

Satsuma: Structure-Based Symmetry Breaking in SAT

Markus Anders ✉
TU Darmstadt, Germany

Sofia Brenner ✉ 
TU Darmstadt, Germany

Gaurav Rattan ✉ 
University of Twente, Enschede, The Netherlands

Abstract

Symmetry reduction is crucial for solving many interesting SAT instances in practice. Numerous approaches have been proposed, which try to strike a balance between symmetry reduction and computational overhead. Arguably the most readily applicable method is the computation of static symmetry breaking constraints: a constraint restricting the search-space to non-symmetrical solutions is added to a given SAT instance. A distinct advantage of static symmetry breaking is that the SAT solver itself is not modified. A disadvantage is that the strength of symmetry reduction is usually limited. In order to boost symmetry reduction, the state-of-the-art tool BREAKID [Devriendt et. al] pioneered the identification and tailored breaking of a particular substructure of symmetries, the so-called row interchangeability groups.

In this paper, we propose a new symmetry breaking tool called SATSUMA. The core principle of our tool is to exploit more diverse but frequently occurring symmetry structures. This is enabled by new practical detection algorithms for row interchangeability, row-column symmetry, Johnson symmetry, and various combinations. Based on the resulting structural description, we then produce symmetry breaking constraints. We compare this new approach to BREAKID on a range of instance families exhibiting symmetry. Our benchmarks suggest improved symmetry reduction in the presence of Johnson symmetry and comparable performance in the presence of row-column symmetry. Moreover, our implementation runs significantly faster, even though it identifies more diverse structures.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis

Keywords and phrases symmetry breaking, boolean satisfiability, graph isomorphism

Digital Object Identifier 10.4230/LIPIcs.SAT.2024.4

Related Version *Preprint*: <https://arxiv.org/abs/2406.13557>

Supplementary Material *Software (source code)*: <https://github.com/markusa4/satsuma> [2]
archived at [swh:1:dir:134ea7952a25d92ba2addf7ff25aa18183550735](https://swh.1:dir:134ea7952a25d92ba2addf7ff25aa18183550735)

Funding The research leading to these results has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (EngageS: grant agreement No. 820148). Sofia Brenner additionally received funding from the German Research Foundation DFG (SFB-TRR 195 “Symbolic Tools in Mathematics and their Application”).

1 Introduction

Symmetries are present in many interesting SAT instances, ranging from hard combinatorial problems to circuit design. Making use of symmetry is paramount in order to efficiently solve many of these instances. Practical approaches for symmetry reduction must always strike a balance between the computational overhead incurred and the strength of the symmetry reduction. Two decades of research have led to many approaches to tackle this problem [17, 1, 21, 36, 22, 20, 34, 26, 30]. At one end of the spectrum, isomorph-free generation techniques [26, 30] apply sophisticated algorithms in conjunction with the solver, such that a



© Markus Anders, Sofia Brenner, and Gaurav Rattan;
licensed under Creative Commons License CC-BY 4.0

27th International Conference on Theory and Applications of Satisfiability Testing (SAT 2024).

Editors: Supratik Chakraborty and Jie-Hong Roland Jiang; Article No. 4; pp. 4:1–4:23

Leibniz International Proceedings in Informatics



LIPIC Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

solver only explores asymmetric branches of the search. While these techniques are successful in solving hard combinatorial instances (e.g., [29]), this comes at the price of substantial overhead: both in terms of computational cost as well as interfering with the other strategies employed by solvers. Hence, one must be sure that the symmetry reduction is worth the additional overhead. It therefore seems impractical to turn these techniques “on-by-default”.

Arguably at the other end are tools producing *static symmetry breaking constraints* [17, 1, 21]. These tools add additional clauses and variables to a given instance, with the aim of reducing the number of symmetric branches explored by the solver. While the symmetry reduction is usually not as strong as for dynamic techniques, such constraints can be computed comparatively cheaply. More importantly, a distinct advantage of static symmetry breaking constraints is that the SAT solver itself is not modified, and hence there is a complete separation of concerns. State-of-the-art static symmetry breaking tools are successfully applied as an “on-by-default” technique [14, 21]. Static symmetry breaking is not only used in SAT, but also in various other areas of constraint programming [25, 35, 19, 8].

