

Complete Symmetry Breaking for Finite Models

Marek Dančo¹, Mikoláš Janota¹, Michael Codish², João Jorge Araújo³

¹Czech Technical University in Prague, Czechia

²Ben-Gurion University of the Negev, Beer-Sheva, Israel

³Department of Mathematics, NOVA FCT, Portugal

dancomar@cvut.cz

Abstract

This paper introduces a SAT-based technique that calculates a compact and complete symmetry-break for finite model finding, with the focus on structures with a single binary operation (magmas). Classes of algebraic structures are typically described as first-order logic formulas and the concrete algebras are models of these formulas. Such models include an enormous number of isomorphic, i.e. symmetric, algebras.

A complete symmetry-break is a formula that has as models, exactly one canonical representative from each equivalence class of algebras. Thus, we enable answering questions about properties of the models so that computation and search are restricted to the set of canonical representations.

For instance, we can answer the question: How many non-isomorphic semigroups are there of size n ? Such questions can be answered by counting the satisfying assignments of a SAT formula, which already filters out non-isomorphic models. The introduced technique enables us calculating numbers of algebraic structures not present in the literature and going beyond the possibilities of pure enumeration approaches.

Introduction

Finite model finding has a longstanding tradition in automated reasoning: often a user is interested in a model rather than proving a theorem (McCune 1994). Models serve as counterexamples to invalid conjectures (Blanchette 2010) but are also interesting on their own. Indeed, Forsythe (1955) enumerates semigroups of order 4 on a computer as early as 1955. A large body of research exists that tackles finite model finding, enumeration, and counting, using constraint programming (CP) and SAT techniques (Audemard and Benhamou 2002; Distler and Kelsey 2014; Distler, Shah, and Sorge 2011; Audemard and Henocque 2001; Zhang 1996; Zhang and Zhang 1995; McCune 2003; Claessen and Sörensson 2003; Araújo, Chow, and Janota 2023; Reger, Riemer, and Suda 2019).

This paper computes compact and complete symmetry breaking constraints for a wide range of algebraic structures with a single binary operation, known as magmas. Classes of algebraic structures, such as semigroups, are typically described as first order logic formulas and the concrete algebras are models of these formulas. These models include an

enormous number of isomorphic, i.e. symmetric, algebras. For a finite algebra over a domain of size n , each permutation of the n domain elements introduces a symmetry. So there are a super-exponential number of symmetries. A complete symmetry breaking constraint for an algebraic structure is satisfied exactly for the canonical representations of the structure.

Complete symmetry breaks for algebraic structures, as computed in this paper, are not found in the literature. Applying these enables to answer, via computation, questions about properties of the models so that computation and search are restricted to the set of canonical representations.

For instance, we can answer the question: How many non-isomorphic semigroups of size n are there? — As long as we can calculate the complete symmetry-break. Questions of this type can be answered by counting the satisfying assignments of a SAT formula encoding both the property of the algebraic class and the complete symmetry break. Like this, it already filters out non-isomorphic models. This enables us to use out-of-the-box *model counters* (Gomes, Sabharwal, and Selman 2021) to calculate the non-isomorphic structures *without* enumerating them.

Our approach can be seen as “constrain and generate”, whereas pure enumeration is “generate and prune”. Applying our approach enables us to calculate numbers of algebraic structures not present in the literature and that go beyond the possibilities of pure enumeration approaches.

We mention also partial symmetry breaks. These are constraints which are satisfied by at least one element from each equivalence class of objects. Partial symmetry breaks are typically smaller in size than complete breaks but admit redundant (isomorphic) solutions.

The symmetry breaks we compute are lex-leader symmetry breaks. Namely, from each class of isomorphic (or symmetric) objects, we select a canonical representative which is the minimal object, with respect to a lexicographic ordering, in its class. So, in theory, a complete lex-leader symmetry break is obtained by introducing a constraint that imposes that an algebraic structure must not be larger than any of its permutations. A symmetry break defined in this way involves a super-exponential number of constraints and is too large to be of practical use.

In this paper we show that symmetry breaks can be often more compact in practice. However, we also show that un-

less $P=NP$, we do not expect to find a symmetry break for unrestricted algebraic structures of polynomial size.

This paper adapts the approach described in Itzhakov and Codish (2016) applied to compute compact and complete symmetry breaks for graph search problems. In fact, as graphs can be posed as an algebraic structure, our approach is a generalization of this previous work. In the case of graphs, Itzhakov and Codish (2016) compute a complete symmetry breaking constraint for order 10 graph search problems consisting of only 7,853 lexicographic order constraints instead of all $10! = 3,628,800$ constraints. This is our motivation in the context of the current paper. In our context, for example, we compute a complete symmetry break for AG-groupoids of size 7 that involves 240 lexicographic order constraints instead of $7! = 5,040$. Applying this symmetry break enables counting of all AG-groupoids of size 7, a number which is not found in current literature.

Reasoning about algebraic structures is significantly harder than about graphs, algebras correspond to matrices where the contents and row-column indices interact (one may for example have an axiom $x * x = x$). For binary operations, the search-space is bounded by n^{n^2} , which gets unwieldy as quickly as for $n = 5$ ($\approx 3e+17$), which is why $n = 6$ is already a challenge for many algebra classes.

To summarize the main contributions of this paper:

1. We generalize the approach described by Itzhakov and Codish (2016) to apply to arbitrary classes of algebraic structures described as first order logic formulas.
2. We compute compact complete symmetry breaks for a variety of algebraic structures.
3. We apply the compact complete symmetry breaks that we compute to count non-isomorphic structures of several representative classes of algebras obtaining new results.
4. We give novel theoretical insights in the complexity of complete and partial symmetry breaks for finite models.

Preliminaries

In this paper we study finite mathematical structures with a single binary operation, known as *magmas* (aka *groupoids*). For example, finite groups or semigroups are magmas. Magmas are denoted by a pair (D, \circ) where D is the domain and \circ a binary operation on D . We rely on the well-established term of isomorphism and isomorphic copy.

Definition 1 (isomorphism). A bijection $f : D_1 \rightarrow D_2$ is an isomorphism from a magma $(D_1, *)$ to (D_2, \diamond) if $f(a * b) = f(a) \diamond f(b)$, for all $a, b \in D_1$. Two magmas are *isomorphic* iff there exists at least one isomorphism between them.

