

Automatic Synthesis of Random Generators for Numerically Constrained Algebraic Recursive Types^{*}

Ghiles Ziat¹, Vincent Botbol², Matthieu Dien³, Arnaud Gotlieb⁴, Martin Pépin¹^[0000-0003-1892-3017], and Catherine Dubois⁵

¹ Université Paris Cité, IRIF, 75205 Paris Cedex 13, France, {ghiles.ziat, martin.pepin}@irif.fr,

² Nomadic Labs, vincent.botbol@nomadic-labs.com

³ Normandie Université, UNICAEN, ENSICAEN, CNRS, GREYC, 14000 Caen, France, matthieu.dien@unicaen.fr

⁴ SIMULA RESEARCH LAB, arnaud@simula.no

⁵ Ecole Nationale Supérieure d'Informatique pour l'Industrie et l'Entreprise, Samovar, catherine.dubois@ensiie.fr

Abstract. In program verification, constraint-based random testing is a powerful technique which aims at generating random test cases that satisfy functional properties of a program. However, on recursive constrained data-structures (e.g., sorted lists, binary search trees, quadrees), and, more generally, when the structures are highly constrained, generating uniformly distributed inputs is difficult. In this paper, we present Testify: a framework in which users can define algebraic data-types decorated with high-level constraints. These constraints are interpreted as membership predicates that restrict the set of inhabitants of the type. From these definitions, Testify automatically synthesises a partial specification of the program so that no function produces a value that violates the constraints (e.g. a binary search tree where nodes are improperly inserted). Our framework augments the original program with tests that check such properties. To achieve that, we automatically produce uniform random samplers that generate values which satisfy the constraints, and verifies the validity of the outputs of the tested functions. By generating the shape of a recursive data-structure using *Boltzmann sampling* and generating evenly distributed finite domain variable values using constraint solving, our framework guarantees size-constrained uniform sampling of test cases. We provide use-cases of our framework on several key data structures that are of practical relevance for developers. Experiments show encouraging results.

^{*} this research was partially supported by the ANR PPS project ANR-19-CE48-0014 and the “DYNNET” project, co-funded by the Normandy County Council and the European Union (ERDF-ESF 2014-2020).

36 1 Introduction

37 Software Testing is one of the most widespread program verification techniques,
38 and is also one of the most practical to implement. One interesting instance of
39 it is Property-based Testing (PBT), where programs are tested by generating
40 random inputs and results of the output are compared against software specifica-
41 tions. This technique has been extensively studied, for testing correctness [20,31],
42 exhaustiveness [34], complexity [11] etc. However, this technique requires the de-
43 veloper to manually write the tests, that is the properties to be checked and the
44 random generators. The latter can be particularly complicated to design, espe-
45 cially in the case of complex and constrained algebraic data structures.

46 In this field, *constraint-based random testing* [23] (commonly used in PBT
47 [26,13]) is a powerful technique which aims at generating random test cases that
48 satisfy functional properties of a program under test. By specifying a property
49 that a program has to satisfy and by using uniformly-distributed inputs gener-
50 erators, it is possible to uncover subtle robustness faults that may be not be
51 discovered otherwise. For instance, [1] explored the usage of PBT for testing
52 a steam boiler, [28] explored its usage for wireless sensor network applications.
53 It is worth noticing that generating inputs according to a uniform probability
54 distribution is crucial to ensure that all the distinct program behaviours have
55 the same chance to be triggered, even those which are the most constrained. The
56 technique has been successfully applied in the field of unit testing for imperative
57 programs [22] as well as various programming languages including Haskell [15],
58 Prolog [3] and proof-assistant methodologies and tools such as Coq [32] or Is-
59 abelle [10]. Sampling constraint systems solutions according to a uniform dis-
60 tribution is a well-known difficult problem. Initially studied in the context of
61 hardware testing [30], the problem has been studied in [21] and more extensively
62 in [2]. Other random generation schemes are either not uniform, or very slow
63 *e.g.* rejection sampling is generally uniform by construction, but fits very poorly
64 with generation under constraints.

65 Recently, in [38] the authors introduced an automated framework capable
66 of providing tests for functions that manipulate constrained values without re-
67 quiring manual input from the programmer. The framework introduces a type
68 language, with algebraic data-types, and constrained types *i.e.* types augmented
69 with a membership predicate that is used to filter invalid representations. To
70 verify that a function does not create invalid representations, the authors opted
71 for a random testing approach. The main interest of the framework is that both
72 the generators and the specifications are automatically extracted from the con-
73 straints specified by the user, which greatly alleviate the user's workload. Gen-
74 erators are uniform random value samplers used to provide input for functions,
75 and specifications that are predicates that verify that a given value satisfies the
76 constraint attached to its type, are used to check whether a function's output
77 violates the constraint or not. Their tool is implemented as a pre-processor for
78 OCaml programs, *i.e.* before compiling, programs are rewritten into augmented
79 programs where a test suite has been added. During the pre-processing step,
80 from each constrained type declaration τ is extracted a CSP p . Then, each p is

81 solved only once, that is to say that a characterisation of the set of solutions of
82 p , called coverage, is calculated. Each coverage is then compiled into code which
83 uniformly generates solutions which are then converted back into values of the
84 type τ . However, to be able to solve a CSP only once per constrained type,
85 the authors limit themselves to types involving a fixed number of numerical
86 atoms (*e.g.* tuples), which automatically excludes recursive types. This makes
87 it impractical as, for instance, in OCaml, real-world programs rely heavily on
88 recursive data-types (lists, trees, sets, etc.).

89 This paper investigates the automatic synthesis of uniform pseudo-random
90 generators, as in [38], but for recursive constrained types.

91 1.1 Contributions

- 92 – A programmable method to restrict the values a recursive type can take.
- 93 – An algorithm that uses Boltzmann generation and constraint solving to au-
94 tomatically derive uniform generators for recursive constrained types.
- 95 – An experimental study of the performances of our technique.

96 1.2 Outline

97 This paper is organised as follows: Sec.2 presents use cases of our methods on
98 some examples of realistic code. Sec.3 defines our solving technique which mixes
99 Boltzmann generation and global constraint solving. Sec.4 recalls some elemen-
100 tary notions about Boltzmann sampling and details some specifics about our
101 use-case. Sec.5 presents our prototype and gives some details about its function-
102 ing, current capabilities and restrictions. We also give some details about our
103 implementation and measure experimentally the performances of the generators
104 we derive for recursive constrained types. Sec.6 describes some related work. Fi-
105 nally, Sec.7 summarises our work and discusses possible future improvements.

