s10601-014-9169-3>, [doi:10.1007/S10601-014-9169-3](https://doi.org/10.1007/S10601-014-9169-3).

691        **32** Petar Tahchiev, Felipe Leme, Vincent Massol, and Gary Gregory. *JUnit in Action, 2nd  
 692        Edition*. Manning Publications Company, 2011. URL: <https://www.manning.com/books/junit-in-action-second-edition>.

693        **33** Willem Jan van Hoeve. The alldifferent constraint: A survey. *CoRR*, cs.PL/0105015, 2001.  
 694        URL: <https://arxiv.org/abs/cs/0105015>.

695        **34** Willem-Jan van Hoeve and Irit Katriel. Global constraints. In Francesca Rossi, Peter van Beek,  
 696        and Toby Walsh, editors, *Handbook of Constraint Programming*, volume 2 of *Foundations of  
 697        Artificial Intelligence*, pages 169–208. Elsevier, 2006. [doi:10.1016/S1574-6526\(06\)80010-6](https://doi.org/10.1016/S1574-6526(06)80010-6).

698        **35** Dieter Vandesande, Wolf De Wulf, and Bart Bogaerts. Qmaxsatpb: A certified maxsat solver.  
 699        In Georg Gottlob, Daniela Inclezan, and Marco Maratea, editors, *Logic Programming and  
 700        Nonmonotonic Reasoning - 16th International Conference, LPNMR 2022, Genova, Italy,  
 701        September 5-9, 2022, Proceedings*, volume 13416 of *Lecture Notes in Computer Science*, pages  
 702        429–442. Springer, 2022. [doi:10.1007/978-3-031-15707-3\\\_\\\_33](https://doi.org/10.1007/978-3-031-15707-3_\_33).

703        **36** T.W. Williams, M.R. Mercer, J.P. Mucha, and R. Kapur. Code coverage, what does it mean  
 704        in terms of quality? In *Annual Reliability and Maintainability Symposium. 2001 Proceedings.  
 705        International Symposium on Product Quality and Integrity (Cat. No.01CH37179)*, pages  
 706        420–424, 2001. [doi:10.1109/RAMS.2001.902502](https://doi.org/10.1109/RAMS.2001.902502).

707        **37** Dominik Winterer, Chengyu Zhang, and Zhendong Su. On the unusual effectiveness  
 708        of type-aware operator mutations for testing SMT solvers. *Proc. ACM Program. Lang.*,  
 709        4(OOPSLA):193:1–193:25, 2020. [doi:10.1145/3428261](https://doi.org/10.1145/3428261).

710        **38** Dominik Winterer, Chengyu Zhang, and Zhendong Su. Validating SMT solvers via semantic  
 711        fusion. In Alastair F. Donaldson and Emina Torlak, editors, *Proceedings of the 41st ACM  
 712        SIGPLAN International Conference on Programming Language Design and Implementation,  
 713        PLDI 2020, London, UK, June 15-20, 2020*, pages 718–730. ACM, 2020. [doi:10.1145/3385412.3385985](https://doi.org/10.1145/3385412.3385985).

714        **39** Hui Ye, Shaoyin Cheng, Lanbo Zhang, and Fan Jiang. Droidfuzzer: Fuzzing the android  
 715        apps with intent-filter tag. In René Mayrhofer, Luke Chen, Matthias Steinbauer, Gabriele  
 716        Kotsis, and Ismail Khalil, editors, *The 11th International Conference on Advances in Mobile  
 717        Computing & Multimedia, MoMM '13, Vienna, Austria, December 2-4, 2013*, page 68. ACM,  
 718        2013. [doi:10.1145/2536853.2536881](https://doi.org/10.1145/2536853.2536881).

723 40 Andreas Zeller. Yesterday, my program worked. today, it does not. why? In Oscar Nier-  
724 strasz and Michel Lemoine, editors, *Software Engineering - ESEC/FSE'99, 7th European*  
725 *Software Engineering Conference, Held Jointly with the 7th ACM SIGSOFT Symposium*  
726 *on the Foundations of Software Engineering, Toulouse, France, September 1999, Proceed-*  
727 *ings*, volume 1687 of *Lecture Notes in Computer Science*, pages 253–267. Springer, 1999.  
728  
doi:10.1007/3-540-48166-4\16.