Definition 2 (isomorphic copy). Consider a magma $(D_1, *)$ and a bijection $f : D_1 \rightarrow D_2$ then the *isomorphic copy* $(D_2, \diamond)_f$ is defined as $a \diamond b = f(f^{-1}(a) * f^{-1}(b))$. For a magma A we write $f(A)$ for its isomorphic copy.

Throughout the paper, we consider magmas on a finite domain $D = \{0, \dots, n-1\}$ for $n \in \mathbb{N}^+$, where n is denoted as the *size of the magma*. Then, we are only concerned with isomorphisms between magmas on the same domain D , which means that any such isomorphism is a permutation π on D ,

i.e., an element of the relevant symmetric group, which we will denote as S_n .

A finite magma can be naturally represented as a two-dimensional multiplication table, which enables us to order them lexicographically, comparing their vectorization. The table may be vectorized in an arbitrary fashion, as long as this vectorization is fixed.

Definition 3 (magma vectorization). Let A be a magma. Then $\text{vec}(A)$ is a fixed vectorization of the two-dimensional multiplication table corresponding to A .

Throughout the paper we consider three alternative vectorizations of magmas. Let $\text{vec}_r(A)$ denote the “row-by-row” concatenation of the rows of A left-to-right, top-down, i.e., the way people are used to read and write in Western civilization. Let $\text{vec}_d(A)$ denote the “diagonal first” vectorization in which the elements of the diagonal occur first, followed by the row-by-row vectorization, skipping the diagonal elements. Let $\text{vec}_c(A)$ denote the “concentric” vectorization where cells are ordered by the sum of row and column indices and ties are broken by row-by-row ordering. So for example, xor , i.e. $\begin{smallmatrix} 01 \\ 10 \end{smallmatrix}$, is vectorized as $0, 1, 1, 0$ in $\text{vec}_r(\cdot)$, $\text{vec}_c(\cdot)$ and as $0, 0, 1, 1$ in $\text{vec}_d(\cdot)$. For $x + y \pmod 3$, i.e. $\begin{smallmatrix} 012 \\ 120 \\ 201 \end{smallmatrix}$, $\text{vec}_r(\cdot)$ gives $0, 1, 2, 1, 2, 0, 2, 0, 1$, $\text{vec}_c(\cdot)$ gives $0, 1, 1, 2, 2, 0, 2, 0, 1$, and $\text{vec}_d(\cdot)$ gives $0, 2, 1, 1, 2, 1, 0, 2, 0$. These vectorizations are mutually incompatible but each individually impose a total order on algebras via the lexicographic ordering of the corresponding vectors.

Definition 4 (\preceq). Assume a vectorization $\text{vec}(\cdot)$ for magmas of size n . Define a total order $\preceq_{\text{vec}(\cdot)}$ by defining $A \preceq_{\text{vec}(\cdot)} B$ true iff $\text{vec}(A)$ is lexicographically smaller or equal than $\text{vec}(B)$. Note that $\text{vec}(A)$ is a vector of length n^2 . We usually omit the subscript and write \preceq when $\text{vec}(\cdot)$ is clear from the context.

Definition 5 (lex-leader). We say that a magma $A = (D, *)$ is a *lex-leader* with respect to a given vectorization $\text{vec}(\cdot)$, iff for all magmas B isomorphic to A , it holds that $A \preceq_{\text{vec}(\cdot)} B$.

Canonizing Sets Preliminaries

Let \mathcal{M}_n denote the set of all order n algebraic structures of a selected type (e.g. quasigroups).

The following is adapted from (Itzhakov and Codish 2016) where the definition and algorithm are stated for simple graphs.

Definition 6 (canonicity). Let $A \in \mathcal{M}_n$, $\Pi \subseteq S_n$, and denote the predicate $\text{min}_\Pi(A) = \bigwedge_{\pi \in \Pi} A \preceq \pi(A)$. We say that A is canonical if $\text{min}_{S_n}(A)$. We say that the set Π is *canonizing* if $\forall A \in \mathcal{M}_n. \text{min}_\Pi(A) \leftrightarrow \text{min}_{S_n}(A)$.

In a nutshell, a canonizing set Π is a means of *quantifier elimination* (Bradley and Manna 2007) by replacing $\forall \pi \in S_n$ by $\bigwedge_{\pi \in \Pi}$ with the aim of obtaining Π smaller than $n!$. Algorithm 1 gradually augments Π through counterexamples until it becomes canonizing. It starts with some initial set of permutations Π (for simplicity, assume that $\Pi = \emptyset$). Then, incrementally apply the step specified in lines 2–3 of Algorithm 1, as long as the stated condition holds.

Algorithm 1 Compute Canonizing Set

```
1: Init:  $\Pi = \emptyset$ 
2: while  $\exists A \in \mathcal{M}_n \exists \pi \in S_n$  s.t.  $\min_{\Pi}(A)$  and  $\pi(A) \prec A$ 
   do
3:    $\Pi = \Pi \cup \{\pi\}$ 
4: end while
5: return  $\Pi$ 
```

Lemma 7 (Itzhakov and Codish (2016)). *Algorithm 1 terminates and returns a canonizing set Π .*

We say that a canonizing set Π of permutations is *redundant* if for some $\pi \in \Pi$ the set $\Pi \setminus \{\pi\}$ is also canonizing. Algorithm 1 may compute a redundant set. For example, if a permutation added at some point becomes redundant in view of permutations added later. An algorithm *Reduce Algorithm* is straightforward to implement and we denote it as **Algorithm 2**. It iterates on the elements of a canonizing set to remove redundant permutations (similarly to the iterative algorithm for minimally unsatisfiable set or monotone predicates in general (Nöhner, Biere, and Egyed 2012; Marques-Silva, Janota, and Belov 2013)).

Canonizing Sets for Algebras

Encodings

To apply Algorithm 1 it is needed to encode into SAT that we are looking for an interpretation A that is a model of some given first order logic axioms with domain size n . This is done by standard means, cf. (McCune 1994; Claessen and Sörensson 2003). Notably, for each triple of domain elements (d_1, d_2, d_3) a propositional variable x_{d_1, d_2, d_3} is introduced, representing the fact that $d_1 * d_2 = d_3$ in the model A ; for ease of notation, we write $A[d_1, d_2] \approx d_3$ to refer to this variable. These variables encode the value of $x * y$ in *one-hot* encoding, expressed as a cardinality constraint, encoded into CNF by standard means (Roussel and Manquinho 2021), i.e. $1 = \sum_{d \in D} A[d_1, d_2] \approx d$, for $d_1, d_2 \in D$.