106 2 A Declarative Programming Approach

107 We propose a testing framework that allows programmers to specify constraints
108 on recursive data structures. From these constraints, the framework extracts
109 a Constraint Satisfaction Problem (CSP) which is solved in such a way that
110 uniform random instances (*i.e.*, test cases) are generated. These instances are
111 then used for testing functions in order to find defects.

112 2.1 Preliminaries

113 A pseudo-random generator g for an algebraic data-type τ is a function g of
114 type $\mathcal{S} \rightarrow \tau$. Here, \mathcal{S} is the random state used by the pseudo random number
115 generator. A constrained type is a pair $\langle \tau, p \rangle$, with τ an algebraic data-type and
116 $p : \tau \rightarrow \text{bool}$ a predicate over values of type τ . The set of its inhabitants is

117 defined as $\{t \in \tau \mid p(t) = true\}$. A pseudo-random generator g for a constrained
118 type $\langle \tau, p \rangle$ is a function $g : \mathcal{S} \rightarrow \tau$ s.t $\forall s \in \mathcal{S}, p(g(s)) = true$.

119 Here, we face two main challenges for automating random testing of recursive
120 data-types. First, we have to equip the developer with convenient means for
121 specifying constraints attached to a given data-type. For example, we want to
122 express that a list of integers is sorted or that a tree is a binary search tree
123 (i.e., the left child node value is always smaller than the right one). Second,
124 building an uniform random value generator for constrained recursive data-types
125 is highly challenging. Recursive types can dynamically grow to an arbitrarily
126 large size and, deriving generators for such types requires the resolution of a
127 complex constraint system. In particular, we have to manage CSPs with an a-
128 priori unknown number of variables and constraints. The grammar of Ocaml
129 types and constraints annotations are given in Figure 7. In the following, we
130 give two illustrative examples.

131 2.2 Example 1: Inserting an element into a set of integers

132 Let start with `list`, a recursive data-type associated to lists of integers, for
133 which a possible type declaration is given in Fig.1. Using `list` to specify a

```
1 type list = Empty | Cons of int * list
```

Fig. 1: OCaml type declaration of lists of integers

133 Set data structure can easily be done using Testify, by using the annotation
134 `[@satisfying _]` and the `[alldiff]` constraint as illustrated by Fig. 2.

```
1 type uniquelist =  
2   | Empty  
3   | Cons of int * uniquelist [@satisfying alldiff]
```

Fig. 2: OCaml type declaration of sets of integers using lists

135 We can automatically test the functions that manipulate instances of the
136 `uniquelist` type by checking if they break the properties attached to it. For
137 that, we have to define a generator and a specification for the corresponding
138 type. To randomly generate instances, we first draw at random an instance of
139 size n using Boltzmann generation (see Sec.4), then we build a CSP $(\mathcal{X}, \mathcal{D}, \mathcal{C})$
140 containing n finite domain variables and solve it using Path-oriented Random
141 Testing (PRT) (see Sec.3) and eventually we build a random generator g able to
142 produce uniformly distributed sets of size n .

144 A key aspect of Testify is based on the usage of *global constraints*, which
145 are arithmetic-logic constraints holding over a non-fixed number of variables.
146 In the example of Fig.2, we translate the declaration `[alldiff]` into a `all_-`
147 `different` global constraint implementation and used it to generate uniformly
148 distributed solutions that can be used to populate test cases. For other recursive
149 data-types, we use combination of multiple global constraints and arithmetic

150 constraints. Possible recursive data-types that can be implemented and tested
 151 in our framework include functions that generate and manipulate (un-)ordered
 152 lists and sets, trees, binary search trees, quadtrees, etc.

153 Fig.3 shows an example of a function which implements the insertion of an
 154 element within a set of integers and the code that is automatically generated for
 155 the testing of this function⁶.

```

1 let rec add (x:int) (l:uniquelist) : uniquelist =
2   match l with
3   | Empty -> Cons(x,Empty)
4   | Cons(h,tl) -> if x <> h then Cons(h,(add x tl) else l)
5
6 (* generated code*)
7 let add_test () =
8   let size = Random.int () in let rand_x = Random.int () in
9   let rand_l = unique_list size in
10  assert (alldiff_checker (add rand_x rand_l))

```

Fig. 3: Insertion of an element into a set, and the generated corresponding test

156 Here, testing the function means verifying that every output pro-
 157 duced is indeed sorted (`assert (alldiff_checker (add rand_x rand_l))`). Note
 158 that we have used the return type annotation to automatically derive a (partial)
 159 specification for the function, but the generator we automatically synthesise can
 160 also be used to test any hand-written specification.

161 2.3 Example 2: Binary Search trees

162 Binary Search Trees (BST) are binary trees that additionally satisfy the following
 163 constraint: the key in each node is greater than or equal to any key stored in
 164 the left sub-tree, and less than or equal to any key stored in the right sub-tree.
 165 Stated differently, the keys in the tree must be in increasing order in a depth-
 166 first search traversal, in infix order. From this observation, we propose, using
 167 our framework, a possible OCaml declaration for BSTs illustrated in Fig.4.

```

1 type bst =
2   | Node of bst * (int[@collect]) * bst
3   | Leaf [@@satisfying fun x -> increasing x]

```

Fig. 4: Testify type annotation for binary search trees

168 Rather than defining global constraint for all user-declared data-types, we
 169 break the problem in two parts. On the one hand, we define or reuse known
 170 global constraints for lists, and on the other hand we define a way to browse

⁶ The predicate `alldiff_checker` checks that the result list does not contain dupli-
 cates. It should not be mixed with the version of `alldiff` used in the type declaration
 which is used to generate randomly distributed solutions of that constraint

171 data structures, in a certain order, by collecting the components that are subject
 172 to a global constraint. This is done using the `(int[@collect])` annotation.

173 Also, the order in which the structure is explored is crucial as it determines
 174 the order in which the variables will be passed to the global constraint. By
 175 default, a depth first order is assumed. For constructors with several arguments
 176 (e.g. `Node`), and for tuples, the order in which the traversal is made is mapped
 177 on the declaration order of the tuple component, that is in traversal order. Fig.5
 178 shows the code generated that traverses the tree.

```

1 let rec collect = function
2 | Node (a, b, c) ->
3   List.flatten [collect a; Collect.int b; collect c]
4 | Leaf -> []

```

Fig. 5: Generated collector for binary trees

179 Here, the primitive `Collect.int` is a primitive of our framework that takes
 180 an integer and builds the singleton list with this element. This way, we first
 181 visit the left sub-tree, the root and the right sub-tree. Using pre-order or post-
 182 order would give different results. This means that the constructor `Node` must
 183 be declared in the above order and, for example, the following would be in-
 184 valid: `Node of (int[@collect]) * binary_tree * binary_tree`. However, this re-
 185 striction can easily be lifted by providing an annotation which would allow the
 186 programmer to explicitly specify the traversal order. Similarly, a global anno-
 187 tation `[@bfs]` (resp. `[@dfs]`) could be used to specify that the structure must
 188 be traversed using a breadth first search (resp. depth first search). This will be
 189 studied in future work.