729 **A Overview of bugs found during experimental evaluation**

730 We identify 2 OR-Tools bugs, 4 MiniZinc bugs and 8 CPMpy bugs, and give a short description  
 731 in this section.

732 **Bug 1**

733 Some mutated models are declared unsatisfiable when solving them using Gecode through  
 734 its MiniZinc interface. Solving with another solver confirms that the models are in fact  
 735 satisfiable. This is a critical soundness bug.

736 **Bug 2 & 3**

737 The next 2 bugs are also considered soundness bugs in MiniZinc but are not as severe as  
 738 the first one. There are some models where MiniZinc does not output a value for all the  
 739 variables after solving. This happens for most but not all of the available solvers within  
 740 MiniZinc. The reason we count 2 different bugs is that a third similar bug has already been  
 741 solved after HURRICANE found it earlier on, but this didn't resolve the ones we found here.  
 742 Further distinction lies in the fact that Bug 2 occurs when solving to satisfiability and Bug 3  
 743 happens when solving to optimality.

744 **Bug 4**

745 When using MiniZinc python some models do not respect the given time limit when solving.  
 746 This is due to the compiler optimisation phase getting stuck.

747 **Bug 5**

748 A bug in CPMpy's MiniZinc interface, that causes a crash when a nested sum appears in the  
 749 arguments of the global constraint: ALLDIFFERENTEXCEPT0.

750 **Bug 6**

751 A bug in CPMpy's MiniZinc interface, that causes a crash when the COUNT global constraint  
 752 appears as an argument in a weighted sum.

753 **Bug 7**

754 The helper function canonical\_comparison contained a bug where weighted sums were  
 755 incorrectly transformed. This is a soundness bug.

756 **Bug 8**

757 Inconsistent implementation of the relational semantics for constraint modeling languages  
 758 meant that handling of partial functions such as ELEMENT leads to missing solutions where  
 759 the constraint is undefined, but occurs in a nested context.

760 **Bug 9**

761 CPMpy's helper function is\_bool did not recognise a specific datatype to be Boolean.

**Bug 10**

762 The internal transformation `canonical_comparison` can create weighted sums with zero  
764 arguments, leading to a crash later in the transformation pipeline.

**Bug 11**

766 An assertion error gets triggered in the internal function `canonical_comparison`, when a  
767 CPMpy sum operator is encountered that only contains integers and no variables.

**Bug 12**

769 An equation between an integer and a Boolean expression was treated as reification by the  
770 `flatten` transformation of CPMpy.

**Bug 13**

772 Crash in the OR-Tools solver causing the Python runtime environment to crash.

**Bug 14**

774 A soundness bug in OR-Tools' presolve where the ordering of constraints influences whether  
775 a model was declared to be satisfiable or not.

**A.1 Occurrences of each bug**

776 In Table 5 and Table 6, we show the unaggregated data of how many times each bug was  
778 found by HURRICANE during our experimental evaluation.

■ **Table 5** Bugs found by different verification methods when running with MiniZinc

Verif	#mut	B1	B2	B3	B4	B5	B6	B7	B8	B9	#bugs	#models
All sol	1	-	-	-	-	1	-	-	23	-	24	3492
	2	-	-	-	-	2	-	-	59	-	61	3594
	5	2	-	-	-	5	1	36	86	13	143	3242
	10	-	-	-	1	1	-	38	28	16	84	839
counting	1	-	-	-	-	1	-	-	24	-	25	3633
	2	-	-	-	-	-	-	-	61	-	61	3655
	5	2	-	-	-	3	1	40	94	15	155	3496
	10	1	-	-	-	1	-	39	28	15	84	839
One sol	1	-	-	66	-	15	-	-	133	-	214	65029
	2	-	-	68	-	29	4	1	429	2	533	64725
	5	2	-	91	-	108	8	12	903	32	1156	61554
	10	-	-	6	1	7	-	3	60	3	80	3187
sat	1	-	-	-	-	26	-	-	-	-	26	88981
	2	-	-	-	-	61	6	-	-	1	68	87419
	5	1	-	-	1	6	1	-	-	2	11	6554
	10	1	-	-	-	6	1	-	-	3	11	3165
opt	1	-	-	-	-	-	-	-	-	-	-	57242
	2	-	-	-	-	-	-	-	-	-	-	57134
	5	2	-	-	-	-	-	28	-	-	30	54038
	10	2	7	-	-	-	-	155	-	-	164	49161

779 **B Reformulations as mutations**

780 We summarize the constraint reformulations implemented in CPMpy which are used in our  
 781 mutational testing framework.