First, we present the encoding of the while loop condition at line 2 in Algorithm 1. The condition has four parts: $\exists A \in \mathcal{M}_n, \exists \pi \in S_n, \min_{\Pi}(A)$, and $\pi(A) \prec A$. The encoding is a conjunction: $\varphi_1 \wedge \varphi_2 \wedge \varphi_3 \wedge \varphi_4$ of four corresponding propositional formulas. The overall complexity of the encoding is $O(n^7)$, however, simplifications are used for fixed permutations see Dančo (2024, Sec 3.2.2).

The first conjunct comprises the constraints imposed by the FOL specification which specifies \mathcal{M}_n .

The second conjunct models that π is an (unknown) permutation. To this end, we introduce propositional variables $\pi[d_1] \approx d_2$ for $d_1, d_2 \in D$, to represent the permutation π , under the constraints that they indeed behave like a permutation: $1 = \sum_{d \in D} \pi[d] \approx d' = \sum_{d \in D} \pi[d'] \approx d$, for $d \in D$.

The fourth conjunct is the encoding of the constraint $\pi(A) \prec A$ where both A and π are unknown (existentially quantified). This is the crux of the condition (we will come back to the third conjunct later). Itzhakov and Codish (2016)

show how to encode this constraint for graphs but for algebras, this encoding is much more involved because the matrix contains values and permutations apply also to these values.

To reason about $\pi(A)$, it is useful to realize that if $\pi(A)$ has the value d in cell r, c , the multiplication table A will have the value $\pi^{-1}(d)$ in cell $\pi^{-1}(r), \pi^{-1}(c)$ (π is effectively a renaming). Since the SAT encoding does not enable us to treat the permutation π as a first-order citizen, we will have to go through all possible combinations of pre-images. To simplify the presentation, we introduce an auxiliary subformula pre , expressing that r', c', d' are pre-images of r, c, d under π , respectively, and the value of A is d' at the cell r', c' .

$$\text{pre}(r', r, c', c, d', d) \equiv \pi[r'] \approx r \wedge \pi[c'] \approx c \wedge \pi[d'] \approx d \wedge A[r', c'] \approx d'$$

Note that if $\text{pre}(r', r, c', c, d', d)$ is true then $\pi(A)$ has the value d at position r, c because, effectively, π replaces the primed values with their non-primed version.

We consider some fixed vectorization of the multiplication table, which is represented as a sequence of row-column index pairs, $(r_1, c_1), \dots, (r_{n^2}, c_{n^2})$. Now we can define auxiliary propositional variables that express that the value of A in cell r_i, c_i is greater/equal than $\pi(A)$ in the same cell by going through all combinations of pre-images and values.

$$\text{gt}_i^d \equiv A[r_i, c_i] \approx d \Rightarrow \bigvee_{\substack{r', c', m', m \in D \\ m < d}} \text{pre}(r', r_i, c', c_i, m', m)$$

$$\text{gt}_i \equiv \bigwedge_{d \in D} \text{gt}_i^d$$

$$\text{eq}_i^d \equiv A[r_i, c_i] \approx d \Rightarrow \bigvee_{r', c', d' \in D} \text{pre}(r', r_i, c', c_i, d', d)$$

$$\text{eq}_i \equiv \bigwedge_{d \in D} \text{eq}_i^d$$

Finally, we introduce auxiliary propositional variables r_1, \dots, r_{n^2-1} to reason about suffixes of A , i.e. if r_i is true than $\text{vec}(A) < \text{vec}(\pi(A))$ starting at index i , which lets us express the inequality between $\pi(A)$ and A .

$$\begin{aligned} \pi(A) \prec A &\equiv (\text{gt}_1 \vee (\text{eq}_1 \wedge r_1)) \wedge \\ &\bigwedge_{2 \leq i \leq n^2-1} r_{i-1} \Rightarrow (\text{gt}_i \vee (\text{eq}_i \wedge r_i)) \\ &\wedge (r_{n^2-1} \Rightarrow \text{gt}_{n^2}) \end{aligned} \tag{1}$$

The third conjunct is the encoding of the constraint $\min_{\Pi}(A) = \bigwedge_{\pi \in \Pi} A \preceq \pi(A)$. Here, Π is a set of known permutations and A is the unknown matrix from the other conjuncts. We skip the details of the encoding. For each known π in Π we have a conjunct $A \preceq \pi$ which is the negation of the encoding from the fourth conjunct, equation (1), and simpler as the permutation is known.

Complexity Insights

For graphs, complexity of lex-leader has been studied. Babai and Luks (1983) show that lex-leader adjacency matrix construction is NP-hard. Crawford et al. (1996) show that deciding whether a given *incidence* matrix is a lex-leader is NP-complete. Here we show that lex-leaders of graphs can be “simulated” in algebras. For a given undirected graph $G = (V, E)$, without self-loops, define the FOL theory F_G :

$$\forall x. c * x = x * c = c \quad (\text{force } c \text{ to } 0) \quad (2)$$

$$\forall x. x * x = c \quad (\text{nothing else is idempotent}) \quad (3)$$

$$\bigwedge_{a,b \in V \cup \{c\}} a \neq b \quad (\text{distinct vertices and } 0) \quad (4)$$

$$v_1 * v_2 \neq c \quad \{v_1, v_2\} \in E \text{ (non-0 on edges)} \quad (5)$$

$$v_1 * v_2 = c \quad \{v_1, v_2\} \notin E \text{ (0 on non-edges)} \quad (6)$$

Observation 8. Any lex-leader under the row-by-row ordering of a model with the domain $\{0, \dots, |V|\}$ of F_G corresponds to a lex-leader of the adjacency matrix of G .

Proof sketch. Since c is the only idempotent and A is a lex-leader, $A(c) = 0$, from (2), (3). Further, A will only place a non-zero value on cells with $A(v_1) *_{A} A(v_2)$ with $\{v_1, v_2\} \in E$ due to (5), (6). Therefore, the sub-table $1..|V| \times 1..|V|$ describes the adjacency matrix of G , where $A(v_1) *_{A} A(v_2) \neq 0$ iff $\{v_1, v_2\}$ is an edge in G . \square

The observation leads us to believe that deciding whether an algebra is a lex-leader is NP-hard (but we have not proven it). On the other, if there existed a canonizing set ω with polynomially number of permutations, there would be a polynomial algorithm to decide whether a given algebra is a lex-leader by testing only the permutations in ω . Hence, in the general case, we expect super-polynomial lower bounds on the size of canonizing sets.