190 3 Constrained Type Solving

191 A *Constraint Satisfaction Problem* (CSP) is a triple $(\mathcal{X}, \mathcal{D}, \mathcal{C})$ where \mathcal{X} is a set
 192 of variables, \mathcal{D} is a function associating a finite domain (considered here as a
 193 subset of \mathcal{Z} without any loss of generality) to every variable and \mathcal{C} is a set of
 194 constraints, each of them being $\langle var(c), rel(c) \rangle$, where $var(c)$ is a tuple of
 195 variables $(X_{i_1}, \dots, X_{i_r})$ called the scope of c , and $rel(c)$ is a relation between
 196 these variables, i.e., $rel(c) \subseteq \prod_{k=1}^r D(X_{i_k})$. For each constraint c , the tuples of
 197 $rel(c)$ indicate the allowed combinations of value assignments for the variables
 198 in $var(c)$. Given a CSP $(\mathcal{X}, \mathcal{D}, \mathcal{C})$, an *assignment* is a vector (d_1, \dots, d_n) , which
 199 associates to each variable $X_i \in \mathcal{X}$ a corresponding domain value $d_i \in D(X_i)$.
 200 An assignment satisfies a constraint c if the projection of \mathcal{X} onto $var(c)$ is a
 201 member of $rel(c)$. The set of all satisfying assignments is called the solution set,
 202 noted $sol(\mathcal{C})$. A constraint c is said to be *satisfiable* if it contains at least one
 203 satisfying assignment, it is inconsistent otherwise. A CSP $(\mathcal{X}, \mathcal{D}, \mathcal{C})$ is satisfiable
 204 if it contains at least one assignment which satisfies all its constraints (i.e.,
 205 $sol(\mathcal{C}) \neq \emptyset$). A global constraint is an extension of CSPs with relations concerning
 206 a non-fixed number of variables. For instance, the *sort*(Xs, Ys) global constraint

207 [29] takes as inputs two lists of n finite domain variables Xs, Ys (where n is
 208 unknown) and states that for each satisfying assignment $(d_1, \dots, d_n, d'_1, \dots, d'_n)$ of
 209 the constraint, we have $\forall j, \exists i$ s.t. $d'_j = \sigma(d_i)$ and $d'_1 \leq \dots \leq d'_n$, where σ is a
 210 permutation of $[1..n]$. Filtering a global constraint $c(X_1, \dots, X_n)$ with the *bound-*
 211 *consistency* local filtering property means to find D' such that for all i , the
 212 extrema values of $D'(X_i)$ are parts of satisfying assignments of c .

213 3.1 Path-Oriented Random Testing

214 *Path-oriented Random Testing* (PRT) is basically a divide-and-conquer algo-
 215 rithm, introduced in [22], which aims to generate a uniformly distributed subset
 216 of solutions of a CSP. Starting from an initial filtering step result, the general
 217 idea is to fairly divide the resulting search space into boxes of equal volumes and,
 218 after having discarded inconsistent boxes using constraint refutation, to draw at
 219 random satisfying assignments.

220 More precisely, applying constraint filtering results in domains that can be
 221 over-approximated by a larger box (i.e., an hyper-cuboid) that contains all the
 222 filtered domains. Based on an external division parameter k , PRT then fairly
 223 divides the box into k subdomains of *equal* volume. When a subdomain cannot
 224 be divided according to the division parameter k , then it is simply extended until
 225 its area can be divided. The iteration of the process leads to a fair partition of
 226 the search space into k^n subdomains where n is the number of variables of
 227 the CSP. Then constraint refutation can be used to discard (some) subdomains
 228 which are inconsistent with the rest of the CSP. As all subdomains have the
 229 same volume, it becomes possible to sample first the remaining subdomains
 230 and then, second, to randomly draw values from these subdomains. Note that,
 231 when all the subdomains are shown to be inconsistent, then the CSP is shown
 232 to be inconsistent. This contrasts with reject-based methods which will trigger
 233 assignment candidates and will reject them afterwards, without terminating in
 234 a reasonable amount of time.

```

Input: CSP:  $(\mathcal{X}, \mathcal{D}, \mathcal{C})$ ,  $k, N$ : #Sol. - Output:  $t_1, \dots, t_N$  or  $\emptyset$  (Inconsistent)
 $\mathcal{D}' := \text{boxfilter}_{bc}(\mathcal{X}, \mathcal{D}, \mathcal{C})$ ;  $(H_1, \dots, H_p) := \text{Fairly\_Divide}(\mathcal{D}', k)$ ;  $T := \emptyset$ ;
while  $N > 0$  and  $p \neq 0$  do
  Pick up uniformly  $H$  at random from  $H_1, \dots, H_p$ ;
  if  $H$  is inconsistent w.r.t.  $\mathcal{C}$  then
    | remove  $H$  from  $H_1, \dots, H_p$ 
  else
    | Pick up uniformly  $t$  at random from  $H$  and remove it;
    | if  $\mathcal{C}$  is satisfied by  $t$  then
      | add  $t$  to  $T$ ;  $N := N - 1$ ;
    | end
  end
end
return  $T$ ;

```

Algorithm 1: Path-Oriented Random Testing adapted from [22] to the uniform random generation of N solutions of a CSP

235 The PRT algorithm, adapted from [22] to the case of CSP solution sam-
 236 pling, is given in Figure 1. It takes as inputs a CSP, a division parameter k ,
 237 and N a non-negative integer. Here, we make the hypothesis that, if the CSP
 238 is consistent, it contains more than N solutions. The algorithm outputs a se-
 239 quence of N uniformly distributed random assignments which satisfy the CSP.
 240 If the CSP is unsatisfiable, then PRT returns \emptyset . After an initial filtering step
 241 using bound-consistency, the algorithm partitions the resulting surrounding box
 242 in subdomains of equal volume (`Fairly_Divide` function). Then, for each lo-
 243 cally consistent subdomain H in the partition, value assignments are randomly
 244 selected and checked against the constraints of the CSP. Those which do not
 245 satisfy the constraints are simply rejected. As shown in [22], this process ensures
 246 the uniform generation of tuples in the solution space.