782 **Unnesting and normalization of lists**

783 This transformation is the first in the transformation pipeline of any solver implemented in  
 784 CPMpy and all subsequent transformation expect as input a flat list of constraints. This  
 785 Additionally any conjunction at the top-level of the constraint model will be split up into  
 786 separate constraints

787  $M_{unnest}([c_1, [c_2, c_3], [c_4 \wedge c_5]])$

789 with  $c_n, n \in 1..5$  being arbitrary constraints, results in

790  $[c_1, c_2, c_3, c_4, c_5]$

792 **Flattening**

793 Makes sure no nested constraints remain in the expression tree. This reformulation introduces  
 794 a fresh variable to be equated with a (numerical) expression and un-nests each constraint  
 795 accordingly. The output of this reformulation is a set of Boolean expressions within a  
 796 restricted grammar defined by CPMpy's developers. For example, given the expression list

797  $[ALLDIFFERENT(MIN(w, x), y, z)]$  (2)

■ **Table 6** Bugs found by different verification methods when running with OR-Tools

Verif	#mut	B7	B8	B9	B10	B11	B12	B13	B14	#bugs	#models
All sol	1	-	26	-	-	-	-	-	-	26	4102
	2	-	64	-	-	-	-	-	-	64	3786
	5	37	95	13	-	-	-	-	-	145	3332
	10	116	70	38	-	-	-	-	1	225	2221
Counting	1	-	26	-	-	-	-	-	-	26	4152
	2	-	69	-	-	-	-	-	-	69	4128
	5	42	117	16	-	-	-	1	1	177	3718
	10	139	78	43	-	-	-	-	7	267	2553
One sol	1	-	5695	-	-	-	-	-	-	5695	2226130
	2	6	10761	79	-	-	-	-	-	10846	1400180
	5	84	7449	212	1	-	-	-	-	7746	412874
	10	71	1250	80	1	6	-	-	-	1408	56001
Sat	1	-	-	-	-	-	-	-	-	-	1958248
	2	4	-	19	-	-	-	-	-	23	1292747
	5	40	-	55	1	-	-	-	-	96	379361
	10	28	-	29	1	3	-	-	-	61	49044
Opt	1	-	-	-	-	-	-	-	-	-	4973786
	2	-	-	-	-	-	-	-	-	-	3971747
	5	811	-	-	-	-	-	-	-	811	1471156
	10	820	-	-	-	-	2	-	-	822	234891

798 the result of the flattening is

$$799 \quad [\text{ALLDIFFERENT}(e, y, z), e = \text{MIN}(w, x)] \quad (3)$$

800 with  $e$  an auxiliary variable with the right bounds.

### 801 **Decomposing global constraints**

802 This function is one of the elementary operations in constraint modeling languages. While  
 803 many CP-solvers support a variety of global constraints, these advanced relations between  
 804 variables are oftentimes not supported by solvers from other solving paradigms. Hence,  
 805 when a model containing a global constraint has to be solved by for example an SMT-solver,  
 806 it needs to be decomposed into simpler expressions first. This reformulation does exactly  
 807 that. For example, if `ALLDIFFERENT` is not supported by the solver, it is decomposed to a  
 808 conjunction of pairwise disequality constraints.

### 809 **Unnesting of reified constraints**

810 This transformation is applied to ensure no unsupported expressions remain reified. For some  
 811 of the backend solvers in the CPMpy library, reification is only supported on a subset of  
 812 expressions. This reformulation is applied after flattening, and ensures further unnesting such  
 813 that only reifications of supported constraints remain. For example, given the unsupported  
 814 expression  $b \rightarrow \text{MAX}(x, y, z) \leq 10$ , a valid transformation in order to remove the reification  
 815 of the `MAX` is

$$816 \quad (b \rightarrow a \leq 10) \wedge (\text{MAX}(x, y, z) = a) \quad (4)$$

## 23:24 Mutational Fuzz-testing for Constraint Modeling Systems

817 with  $a$  an auxiliary variable with the appropriate bounds. Input constraints must not contain  
818 unsupported global constraints, and must be flattened first.