In contrast, the following example shows the other extreme and that is, if we focus on a specific class of magmas, set of permutations can be drastically reduced.

Example 1. Consider the FOL unit clause $x * y = z * w$. For any of its model A we have $A(x, y) = d$ for all $x, y \in D$ for some fixed $d \in D$. For every domain size there exists exactly one isomorphism class and the lex-leader multiplication table consists of all zeros (with respect to any ordering of cells). There exists a minimal canonizing set for any domain of size n containing exactly the permutation consisting of a single cycle $(n - 1 \ n - 2 \ \dots \ 1 \ 0)$.

Partial Symmetry Breaks

Partial symmetry breaks for finite models are considered in the literature, notably *Least number heuristic (LNH)* (Zhang 1996; Zhang and Zhang 1995; Audemard and Henocque 2001; Araújo, Chow, and Janota 2023) is a symmetry break for finite model finding that can be used both dynamically and statically (Claessen and Sörensson 2003; Reger, Rienner, and Suda 2019). The intuition behind the break is that in a partially filled table, values that do not yet appear are indistinguishable and therefore the smallest one can be used to represent all of them. So for example, the break will restrict

\diamond	0	1	2	3
0	0	0	0	0
1	0	0	0	0
2	0	0	0	0
3	0	0	0	2

\circ	0	1	2	3
0	0	0	0	0
1	0	0	0	0
2	0	0	0	0
3	0	0	0	1

Figure 1: A that satisfies LNH and $\tau(A)$ with $\tau(A) \prec A$.

the values of the cell $(0, 0)$ to $\{0, 1\}$ because 1 is the smallest of all the yet unseen values (0 is taken to be seen already because is in the index of the cell).

Any set of permutations $\omega \subseteq S_n$ gives us a partial symmetric break for an algebra A by constructing the constraint $\bigwedge_{\pi \in \omega} A \preceq \pi(A)$. A transposition is a permutation which interchanges two elements and leaves all others fixed. We show that LNH is strictly weaker than breaking by transpositions only, i.e. by considering only $\binom{n}{2}$ elements of the symmetric group S_n . We note that breaking symmetries with transpositions is a popular technique for graph search problems and is considered in (Codish et al. 2018).

Definition 9. We say that a magma A is *minimal with respect to the set of all transpositions* iff for all transpositions $\tau \in S_n$ we have $A \preceq \tau(A)$.

Proposition 10. Any magma minimal with respect to the set of all transpositions also satisfies the LNH.

Proof. Suppose we have a magma A and the smallest i such that there exists a cell $A(r_i, c_i)$ with $A(r_i, c_i) = d$ where d is larger than the maximal designated number m . Then this model does not satisfy the LNH. Moreover, for this model also holds $\tau(A) \prec A$ for $\tau = (m \ d)$, therefore A is not minimal with respect to the set of all transpositions. \square

Example 2. Suppose $A = \langle D, \diamond \rangle$ is a magma of order 4 with $3 *^A 3 = 2$ and $x *^A y = 0$ for all other $x, y \in D$. Then A is permitted by the LNH. However, A is not minimal with respect to the set of all transpositions since for $\tau = (1 \ 2)$ and $\tau(A) = \langle D, \circ \rangle$ we have $\tau(A) \prec A$. A and $\tau(A)$ are depicted in Figure 1.

Observation 11. There is a magma, which is not transposition-minimal (and therefore not a lex-leader) but permissible under LNH (by observing Example 2).

Experiments

We implemented Algorithms 1 and 2 in Python3 using the PySAT package (Ignatiev, Morgado, and Marques-Silva 2018) with CaDiCaL 1.9.5 (Biere et al. 2020) as the back-end SAT solver. The knowledge-compilation-based tool d4 (Lagniez and Marquis 2017) was used for model counting. We chose d4 for its support of *projected* model counting, which is needed due to auxiliary variables in the encoding. The experiments were run on a server with four AMD EPYC 7513 32-Core Processor@2.6GHz and with 504 GB of memory. For each problem, we set a memory limit of 50GB and a timeout of 24 hours. We use the row-by-row ($\text{vec}_r(\cdot)$) and diagonal ($\text{vec}_d(\cdot)$) vectorization as the

Algebra class	Definition in FOL
A1:AG-groupoid	$(xy)z = (zy)x.$
A2:Comm. quasigroup	Quasigr. + $xy = yx.$
A3:Group	Monoid + $xx' = e, x'x = e$
A4:Impl. zroupoid	$(xy)z = (((ze)x)((yz)e))e,$ $(ee)e = e.$
A5:Inverse semigroup	$(xy)z = x(yz), x = xx'x,$ $x'' = x, (xx')(yy') = (yy')(xx').$
A6:Inverse property loop	Loop + $x'(xy) = y, (yx)x' = y.$
A7:Loop	Quasigr. + $ex = x, xe = x.$
A8:Magma	No requirement.
A9:Medial quasigroup	Quasigr. + $(xy)(wz) = (xw)(yz).$
A10:Monoid	$x(yz) = (xy)z, xe = x, ex = x.$
A11:Quasigr.	$xy = xz \Rightarrow y = z, yx = zx \Rightarrow y = z.$
A12:Rectang. groupoid	$(wx = yz) \Rightarrow (wx = wz).$
A13:Right involutory magma	$(xy)y = x.$
A14:Semigr.	$x(yz) = (xy)z.$

Table 1: FOL definitions of the used algebraic structures.

bases of the lexicographic ordering (see Preliminaries), the concentric vectorization behaved similarly to the row-by-row one and is not included in the experiments. For more detailed analysis of the experiments see Dančo (2024).

Several representative algebra classes with domain sizes $n \in 2..10$ are considered—we recall that the search-space of size n^{n^2} becomes unwieldy already for $n = 5$. The axiomatization of the algebra classes can be found in Table 1—for the sake of succinctness, the infix operator $*$ is omitted. Strictly speaking, some of these classes of algebras are not magmas because they are equipped with $'$. However, our approach is applicable because in all these cases, the unary operation $'$ is uniquely determined by $*$.