247 3.2 Extension with Global Constraints

248 Handling global constraints is a natural extension of PRT as it allows us to
 249 handle recursive constrained data-types. As the shape of the data structure
 250 is unknown at constraint generation time, the number of variables to be han-
 251 dled is also unknown in the general case. Thus, using global constraints in this
 252 context is particularly useful as it allows us to avoid the decomposition of a
 253 global constraint into the conjunction of several simpler constraints. This re-
 254 sults in both a stronger and faster pruning. In order to handle recursive con-
 255 strained data-types, we had to provide a dedicated interface for accessing the
 256 deductions from global constraint solving. To facilitate the access to global
 257 constraints, we created an API which provides results of PRT over different
 258 global constraint combinations. The API provides access to predicates such as
 259 `increasing_list(+int LEN,+int GRAIN,-var L)` in which L is instantiated to
 260 a list of LEN uniformly distributed random integers ranked in increasing order,
 261 and the random generator is initialised with `GRAIN`. Optionally, the predicate
 262 can be called with domain constraints in order to constrain the returned list
 263 of values in specific subdomains. Other similar predicates are provided as part
 264 of the API, namely `increasing_strict_list/3 (+int LEN,+int GRAIN,-var`
 265 `L)` which returns a list of strictly increasing integers; `decreasing_list/3 (resp.`
 266 `decreasing_strict_list/3)` which provides a list of integers in (resp. strict)
 267 decreasing order or else `alldiff_list/3` which returns a list of uniformly dis-
 268 tributed random distinct integers. PRT can also be used in combination with
 269 any available global constraint and arithmetico-logic constraint. The following
 270 example, given in Fig.6 illustrates how PRT is used in this respect.

271 In this example, PRT is used with one global constraint, namely `sort(Xs, Ys)`,
 272 and some domain and arithmetic constraints to populate a constrained binary
 273 search tree (BST) of size 6. In this example, the shape of the BST is unknown
 274 and some constraints hold over the keys: the domain of the key-variables is
 275 constrained (from an externally specified source), *e.g.*, key $X_1 \in -2..8$, key
 276 $X_2 \in -3..5$, etc. and any key of the BST corresponds to the sum of its chil-
 277 dren (if any), *e.g.*, $Y_{father} = Y_{child_1} + Y_{child_r}$. Note that the keys have to be set
 278 in increasing order to correspond to a valid BST. Note also that we ignore in

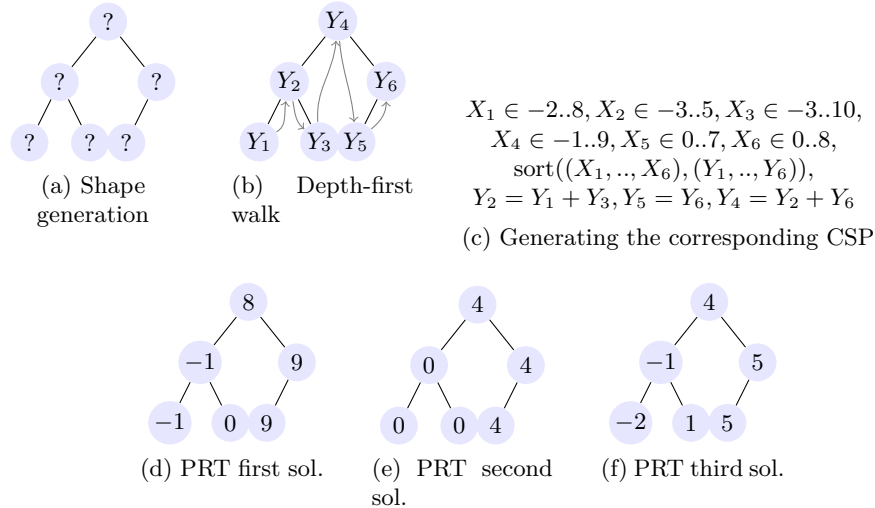


Fig. 6: Generation of a constrained BST of size 6. Division parameter=2, length of seq.=3, 60 subdomains over 64 have been discarded after the first filtering.

279 which order will the keys be positioned in the tree. The first step of our method
 280 corresponds to the generation of a uniformly distributed random shape of the
 281 BST (Fig.6(a)) using the Boltzmann method, described in Sec.4. Then, a depth-
 282 first walk along the tree assigns variable identifiers to the nodes and collects the
 283 constraints that must hold over the constrained data structure (Fig.6(b)). The
 284 generated CSP (Fig.6(c)) can then be solved by using PRT, which generates, in
 285 this example, three uniformly distributed random solutions (Fig.6(d)(e)(f)). It
 286 is worth noticing that other uniform random solutions sampling methods such
 287 as [21,36] could have been used in this context. PRT was chosen because of its
 288 availability and simplicity. However, non-uniform random sampling such as a
 289 simple heuristic selecting at random variable and values to be enumerated first
 290 would not have been appropriate in this context as the goal was to test the
 291 robustness of user-defined functions in functional programming.

292 4 Boltzmann Sampling

293 The Boltzmann method was introduced in [17] as an algorithmic method to
 294 derive efficient sampler from *combinatorial classes*. Combinatorial classes are
 295 just sets of discrete structures with a size (a non-negative integer) and such that
 296 the number of structures having the same size is finite. For example, the binary
 297 trees whose size is the number of leafs is a combinatorial class, but binary trees
 298 whose size is the length of leftmost branch is not because the number of binary
 299 trees with a leftmost branch of fixed length k is infinite. We briefly present the
 300 method here and refers the reader to [17] for more details.

301 In the context of that paper (similarly to [12]), that method directly trans-
 302 lates into an automatic way to derive a uniform random generator of terms for

$\langle decl \rangle ::= \text{'type' } \langle type\ identifier \rangle \text{'=' } \langle type \rangle$	type declaration
$\quad \{ \text{'[@satisfying' } \langle constraints \rangle \text{']' } \}$	Testify's annotation
$\langle type \rangle ::= \langle coretype \rangle$	
$\quad \langle sumtype \rangle$	
$\langle coretype \rangle ::= \text{'int' } \text{'float' } \text{'char' } \dots$	basic types
$\quad \langle coretype \rangle \{ \text{'*' } \langle coretype \rangle \}$	product
$\quad \langle type\ identifier \rangle$	
$\langle sumtype \rangle ::= \langle constructor\ identifier \rangle \{ \text{'[@collect]' } \} \text{'of' } \langle coretype \rangle$	
$\quad \langle sumtype \rangle \{ \text{' ' } \langle sumtype \rangle \}$	
$\langle constraints \rangle ::= \text{'alldiff' } \text{'increasing' } \text{'decreasing'}$	SICStus global constraints
$\quad \langle arith \rangle$	arithmetic constraints like in [38, Fig. 1]
$\quad \langle constraints \rangle \{ \text{'\&\&' } \langle constraints \rangle \}$	

Fig. 7: Syntax of OCaml algebraic data-types (ADT) with Testify's annotation

303 the type language whose syntax is given in Fig.7. In our case, the produced
304 generators only generate a *shape* of tree structure in a first step and the content
305 of this shape is provided in a second step by a constraint solver which makes
306 sure to fill the shape with values that satisfy the specified constraints. For each
307 constrained recursive type declaration, we must therefore generate a glue func-
308 tion between the shapes generated by the Boltzmann sampling method and the
309 solutions returned by the solver used. This function is illustrated in the case of
310 binary search trees in Fig.8

```

1  let rec fill_binary_tree shape solutions =
2    match shape with
3    | Label ("Node", [x1; x2; x3]) ->
4      let x1 = fill_binary_tree x1 solutions in
5      let x2 = Testify_runtime.to_int x2 solutions in
6      let x3 = fill_binary_tree x3 solutions in
7      Node (x1, x2, x3)
8    | Label ("Leaf", []) -> Leaf

```

Fig. 8: Generated function for filling the shapes for binary search trees.