### 819 Only half-reification

820 It removes all “full reification constraints” from the expression tree and ensures all reifications  
821 end up in the form  $b \rightarrow bexpr$ . This transformation always has to be preceded by the previous  
822 *only boolean variables reify* transformation. For each constraint of the type  $b \leftrightarrow bexpr$ , two  
823 half-reification constraints are introduced:  $b \rightarrow bexpr$  and  $\neg b \rightarrow \neg bexpr$ . This transformation  
824 also simplifies the negated Boolean expression whenever possible. For example, given  $b \leftrightarrow x \wedge y$   
825 as input, the transformation returns  $\{b \rightarrow x \wedge y \text{ and } \neg b \rightarrow (\neg x \vee \neg y)\}$ .

### 826 Normalization of reifications

827 This transformation rewrites any reification such that the Boolean variable occurs on the  
828 left hand side. E.g., constraints of the type  $bexpr \rightarrow b$  are rewritten to  $\neg b \rightarrow \neg bexpr$ ,  
829 full-reification constraints  $bexpr \leftrightarrow b$  are swapped to  $b \leftrightarrow bexpr$ . Similar to the previous  
830 transformation, negated Boolean expressions are simplified when possible. Input constraints  
831 must be flat.

### 832 Linearize

It ensures any flattened constraint is transformed into a canonicalized linear constraint, i.e.,  
a comparison with a weighted sum of integer or Boolean variables on the left-hand side and  
a constant on the right-hand side. The output is thus always of the form

$$\sum w_i x_i \langle cmp \rangle c$$

833 where  $\langle cmp \rangle$  is the one of the comparison operator allowed ( $=, \leq$  or  $\geq$ ), the  $w_i$  are the  
834 integer weights and  $x_i$  the Boolean/integer variables. Before linearizing, unsupported global  
835 constraints must be decomposed, and must contain only boolean implications.

### 836 Normalized numerical expressions

837 This transformation is targeted to be used with solvers that don’t support comparisons ( $<$ ,  
838  $\leq$ ,  $\geq$ ,  $>$ ,  $\neq$ ) between an expression and a constant. An auxiliary variable is thus required to  
839 transform it into a simple comparison. For example, if  $\text{MAX}(x, y, z) \leq 10$  is not supported,  
840 it will be transformed into

$$841 (\text{MAX}(x, y, z) = e) \wedge (e \leq 10) \quad (5)$$

842 by using the auxiliary variable  $e$  (with appropriate bounds). Input constraints must be flat.

### 843 Converting negated Boolean variables

844 After linearization of a set of constraints, it helps make the constraints more compatible  
845 with the API of a typical Mixed Integer Programming solver. Pseudo-Boolean constraints  
846 (weighted-sum over Boolean variables) are converted such that only positive Boolean variables  
847 remain on the left-hand side of the comparisons. For example, the expression  $\neg p + q + r \geq 1$   
848 is re-written as  $-p + q + r \geq 0$  by creating a negative weight and allowing no negation  
849 operator in the formula. Input constraints must be linear.

850 **Conversion of flat expressions to CNF**

851 It is required when using SAT-solvers as backend solvers. This transformation rewrites any  
852 Boolean operator with Boolean variables as arguments to CNF. For example,  $(w \wedge x) \vee (y \wedge z)$   
853 is re-written in

854 
$$(w \vee y) \wedge (w \vee z) \wedge (x \vee y) \wedge (x \vee z) \quad (6)$$

855 Input must ensure only boolean implications

856 **Push negation to leaves**

857 This one simplifies the number of nodes in the expression tree. The reformulation applies  
858 simple equivalence rules such as DeMorgan's laws to make sure the only negation operators  
859 left in the tree are bound to Boolean variables or global constraints. For example, it would  
860 transform the expression  $\neg(a \vee b)$  into  $\neg a \wedge \neg b$ , or the expression  $\neg(a \leq b)$  into  $a > b$ .  
861 The negation of a global constraint such as  $\neg\text{ALLDIFFERENT}(a, b, c)$  can not be simplified  
862 any further, except by decomposing the global constraint first. This will happen in the  
863 "decomposing globals" transformation, depending on solver support.

864 **Simplification of Boolean comparisons**

865 This operation can be done when a Boolean expression is compared to a constant. In that  
866 case, it is trivial to convert the Boolean expression at hand to itself or to its negation. For  
867 example, comparison  $b < 1$ , where  $b$  is a Boolean variable, can be simplified to  $\neg b$ . And  
868  $b \geq \text{True}$  can be converted to just the literal  $b$ .