A group is an associative quasigroup with identity, making associative magmas (semigroups) and quasigroups the most studied magmas. Consequently, we mainly focus on these structures and several of their subclasses, including inverse semigroups, which is, after groups, the most important and studied class of semigroups. AG-groupoids, originally called *left almost semigroups*, were included in our study, as they were previously tackled by Distler, Shah, and Sorge (2011) as an important algebra class. Zroupoids were introduced as a generalization of two important algebras: De Morgan algebras (a generalization of Boolean algebras) and

type	n	#	time	#models	mc-time
A1	7	240	1,013	643,460,323,187	72,825
A2	7	96	327	6,381	186
A3	10	18	149	2	31
A4	7	173	371	600,767,308,670	48,816
A4	6	35	22	34,810,736	43
A5	10	317	18,202	169,163	10,697
A6	10	45	312	47	87
A7	7	78	40	23,746	53
A8	4	23	2	178,981,952	145
A9	10	29	2,844	19	372
A10	8	218	1,236	1,668,997	1,830
A11	6	124	101	1,130,531	717
A12	6	158	241	52,574,246	6,146
A13	6	554	411	267,954,164	12,070
A14	7	218	872	1,627,672	983

Table 2: Highest n for which model count was obtained for algebras in Table 1. The column # is the size of the canonizing set and “time” the total time of its calculation. The last two columns report on the model counting by d4. Novel results are highlighted in bold.

join-semilattices with zero (Sankappanavar 2012). Numbers for groups, magmas, and semigroups are known for larger values of n due to theoretical insights.

Rectangular groupoids generalize two other algebraic classes: rectangular semigroups and central groupoids, cf. (Boykett 2013). Right involutory magmas were introduced recently, in the context of set-theoretic solutions for the Yang-Baxter equation (Chirvasitu and Militaru 2023).

Table 2 summarizes the main achieved results. For each algebra class we list the highest size (n) where we were able to calculate the number of non-isomorphic models. In the case of implication zroupoids (A4), we list both $n = 6$ and $n = 7$ because none of these are found in the literature. The size of the canonizing set (column “#”) is measured in the number of permutations it contains.

These results highlight the power of our approach. For instance Distler, Shah, and Sorge (2011) devise a dedicated techniques to count AG-groupoids (A1) but are only able to get to size 6 (where there are already 40,104,513 algebras) but it is still based on enumeration. In contrast, our approach is able to count them for $n = 7$, which is 4 orders of magnitude larger and beyond the possibilities of enumeration.

The next values of n were not reached due to either failing to compute a canonizing set for the corresponding algebra or the model counter exceeding the time limit, with the former being the predominant case. Canonizing set computation was slightly more often limited by time than by memory (29 to 23), with memory limitations typically arising for $n = 10$ due to the complexity of the encoding. We computed canonizing sets for magmas up to $n = 6$, and for implication zroupoids and AG-groupoids up to $n = 8$.

As expected, the more structure the algebra class has, the smaller the canonizing set. An extreme case are groups (A3), where there are only 2 groups of size 10 and only 18 permutations are sufficient to break all the symmetries. In con-

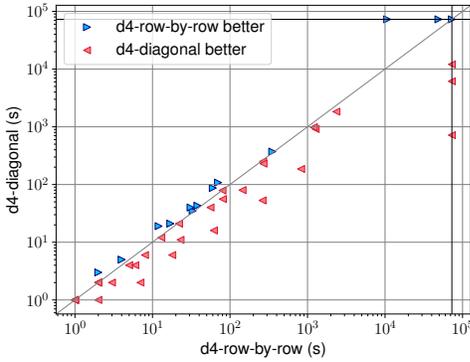


Figure 2: Comparison of the runtime d4 model counter on row-by-row and diagonal ordering.

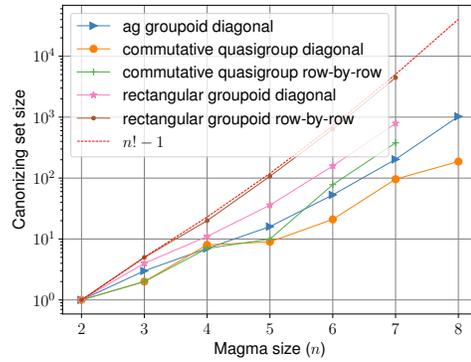


Figure 4: Sizes of canonizing set on selected algebra classes.

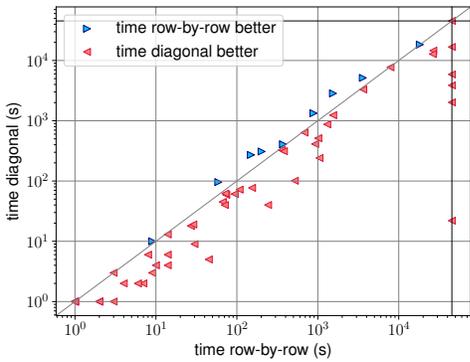


Figure 3: Comparison of the times of canonizing set on row-by-row and diagonal ordering.

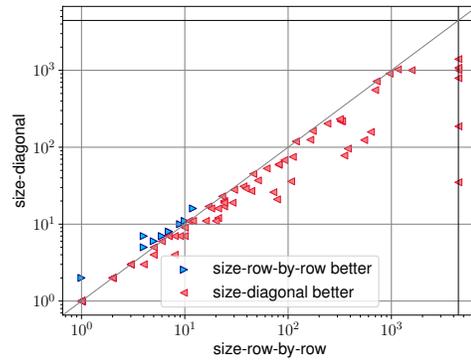


Figure 5: Comparison of the sizes of canonizing set on row-by-row and diagonal ordering.

trast, rectangular groupoids (A12) are “loosely defined”, and quickly lead to a large number of non-isomorphic models as well as the number of permutations in the canonizing set.

For several of these classes it is possible to verify the counts at *The On-Line Encyclopedia of Integer Sequences (OEIS)* (OEIS Foundation Inc. 2024). We also remark that a closed formula is known for magmas (Harrison 1966).

Figures 2–5 provide further insights into the experiments. Figure 2 presents the runtime of the model counter d4 on all the considered algebras, comparing the two considered orderings. Figure 3 shows the total times of the calculation of the canonizing set (including the reduction).