311 Here, we consider types as sets of terms (the inhabitants of the type) whose
312 size is the number of `[@collect]` values they contain. For example, using `type binary_tree`
313 of Sec.2, the term `Node(Node(Leaf, 3, Leaf), 25, Leaf)` has size 2.

314 In the following, we denote $\Gamma \mathcal{A}_x$ a Boltzmann sampler of parameter x for the
315 set \mathcal{A} . Such sampler produces an object $\gamma \in \mathcal{A}$ with a probability $\frac{x^{|\gamma|}}{A(x)}$ where $|\gamma|$
316 is the size of γ and $A(x)$ is a normalizing factor called *generating series*⁷. Note
317 that objects of the same size have the same probability to be drawn.

⁷ The generating series $A(z)$ of a combinatorial class \mathcal{A} is defined by $A(z) = \sum_{\gamma \in \mathcal{A}} z^{|\gamma|}$

318 The second interest of Boltzmann samplers is that they compose well with
 319 sum, product and substitution *i.e.* the constructors of ADTs. Fig.9 shows the
 derivation of such samplers. At the end of the generating process, the object

```

1 type t = a * b (* a and b are 2 types previously defined *)
2 let gen_t x = gen_a x, gen_b x
3
4 type u = A of a | B of b
5 let gen_u x =
6   if random() < A(x)/(A(x) + B(x))
7   then gen_a x else gen_b x
8
9 type alst = Nil | Cons of a * alst (* alst(z) = 1 + z · alst(z) *)
10
11 let rec gen_alst x : alst =
12   if random() < 1/(1 - A(x))
13   then Nil else Cons(gen_a x, gen_aList A(x))

```

Fig. 9: Sampler derivation using Boltzmann

320
 321 drawn has a random size, but we see in the previous code that the choice of
 322 the parameter x influences the size. Note that we can precisely and efficiently
 323 compute x to target a size (see [6] or [33] for the details).

324 Still, the size is random. The last ingredient is to choose a parameter ϵ (which
 325 does not depend of the targeted size n) and keep only objects of size between
 326 $n - \epsilon$ and $n + \epsilon$. Thus, the size of the object is kept up to date during the
 327 generation and the generation is stopped if that size exceeds the upper bound
 328 $n + \epsilon$. At the end, the object may be smaller than $n - \epsilon$ in which case it is rejected
 329 too. However, the theory (see [17]) guarantees that the rejections cost remain
 330 relatively low, *i.e.* the cumulated size of objects sampled to obtain an object of
 331 size in the interval $[n - \epsilon, n + \epsilon]$ is in $\mathcal{O}(n)$. So the complexity of the overall
 332 process is linear in the size of the generated object.

333 An important point to mention is the case of polymorphic types. From a
 334 theoretical point of view they fit in the framework. But from a practical point of
 335 view it is hard to sample a “polymorphic value”. To deal with that limitation, the
 336 Boltzmann samplers are instantiated only for concrete types *e.g.* not for `'a list`
 337 but for `int list`.

338 5 Implementation and Experiments

339 We have implemented the work presented in the previous sections in a tool
 340 available at the url <https://github.com/ghilesZ/Testify>. Our implementa-
 341 tion relies on several state-of-the-art tools. The derivation of OCaml code from
 342 annotated OCaml source files is done using the *ppx* framework, as in [37,5],
 343 which is a form of generic programming [24]. Pre-processors using *ppx* are ap-
 344 plied to source files before passing them on to the compiler. They can be seen as
 345 self-maps over abstract syntax trees. In our case, the source files are traversed to

346 find OCaml type declarations and derive their associated generators. These gener-
347 erators are then used to provide inputs for the functions that must be tested. We
348 have implemented the techniques presented here for the global constraints that
349 we have been able to identify in real data structures (BSTs, Sets, etc) namely
350 *alldiff*, *increasing* and *decreasing* (both strict and large versions). Note that to
351 extend our implementation, *i.e.* add a global constraint,, it is sufficient to add
352 to the constraint solver a propagator for the said global constraint, as both the
353 step of traversing the structure and the random generation procedure presented
354 in section 3 being common to all types.

355 The work done by Testify is divided into two phases: the first is the pre-
356 processing phase during which our tool collects some information on the types
357 needed to build the generators. The second is the testing phase, where the gener-
358 ated code is executed to produce inputs for the functions under test. Note that
359 the pre-processing phase is performed only once while the testing phase can be
360 triggered multiple times, each time one needs to run the tests.

361 We distinguish four kinds of types, for which we provide four different syn-
362 thesis techniques:

- 363 – For non recursive unconstrained types (*e.g.* `int, float * (int * int) ...`)
364 we determine at pre-processing time the function to be used as a generator.
365 For that, we rely on the `qcheck` [15] library, which provides the primitives
366 for building and composing generators.
- 367 – For non recursive constrained types (*e.g.* `int[@satisfying fun x -> x >=0]`),
368 we extract a single CSP which is solved once, still at pre-processing time.
369 From this resolution is extracted a code that draws uniformly solutions of
370 this CSP and rebuild from them a value of the corresponding type. This is
371 the method described in [38].
- 372 – For recursive unconstrained types (*e.g.* lists, binary trees), we build samplers
373 by using the Arbogen [19] tool. This tool implements the Boltzmann method
374 presented in Sec.4. The tuning of the Boltzmann parameter is done at pre-
375 processing time while the shape generation, and the conversion of this shape
376 to a value of the targeted type is done at testing time.
- 377 – Finally, for recursive constrained types (*e.g.* sorted lists, binary search trees),
378 the previous techniques are mixed together to produce efficient generators:
379 first, a targeted size n is drawn, then, a shape of size n is sampled. We
380 then browse the generated shape, collecting constrained values to build a
381 CSP as explained in Sec.3. This CSP is then fed to the SICStus Prolog [4]
382 solver, which builds from it a generator using the *PRT* library [22]. Finally
383 we put together shapes and constrained values. All of these steps are made
384 at testing time, that is every time we have to generate a value we must solve
385 a CSP. This is arguably the bottleneck of our architecture, but experiments
386 still demonstrate the usability of our method.

387 5.1 Experiments

388 In this section, we focus on the performance of our automatically derived gener-
389 ators. We measure the generation times (in seconds) obtained with our method

390 for different constrained recursive types and by varying the size of the gen-
 391 erated structure. The types we are interested in are: lists sorted in ascend-
 392 ing order, association lists with unique keys, lists of pairs in ascending order
 393 $((x, y) \leq (x', y) \Leftrightarrow x \leq x' \wedge y \leq y')$, binary search trees (unbalanced), functional
 394 maps (key-value stores as binary search trees) and quadtrees. These types are
 395 among the most frequent in the literature, and they only involve numerical
 396 constraints, which Testify is able to manage.

Types	Tar- geted	Average	#Objects	time
increasing_list	10	8.50	2889	0.020
	100	93.95	13691	0.004
	1000	948.57	17763	0.003
	10000	9392.45	79	0.757
assoc_list	10	8.49	2534	0.023
	100	93.96	11949	0.005
	1000	947.87	13660	0.004
	10000	9406.92	76	0.786
bicollect	10	6.99	2492	0.024
	100	93.04	6418	0.009
	1000	947.73	16048	0.003
	10000	9456.85	1596	0.037
binary_tree	10	9.00	238690	0.001
	100	94.35	21214	0.001
	1000	948.00	3416	0.006
	10000	9740.00	1500	0.040
map	10	9.00	238690	0.001
	100	94.37	21208	0.001
	1000	947.08	3423	0.006
	10000	9047.00	1276	0.047
quad_tree	10	8.00	3590507	0.001
	100	93.88	228357	0.001
	1000	947.79	23548	0.002
	10000	9489.19	2191	0.027

Fig. 10: Generation time per object according to the size of the structure

397 The experience was to sample as much as possible constrained structures
 398 during one minute. The results are shown in Fig.10. For each type we report the
 399 size of the terms (number of `@collect` values) targeted, the average size of the
 400 generated terms, the number of terms sampled and the average time to sample
 401 one term. The computer running the experiments has an Intel Core i7-6700 CPU
 402 cadenced at 3.40GHz with 8 GB of RAM.

403 As expected, at least for the tree-like types, we observe that the complexity
 404 is quite linear in the size of the sampled terms: the Boltzmann method keeps
 405 its promises and the use of a constraint solver proves to be fast enough to be
 406 used in our context. For most of these structures we manage to generate several
 407 hundred values per second, up to a certain structure size. These results prove the
 408 relevance of our method in the context of testing, as it can allow the user to fine-

409 tune the generators to decide whether he wants to test his functions on several
410 small structures and/or a few large ones. However, we may note that sampling
411 of lists is much slower than sampling of trees. This is due to the fact that the
412 Boltzmann method is not tailored for regular languages (such as list). It would
413 probably be more efficient to use specialised algorithms for regular languages
414 such as the one of [8].

415 6 Related Work

416 In this section we focus on related work dealing with constraint-based generation
417 techniques. Constraint-based generation of test data has been exploited in white-
418 box testing to produce inputs that will follow some execution paths, as well as in
419 functional testing to generate constrained inputs. In [35], Senni applies constraint
420 logic programming to systematically develop generators of structurally complex
421 test data, e.g. red-black trees, in the context of Bounded-Exhaustive Testing.

422 PBT, as exemplified by Quickcheck for Haskell, has been adapted to many
423 programming languages but also to proof assistants to test conjectures before
424 proving them, e.g. [18,10,32,13]. In [18] restricted classes of indexed families
425 of types are provided with surjective generators. In [13], the authors propose
426 the FocalTest framework for testing - conditional - conjectures about functional
427 programs and for automatically generating constrained values. In this work, CP
428 global constraints are not used and thus FocalTest does not take benefit from
429 the corresponding efficient filtering ad hoc procedures.

430 In the context of PBT of Erlang programs, De Angelis *et al* propose in [16]
431 an approach to automatically derive generators of values that satisfy a given
432 specification. Generation is performed via symbolic execution of the specifica-
433 tion using constraint logic programming. A difference between their approach
434 and ours is that we craft a suitable representation of a given type at static time,
435 which is then compiled into an efficient generator. In [16], generators are built
436 at execution time, while testing, which ultimately leads to a slower generation.
437 The Coq plugin QuickChick helps to test Coq conjectures as soon as involved
438 properties are executable. It allows the automatic synthesis of random genera-
439 tors for algebraic data-types, recursive or not, and also the definition of simple
440 inductive properties, e.g. a property specifying binary search trees whose ele-
441 ments are between two bounds, to be turned into random generators of con-
442 strained values [25]. The approach is narrowing-based, like in [14]. Such a binary
443 tree is built lazily while solving the constraints found in the inductive property
444 while in Testify, the shape of the data structure is randomly chosen and then
445 its elements are obtained by solving constraints. This tool comes with differ-
446 ent primitives or mechanisms allowing for some flexibility in the distribution
447 of the sampled values. For example the user can annotate the constructors of
448 an inductive data-type with weights that are used when automatically deriving
449 generators. Furthermore, it also produces proofs of the generators correctness.
450 In [12], the authors adapt a Boltzmann model for random generation of OCaml
451 algebraic data-types, possibly recursive, but not constrained. Generators are au-

452 tomatically derived from type declarations. In [14], Claessen et al. propose an
453 algorithm that, from a data-type definition, a constraint defined as a Boolean
454 function and a test data size, produces random constrained values with a uni-
455 form distribution. However the authors show that this uniformity has a high
456 cost. They combine this perfect generator with a more efficient one based on
457 backtracking. Limiting the class of constraints and combining it with an efficient
458 solving process, Testify can generate constrained values with a uniform distri-
459 bution in a reasonable time. Some work focus on the enumeration or sampling
460 of combinatorial structures, like lambda-terms, using Boltzmann samplers [27],
461 Prolog mechanisms [9] or both [7]. These approaches are dedicated to objects
462 of recursive algebraic data-types with complex constraints, like typed lambda-
463 terms, closed lambda-terms, linear lambda-terms, etc. This kind of constraints
464 is out of reach of our tool whose objective is not only to generate constrained
465 values but also to provide the programmer with syntactic facilities to specify
466 them.