Figure 4 shows for selected algebras how the size of the canonizing set grows with n . Note that the sizes are plotted in log-scale and form almost a perfect straight line, which is indicative of exponential growth. Recall that $n! - 1$ is a trivial upper bound on the canonizing set size and the row-by-row ordering for rectangular groupoids is very close to it. In contrast, for AG-groupoids and commutative quasigroup the slope is much less steep, which enables us to get to $n = 8$. For AG-groupoids, the corresponding DIMACS has 239MB with 1,154 permutations (note that $n! = 40,320$).

Diagonal ordering gives smaller canonizing sets than the row-by-row ordering (Figure 5); it is the subject of future work to understand why. This does not always lead to better performance in the model counter; in several outliers, when the canonizing sets are not too different in size, d4 runs out of memory on the diagonal and not in the row-by-row. However, in the majority cases, the smaller canonizing set provided by the diagonal also leads to faster times in d4 (Figure 2). For example, the new result for implication zroupoids (A4, $n = 7$) and rectangular groupoids (A12, $n = 6$) were only possible under the diagonal ordering.

Related Work

Finite model finding has a longstanding tradition in automated reasoning. Sometimes, a user is interested in a model rather than proving a theorem (McCune 1994). Models serve as counterexamples to invalid conjectures (Blanchette 2010), which also appear in formal methods, e.g. in Software Verification (Torlak and Jackson 2007). Finite models can also be used as a semantic feature for *lemma selection learning* (Urban et al. 2008). A number of CP-based methods exists that enumerate (all) models, cf. Araújo, Chow, and Janota (2023). Finite models are also often constructed by

dedicated algorithms anchored in domain knowledge. The algebraic system GAP (GAP4) has a number of packages for specific types of algebraic structures. The Small Groups library (Besche, Eick, and O’Brien 2002) contains *all* ($\approx 4 \times 10^8$) non-isomorphic groups up to order 2000 (except for order 1024). Similarly, the package Smallsemi (Distler and Mitchell 2022) catalogues semigroups and the package LOOPS catalogues loops (Nagy and Vojtěchovský 2018).

Normal forms are ubiquitous in computer science and mathematics, e.g. the system `nauty` (McKay and Piperno 2014) uses canonical labeling to decide isomorphism of graphs. A large body of research exists on *symmetry breaking* in SAT and CP (Peter et al. 2014; Sakallah 2021). Computational complexity has been studied under various notions of lex-leader (Katsirelos, Narodytska, and Walsh 2010; Narodytska and Walsh 2013; Luks and Roy 2004). We are not aware, however, of any study of lex-leader in the context of constraints stemming from first order logic. Janota et al. (2024) tackle the calculation of the lex-leader for *one fixed given algebra* by using SAT.

Some symmetry breaks are designed to be fast, when used dynamically, or should add a small number of constraints, when used statically (Codish et al. 2018; Itzhakov and Codish 2020). Such symmetry breaking is often partial such as the least number heuristic (see Partial Symmetry Breaks section). Heule (2019) explores optimal complete symmetry breaking for small graphs (≈ 5 vertices) from a theoretical perspective. Kirchweger and Szeider (2021) develop a specific dynamic symmetry breaking, called *SAT Modulo Symmetries*, where a SAT solver is enhanced to look for the lexicographically smallest graph; similarly Li, Bright, and Ganesh (2024) use *orderly generation* with the objective to enumerate graphs with certain properties.

For some structures, closed forms are known for calculating the number of non-isomorphic objects. For instance, a closed formula is known for magmas (Harrison 1966), and in the same paper, the author claims that a closed formula for groups, monoids, and rings might be possible by modifying the techniques he used, but so far, nobody has managed to find those formulas.

To give an idea of how difficult it is to say something about the size sequence, we recall two old conjectures: almost all finite groups have size a power of 2; almost all semigroups are 3-nilpotent (that is, semigroups with zero, 0, satisfying the identity $xyz = 0$). The solution to the conjecture on semigroups, widely believed to be true, was announced by Kleitman, Rothschild, and Spencer (1976), but the proof has a gap that nobody could fix so far.

Given the problems with the closed formula, mathematicians turn to computational methods to find the first terms of the size sequence. Traditionally, taking advantage of the deep knowledge of some class to trim the search tree in a way that usually only works for that class. Probably the greatest achievement has been the computation of the number of order 10 semigroups, as the final piece of a long story: in 1955 it was computed up to size 4 (Forsythe 1955); in 1977, up to size 7 (Juergensen and Wick 1977); in 1994, for 8 (Satoh, Yama, and Tokizawa 1994); in 2010 for order 9 in Distler’s PhD thesis, then published in journal (Distler

and Kelsey 2014); and finally in 2012 for size 10 (Distler et al. 2012) (by using a combination of non-compact lex-leader encoding and deep understanding of the problem). Roughly speaking, once a value was computed, it took about 20 years to get the next one. The OEIS includes many more size sequences and countless pointers to the bibliography. In contrast, we introduce a general, out-of-the-box tool that improves upon the existing methodologies for determining size sequences of algebraic structures.

More broadly, this paper is related to the SAT+CAS program, where SAT is combined with *computer algebra systems*, cf. Bright, Kotsireas, and Ganesh (2018, 2022).

Conclusions and Future Work

This paper designs a method to calculate a compact complete symmetry break for finite models of first order logic theory. Such symmetry breaks open an avenue for SAT-based approaches in computational algebra. We demonstrate the strength of our approach on the problem of counting all non-isomorphic algebras of a fixed size and class. Since we encode into SAT that only canonic (lex-leader) algebras should be considered, the number of models of the produced SAT formula is equal to the number of isomorphism classes. Therefore, we can directly apply model-counting tools to count them *without* enumerating them. Here we apply exact model counters, but approximate model counting (Chakraborty, Meel, and Vardi 2021) would also be applicable if one only aims for approximate values.

Counting the structures of a given size (*size sequence* calculation) is an important and old mathematical discipline—it is no surprise that the first entry of the On-Line Encyclopedia of Integer Sequences (OEIS) is a size sequence of groups (A000001). Our approach provides a theory-agnostic tool that further advances this sub-field of universal algebra.

Model counting is not the only possible application enabled by the complete symmetry break that we calculate. It also enables further SAT-based reasoning answering further questions, e.g. “Is there an algebra with the property X?”, “Do all these algebras have a specific property?”. The paper also opens several theoretical questions. Namely, what are the lower/upper bounds for the sizes of canonizing sets for a specific class of algebras? Experimentally, the class of magmas requires $n! - 1$ permutations to provide a complete symmetry break, can this be proven for all domain sizes n ? We also intend to extend the tool to support multiple function symbols.

Acknowledgments

We thank Chad Brown and Thibault Gauthier for insightful discussions. This work was supported by MEYS through the ERC CZ program under the project *POSTMAN* no. LL1902, and by FCT — Fundação para a Ciência e a Tecnologia, I.P., via the projects UIDB/00297/2020 (doi.org/10.54499/UIDB/00297/2020) and UIDP/00297/2020 (doi.org/10.54499/UIDP/00297/2020) (Center for Mathematics and Applications), co-funded by the European Union under the project *ROBOPROX* (reg. no. CZ.02.01.01/00/22.008/0004590).

References

- Araújo, J.; Chow, C.; and Janota, M. 2023. Symmetries for Cube-And-Conquer in Finite Model Finding. In Yap, R. H. C., ed., *29th International Conference on Principles and Practice of Constraint Programming, CP 2023, August 27-31, 2023, Toronto, Canada*, volume 280 of *LIPICs*, 8:1–8:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Audemard, G.; and Benhamou, B. 2002. Reasoning by Symmetry and Function Ordering in Finite Model Generation. In Voronkov, A., ed., *Automated Deduction - CADE-18, 18th International Conference on Automated Deduction*, volume 2392 of *Lecture Notes in Computer Science*, 226–240. Springer.
- Audemard, G.; and Henocque, L. 2001. The eXtended Least Number Heuristic. In Goré, R.; Leitsch, A.; and Nipkow, T., eds., *Automated Reasoning, First International Joint Conference, IJCAR 2001, Siena, Italy, June 18-23, 2001, Proceedings*, volume 2083 of *Lecture Notes in Computer Science*, 427–442. Berlin, Heidelberg: Springer.
- Babai, L.; and Luks, E. M. 1983. Canonical labeling of graphs. In *Proceedings of the fifteenth annual ACM symposium on Theory of computing*, 171–183. ACM.
- Besche, H. U.; Eick, B.; and O’Brien, E. A. 2002. A Millennium Project: Constructing Small Groups. *Int. J. Algebra Comput.*, 12(5): 623–644.
- Biere, A.; Fazekas, K.; Fleury, M.; and Heisinger, M. 2020. CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling Entering the SAT Competition 2020. In Balyo, T.; Froylyks, N.; Heule, M.; Iser, M.; Järvisalo, M.; and Suda, M., eds., *Proc. of SAT Competition 2020 – Solver and Benchmark Descriptions*, volume B-2020-1 of *Department of Computer Science Report Series B*, 51–53. University of Helsinki.
- Blanchette, J. C. 2010. Nitpick: A Counterexample Generator for Isabelle/HOL Based on the Relational Model Finder Kodkod. In Voronkov, A.; Sutcliffe, G.; Baaz, M.; and Fermüller, C. G., eds., *Short papers for 17th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning, LPAR-17-short*, volume 13 of *EPiC Series in Computing*, 20–25. EasyChair.
- Boykett, T. 2013. Rectangular groupoids and related structures. *Discrete Mathematics*, 313(13): 1409–1418.
- Bradley, A. R.; and Manna, Z. 2007. *The calculus of computation - decision procedures with applications to verification*. Springer.
- Bright, C.; Kotsireas, I. S.; and Ganesh, V. 2018. A SAT+CAS Method for Enumerating Williamson Matrices of Even Order. In McIlraith, S. A.; and Weinberger, K. Q., eds., *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18)*.
- Bright, C.; Kotsireas, I. S.; and Ganesh, V. 2022. When satisfiability solving meets symbolic computation. *Commun. ACM*, 65(7): 64–72.
- Chakraborty, S.; Meel, K. S.; and Vardi, M. Y. 2021. Approximate Model Counting. In Biere, A.; Heule, M.; van Maaren, H.; and Walsh, T., eds., *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, 1015–1045. IOS Press.
- Chirvasitu, A.; and Militaru, G. 2023. A universal-algebra and combinatorial approach to the set-theoretic Yang-Baxter equation. arXiv:2305.14138.
- Claessen, K.; and Sörensson, N. 2003. New Techniques that Improve MACE-style Finite Model Finding. In *Proceedings of the CADE-19 Workshop: Model Computation-Principles, Algorithms, Applications*, 11–27.
- Codish, M.; Miller, A.; Prosser, P.; and Stuckey, P. J. 2018. Constraints for symmetry breaking in graph representation. *Constraints*, 24(1): 1–24.
- Crawford, J. M.; Ginsberg, M. L.; Luks, E. M.; and Roy, A. 1996. Symmetry-Breaking Predicates for Search Problems. In *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning (KR’96), Cambridge, Massachusetts, USA, November 5-8, 1996*, 148–159. Morgan Kaufmann. <https://ix.cs.uoregon.edu/%7Eluks/symmetrybreaking.pdf>.
- Dančo, M. 2024. *The Application of SAT Solving to Finite Model Finding*. Master’s thesis, Charles University, Prague, Czech Republic. <http://hdl.handle.net/20.500.11956/193190>.
- Distler, A.; Jefferson, C.; Kelsey, T.; and Kotthoff, L. 2012. The Semigroups of Order 10. In Milano, M., ed., *Principles and Practice of Constraint Programming, 18th International Conference, CP 2012, Québec City, QC, Canada, October 8-12, 2012, Proceedings*, volume 7514 of *Lecture Notes in Computer Science*, 883–899. Springer-Verlag Berlin Heidelberg.
- Distler, A.; and Kelsey, T. 2014. The Semigroups of Order 9 and Their Automorphism Groups. *Semigroup Forum*, 88: 93–112.
- Distler, A.; and Mitchell, J. 2022. Smallsemi - A library of small semigroups Version 0.6.13. <https://www.gap-system.org/Packages/smallsemi.html>. GAP package.
- Distler, A.; Shah, M.; and Sorge, V. 2011. Enumeration of AG-Groupoids. In *Intelligent Computer Mathematics*, 1–14. Springer Berlin Heidelberg. ISBN 978-3-642-22673-1.
- Forsythe, G. E. 1955. SWAC Computes 126 Distinct Semigroups of Order 4. *Proceedings of the American Mathematical Society*, 6(3): 443–447.
- GAP4. 2021. *GAP – Groups, Algorithms, and Programming, Version 4.11.1*. The GAP Group.
- Gomes, C. P.; Sabharwal, A.; and Selman, B. 2021. Model Counting. In Biere, A.; Heule, M.; van Maaren, H.; and Walsh, T., eds., *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, 993–1014. IOS Press.
- Harrison, M. A. 1966. The Number of Isomorphism Types of Finite Algebras. *Proceedings of the American Mathematical Society*, 17: 731–737.