467 7 Conclusion

468 We have proposed in this paper a technique based on declarative programming,
469 to derive generators of random and uniform values for constrained recursive
470 types. We have proposed a small description language for recursive structure
471 traversal which allows us to build a custom CSP for each term to be generated.
472 The code we generate is efficient, and outperforms a naive generation technique
473 based on rejection, and allows us to generate large recursive structures quickly.
474 Starting from the constraints attached to a type, we first sample the shape of the
475 value to generate and then build a CSP that encodes the valid representations of
476 the terms that have this shape. Then, our tool uses the SICStus Prolog constraint
477 solver to filter invalid representations and produce a uniform solution sampler.
478 Our technique is integrated into the Testify framework, which embeds these
479 generators within a fully automatic test system. The generators derived by our
480 framework are fast enough to allow the user to run tests each time he compiles
481 his code. This would allow him to be able to detect bugs very quickly and
482 fix them before they become potentially harmful. However, we still have a lot
483 of work to do to improve Testify. For example, we can extend the constraint
484 language to be able to handle types with shape constraints (*e.g.* balanced trees).
485 This would require adapting the Boltzmann technique to random sampling of
486 tree structures under constraints. Also, when dealing with a functional language,
487 functions as values cannot be avoided: it will be necessary to have techniques for
488 the derivation of generators for functions, and explore what kind of constrained
489 functions (monotonic, bijective functions, etc.) appear in practice in programs.
490 Moreover, in this paper we have only studied *tree-like* recursive data-structures.
491 Some structures do not fit into this framework (*e.g.* graphs, doubly linked lists)
492 and it would be interesting to see to what extent our methods adapt to these
493 structures. Also, our current implementation tests functions by generating any
494 random input, disregarding their body. This is naturally an important point of

495 improvement. For example, one could imagine a static analysis of the body of the
496 function, to conduct the input generation more precisely, and find bugs faster.
497 Finally, our framework targets OCaml but the methods developed in this paper
498 can be adapted to most programming languages and proof assistants.

499 References

- 500 1. Olfa Abdellatif-Kaddour, Pascale Thévenod-Fosse, and H el ene Waeselynck.
501 Property-oriented testing: A strategy for exploring dangerous scenarios. In Gary B.
502 Lamont, Hisham Haddad, George A. Papadopoulos, and Brajendra Panda, editors,
503 *Proceedings of the 2003 ACM Symposium on Applied Computing (SAC), March 9-*
504 *12, 2003, Melbourne, FL, USA*, pages 1128–1134. ACM, 2003.
- 505 2. Thomas E. Allen, Judy Goldsmith, Hayden Elizabeth Justice, Nicholas Mattei,
506 and Kayla Raines. Uniform random generation and dominance testing for cp-nets.
507 *J. Artif. Intell. Res.*, 59:771–813, 2017.
- 508 3. Cl audio Amaral, M ario Florido, and V itor Santos Costa. PrologCheck - property-
509 based testing in Prolog. In Michael Codish and Eijiro Sumii, editors, *Functional*
510 *and Logic Programming - 12th International Symposium, FLOPS 2014, Kanazawa,*
511 *Japan, June 4-6, 2014. Proceedings*, volume 8475 of *Lecture Notes in Computer*
512 *Science*, pages 1–17. Springer, 2014.
- 513 4. Johan Andersson, Stefan Andersson, Kent Boortz, Mats Carlsson, Hans Nilsson,
514 Thomas Sj oland, and Johan Wid en. SICStus Prolog user’s manual. 1993.
- 515 5. Florent Balestrieri and Michel Mauny. Generic programming in OCaml. *Electronic*
516 *Proceedings in Theoretical Computer Science*, 285:59–100, 12 2018.
- 517 6. Maciej Bendkowski, Olivier Bodini, and Sergey Dovgal. Polynomial tuning of
518 multiparametric combinatorial samplers. In Markus E. Nebel and Stephan G.
519 Wagner, editors, *Proceedings of the Fifteenth Workshop on Analytic Algorithmics*
520 *and Combinatorics, ANALCO 2018, New Orleans, LA, USA, January 8-9, 2018*,
521 pages 92–106. SIAM, 2018.
- 522 7. Maciej Bendkowski, Katarzyna Grygiel, and Paul Tarau. Random generation of
523 closed simply typed λ -terms: A synergy between logic programming and Boltzmann
524 samplers. *Theory Pract. Log. Program.*, 18(1):97–119, 2018.
- 525 8. Olivier Bernardi and Omer Gim enez. A linear algorithm for the random sampling
526 from regular languages. *Algorithmica*, 62(1–2):130–145, feb 2012.
- 527 9. Olivier Bodini and Paul Tarau. On uniquely closable and uniquely typable skele-
528 tons of lambda terms. In Fabio Fioravanti and John P. Gallagher, editors, *Logic-*
529 *Based Program Synthesis and Transformation - 27th International Symposium,*
530 *LOPSTR 2017, Namur, Belgium, October 10-12, 2017, Revised Selected Papers*,
531 volume 10855 of *Lecture Notes in Computer Science*, pages 252–268. Springer,
532 2017.
- 533 10. Lukas Bulwahn. The new quickcheck for isabelle - random, exhaustive and symbolic
534 testing under one roof. In Chris Hawblitzel and Dale Miller, editors, *Certified*
535 *Programs and Proofs - Second International Conference, CPP 2012, Kyoto, Japan,*
536 *December 13-15, 2012. Proceedings*, volume 7679 of *Lecture Notes in Computer*
537 *Science*, pages 92–108. Springer, 2012.
- 538 11. Jacob Burnim, Sudeep Juvekar, and Koushik Sen. Wise: Automated test generation
539 for worst-case complexity. In *Proceedings of the 31st International Conference*
540 *on Software Engineering, ICSE ’09*, page 463–473, New York, NY, USA, 2009.
541 Association for Computing Machinery.

- 542 12. Benjamin Canou and Alexis Darrasse. Fast and sound random generation for
543 automated testing and benchmarking in objective caml. In *Proceedings of the 2009*
544 *ACM SIGPLAN Workshop on ML, ML '09*, page 61–70, New York, NY, USA,
545 2009. Association for Computing Machinery.
- 546 13. Matthieu Carlier, Catherine Dubois, and Arnaud Gotlieb. Focaltest: A constraint
547 programming approach for property-based testing. In José Cordeiro, Maria Virvou,
548 and Boris Shishkov, editors, *Software and Data Technologies - 5th International*
549 *Conference, ICSoft 2010, Athens, Greece, July 22-24, 2010. Revised Selected*
550 *Papers*, volume 170 of *Communications in Computer and Information Science*,
551 pages 140–155. Springer, 2010.
- 552 14. Koen Claessen, Jonas Duregård, and Michał H Pałka. Generating constrained
553 random data with uniform distribution. *Journal of functional programming*, 25,
554 2015.
- 555 15. Simon Cruanes. *QuickCheck inspired property-based testing for OCaml*. <https://github.com/c-cube/qcheck>.
- 557 16. Emanuele De Angelis, Fabio Fioravanti, Adrian Palacios, Alberto Pettorossi, and
558 Maurizio Proietti. *Property-Based Test Case Generators for Free*, pages 186–206.
559 09 2019.
- 560 17. Philippe Duchon, Philippe Flajolet, Guy Louchard, and Gilles Schaeffer. Boltz-
561 mann samplers for the random generation of combinatorial structures. *Combinatorics,*
562 *Probability and Computing*, 13(4-5):577–625, 2004.
- 563 18. Peter Dybjer, Qiao Haiyan, and Makoto Takeyama. Combining testing and proving
564 in dependent type theory. volume 2758, 06 2003.
- 565 19. Frederic Peschanski et al. *Arbogen, a fast uniform random generator of tree struc-*
566 *tures*. <https://github.com/fredokun/arbogen>.
- 567 20. Patrice Godefroid, Nils Klarlund, and Koushik Sen. Dart: Directed automated
568 random testing. *SIGPLAN Not.*, 40(6):213–223, June 2005.
- 569 21. Vibhav Gogate and Rina Dechter. A new algorithm for sampling CSP solutions
570 uniformly at random. In Frédéric Benhamou, editor, *Principles and Practice of*
571 *Constraint Programming - CP 2006, 12th International Conference, CP 2006,*
572 *Nantes, France, September 25-29, 2006, Proceedings*, volume 4204 of *Lecture Notes*
573 *in Computer Science*, pages 711–715. Springer, 2006.
- 574 22. Arnaud Gotlieb and Matthieu Petit. A uniform random test data generator for
575 path testing. *J. Syst. Softw.*, 83(12):2618–2626, 2010.
- 576 23. John Hughes. Quickcheck testing for fun and profit. In Michael Hanus, editor,
577 *Practical Aspects of Declarative Languages, 9th International Symposium, PADL*
578 *2007, Nice, France, January 14-15, 2007*, volume 4354 of *Lecture Notes in Com-*
579 *puter Science*, pages 1–32. Springer, 2007.
- 580 24. Ralf Lämmel and Simon Peyton Jones. Scrap your boilerplate: A practical design
581 pattern for generic programming. *SIGPLAN Not.*, 38(3):26–37, jan 2003.
- 582 25. Leonidas Lampropoulos, Zoe Paraskevopoulou, and Benjamin C. Pierce. Gen-
583 erating good generators for inductive relations. *Proc. ACM Program. Lang.*,
584 2(POPL):45:1–45:30, 2018.
- 585 26. Sophie Laplante, Richard Lassaigne, Frédéric Magniez, Sylvain Peyronnet, and
586 Michel de Rougemont. Probabilistic abstraction for model checking: An approach
587 based on property testing. *ACM Trans. Comput. Log.*, 8(4):20, 2007.
- 588 27. Pierre Lescanne. On counting untyped lambda terms. *Theor. Comput. Sci.*,
589 474:80–97, 2013.
- 590 28. Andreas Löscher and Konstantinos Sagonas. Targeted property-based testing. In
591 *Proc. of the 26th ACM SIGSOFT International Symposium on Software Testing*
592 *and Analysis (ISSTA-17)*, pages 46–56, 07 2017.

- 593 29. Kurt Mehlhorn and Sven Thiel. Faster algorithms for bound-consistency of the
594 sortedness and the alldifferent constraint. In Rina Dechter, editor, *Principles and*
595 *Practice of Constraint Programming - CP 2000, 6th International Conference,*
596 *Singapore, September 18-21, 2000, Proceedings*, volume 1894 of *Lecture Notes in*
597 *Computer Science*, pages 306–319. Springer, 2000.
- 598 30. Yehuda Naveh, Michal Rimon, Itai Jaeger, Yoav Katz, Michael Vinov, Eitan Mar-
599 cus, and Gil Shurek. Constraint-based random stimuli generation for hardware
600 verification. *AI Mag.*, 28(3):13–30, 2007.
- 601 31. Michal Palka, Koen Claessen, Alejandro Russo, and John Hughes. Testing an opti-
602 mising compiler by generating random lambda terms. *Proceedings - International*
603 *Conference on Software Engineering*, 01 2011.
- 604 32. Zoe Paraskevopoulou, Catalin Hritcu, Maxime Dénès, Leonidas Lampropoulos,
605 and Benjamin C. Pierce. Foundational property-based testing. In Christian Urban
606 and Xingyuan Zhang, editors, *Interactive Theorem Proving - 6th International*
607 *Conference, ITP 2015, Nanjing, China, August 24-27, 2015, Proceedings*, volume
608 9236 of *Lecture Notes in Computer Science*, pages 325–343. Springer, 2015.
- 609 33. Carine Pivoteau, Bruno Salvy, and Michele Soria. Algorithms for combinatorial
610 structures: Well-founded systems and Newton iterations. *Journal of Combinatorial*
611 *Theory, Series A*, 119(8):1711–1773, November 2012.
- 612 34. Colin Runciman, Matthew Naylor, and Fredrik Lindblad. Smallcheck and lazy
613 smallcheck automatic exhaustive testing for small values. In *Proceedings of the*
614 *First ACM SIGPLAN Symposium on Haskell*, volume 44, pages 37–48, 01 2008.
- 615 35. Valerio Senni and Fabio Fioravanti. Generation of test data structures using con-
616 straint logic programming. In Achim D. Brucker and Jacques Julliand, editors,
617 *Tests and Proofs - 6th International Conference, TAP@TOOLS 2012, Prague,*
618 *Czech Republic, May 31 - June 1, 2012. Proceedings*, volume 7305 of *Lecture Notes*
619 *in Computer Science*, pages 115–131. Springer, 2012.
- 620 36. Mathieu Vavrille, Charlotte Truchet, and Charles Prud’homme. Solution sampling
621 with random table constraints. In Laurent D. Michel, editor, *27th International*
622 *Conference on Principles and Practice of Constraint Programming, CP 2021, Mont-*
623 *pellier, France (Virtual Conference), October 25-29, 2021*, volume 210 of *LIPICs*,
624 pages 56:1–56:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- 625 37. Jeremy Yallop. Practical generic programming in OCaml. pages 83–94, 01 2007.
- 626 38. Ghiles Ziat, Matthieu Dien, and Vincent Botbol. Automated Random Testing of
627 Numerical Constrained Types. In Laurent D. Michel, editor, *27th International*
628 *Conference on Principles and Practice of Constraint Programming (CP 2021)*,
629 volume 210 of *Leibniz International Proceedings in Informatics (LIPICs)*, pages
630 59:1–59:19, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für In-
631 formatik.