- Heule, M. J. H. 2019. Optimal Symmetry Breaking for Graph Problems. *Math. Comput. Sci.*, 13(4): 533–548.
- Ignatiev, A.; Morgado, A.; and Marques-Silva, J. 2018. PySAT: A Python Toolkit for Prototyping with SAT Oracles. In *SAT*, 428–437.
- Itzhakov, A.; and Codish, M. 2016. Breaking Symmetries in Graph Search with Canonizing Sets. *Constraints*, 21: 357–374.
- Itzhakov, A.; and Codish, M. 2020. Incremental Symmetry Breaking Constraints for Graph Search Problems. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI*, 1536–1543. AAAI Press.
- Janota, M.; Chow, C.; Araújo, J.; Codish, M.; and Vojtěchovský, P. 2024. SAT-Based Techniques for Lexicographically Smallest Finite Models. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI*, 8048–8056. AAAI Press.
- Juergensen, H.; and Wick, P. 1977. Die Halbgruppen von Ordnungen ≤ 7 . *Semigroup Forum*, 14: 69–79.
- Katsirelos, G.; Narodytska, N.; and Walsh, T. 2010. On the Complexity and Completeness of Static Constraints for Breaking Row and Column Symmetry. In *Principles and Practice of Constraint Programming - CP*. Springer.
- Kirchweger, M.; and Szeider, S. 2021. SAT Modulo Symmetries for Graph Generation. In *CP, LIPIcs*, 34:1–34:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Kleitman, D. J.; Rothschild, B. R.; and Spencer, J. H. 1976. The Number of Semigroups of Order n . *Proceedings of the American Mathematical Society*, 55(1): 227–232.
- Lagniez, J.-M.; and Marquis, P. 2017. An Improved Decision-DNNF Compiler. In *IJCAI-17*, 667–673.
- Li, Z.; Bright, C.; and Ganesh, V. 2024. A SAT Solver and Computer Algebra Attack on the Minimum Kochev-Specker Problem (Student Abstract). In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI*, 23559–23560. AAAI Press.
- Luks, E. M.; and Roy, A. 2004. The complexity of symmetry-breaking formulas. *Annals of Mathematics and Artificial Intelligence*, 41(1): 19–45.
- Marques-Silva, J.; Janota, M.; and Belov, A. 2013. Minimal Sets over Monotone Predicates in Boolean Formulae. In Sharygina, N.; and Veith, H., eds., *Computer Aided Verification - 25th International Conference, CAV*, volume 8044 of *Lecture Notes in Computer Science*, 592–607. Springer.
- McCune, W. 1994. A Davis-Putnam program and its application to finite first-order model search: Quasigroup existence problems. Technical report, Argonne National Lab., IL (United States). Mathematics and Computer Science Div.
- McCune, W. 2003. Mace4 Reference Manual and Guide. *CoRR*, cs.SC/0310055.
- McKay, B. D.; and Piperno, A. 2014. Practical graph isomorphism, II. *J. Symb. Comput.*, 60: 94–112.
- Nagy, G.; and Vojtěchovský, P. 2018. LOOPS, Computing with quasigroups and loops in GAP, Version 3.4.1. <https://gap-packages.github.io/loops/>. Refereed GAP package.
- Narodytska, N.; and Walsh, T. 2013. Breaking Symmetry with Different Orderings. In Schulte, C., ed., *Principles and Practice of Constraint Programming*, 545–561. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-40627-0.
- Nöhner, A.; Biere, A.; and Egyed, A. 2012. Managing SAT inconsistencies with HUMUS. In *Sixth International Workshop on Variability Modelling of Software-Intensive Systems*, 83–91. ACM.
- OEIS Foundation Inc. 2024. The On-Line Encyclopedia of Integer Sequences. <https://oeis.org/>. Accessed: 2024-07-15.
- Peter, V. B.; Rossi, F.; Van Beek, P.; and Walsh, T., eds. 2014. *Handbook of constraint programming*. Foundations of Artificial Intelligence. Elsevier Science & Technology.
- Reger, G.; Riemer, M.; and Suda, M. 2019. Symmetry Avoidance in MACE-Style Finite Model Finding. In Herzig, A.; and Popescu, A., eds., *Frontiers of Combining Systems - 12th International Symposium, FroCoS*, volume 11715 of *Lecture Notes in Computer Science*, 3–21. Springer.
- Roussel, O.; and Manquinho, V. M. 2021. Pseudo-Boolean and Cardinality Constraints. In Biere, A.; Heule, M.; van Maaren, H.; and Walsh, T., eds., *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, 1087–1129. IOS Press.
- Sakallah, K. A. 2021. Symmetry and Satisfiability. In Biere, A.; Heule, M.; van Maaren, H.; and Walsh, T., eds., *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, 509–570. IOS Press.
- Sankappanavar, H. P. 2012. De Morgan Algebras: New Perspectives and Applications. *Scientiae Mathematicae Japonicae*, 75(1): 21–50.
- Satoh, S.; Yama, K.; and Tokizawa, M. 1994. Semigroups of Order 8. *Semigroup Forum*, 49: 7–29.
- Torlak, E.; and Jackson, D. 2007. Kodkod: A Relational Model Finder. In Grumberg, O.; and Huth, M., eds., *TACAS*, volume 4424, 632–647. Springer.
- Urban, J.; Sutcliffe, G.; Pudlák, P.; and Vyskočil, J. 2008. MaLAREa SG1- Machine Learner for Automated Reasoning with Semantic Guidance. In *Automated Reasoning, 4th International Joint Conference, IJCAR*. Springer.
- Zhang, J. 1996. Constructing Finite Algebras with FALCON. *Journal of automated reasoning*, 17: 1–22.
- Zhang, J.; and Zhang, H. 1995. SEM: a System for Enumerating Models. In *IJCAI*, 298–303.