In order to improve symmetry reduction, a rather recent development in static symmetry breaking is to detect and make use of so-called *row interchangeability* subgroups [21, 35]. In SAT, this feature was introduced by the state-of-the-art symmetry breaking tool BREAKID [21], but it is also used in symmetry breaking in mixed integer programming (MIP) [35]. Row interchangeability groups stem from a natural modeling of the variables as a matrix in which all *rows* are interchangeable by a symmetry. The idea is to first identify these row interchangeability groups, and then produce tailored symmetry breaking constraints. The current generation of tools identifies row interchangeability by hoping for and exploiting a particular structure in the generators of the symmetry group. However, the method is not guaranteed to work and sometimes incurs significant overhead [5]. Despite this, the gain in symmetry reduction seems to be worth the trade-off [21].

In the realm of constraint programming, symmetry breaking constraints for more structures have been considered: for example, *row-column symmetry* [23] is a natural extension of row interchangeability, where both the rows and columns are interchangeable. These symmetries are common in combinatorics, scheduling, or assignment problems [24, 23], such as the well-known pigeonhole principle. While traditional complete symmetry breaking constraints are unlikely to be efficiently computable for these structures [23], different practical constraints are well-studied in the literature for *manual* breaking of symmetries [23, 28]. Another area in which symmetry breaking has been studied in detail is graph generation [15, 16, 30]. A typical problem in this area is to determine the existence of a graph with a specific property. Symmetries in these problems often simply correspond to *isomorphic graphs*. Even though this is rarely mentioned explicitly, such symmetries can be described by so-called *Johnson groups* [9, 10].

In *automated* symmetry breaking, making use of such results requires us to *identify* the appropriate structures first. However, generalizing the existing identification strategies of contemporary tools to more elaborate structures seems elusive.

1.1 Contribution

We present a new algorithm for the generation of symmetry breaking constraints, and a prototype implementation called SATSUMA. Our goal is to explore whether the approach of “identifying and exploiting specific group structures” can be pushed further.

Techniques. We place the identification of specific symmetry groups at the very heart of SATSUMA. The approach is enabled by our main contribution, a new class of practical detection algorithms. In particular, we provide algorithms identifying row interchangeability

(Section 3.1), row-column symmetry (Section 3.2), and Johnson symmetry (Section 3.3). Furthermore, we detect certain *combinations* of the above groups, as well as groups which are *similar* to the above groups, building essentially a *structural description* of the group. Symmetry breaking constraints are then chosen based on the type of detected structure: for each detected structure, we determine a set of carefully chosen symmetries, for which conventional symmetry breaking constraints are produced.

Our detection algorithms are all based on the highly efficient *individualization-refinement* framework, as is commonly used in practical graph isomorphism algorithms [33]. Our detection algorithms are all heuristics, in that identification of a particular group cannot be guaranteed. However, the success of the heuristics provably depends only on a well-studied graph property (see Section 2).

These algorithms can be applied *without* computing the symmetries of the formula first: they are purely graph-based. We exploit this by first running our tailored detection algorithms, and then only apply general-purpose symmetry detection on parts not yet identified. In order to handle this remainder, SATSUMA reimplements parts of BREAKID. Essentially, our new approach acts as a preprocessor for existing techniques.

Benchmarks. We compare SATSUMA and BREAKID on a range of well-established SAT instance families exhibiting symmetry. In our benchmarks, we observe that our new structure-based implementation

1. leads to improved SAT solver performance on instances with Johnson symmetries,
2. comparable SAT solver performance on instances exhibiting predominantly row interchangeability or row-column symmetry,
3. and incurs less computational overhead on all tested benchmark families (we observe better asymptotic scaling of SATSUMA on some benchmark families).

When SATSUMA detects a structure, our new approach seems to be a win-win: it yields lower computational overhead, and the resulting speed-up for SAT solvers is comparable or better.

2 Preliminaries

2.1 Satisfiability and Symmetry

SAT. In this paper, a SAT formula F in *conjunctive normal form* (CNF) is denoted with $F = \{\{l_{1,1}, \dots, l_{1,k_1}\}, \dots, \{l_{m,1}, \dots, l_{m,k_m}\}\}$. Each element $C \in F$ is called a *clause*, whereas a clause itself consists of a set of *literals*. A literal is either a variable v or its negation $\neg v$. We write $\text{Var}(F) := \{v_1, \dots, v_n\}$ for the set of *variables* of F and use $\text{Lit}(F)$ for its literals.

A symmetry, or *automorphism*, of F is a permutation $\varphi: \text{Lit}(F) \rightarrow \text{Lit}(F)$ satisfying the following properties. First, it maps F to itself, i.e., $F^\varphi = F$, where F^φ means applying φ element-wise to the literals in each clause. Second, for all $l \in \text{Lit}(F)$ it holds that $\neg\varphi(l) = \varphi(\neg l)$. We define the *support* of φ as $\text{supp}(\varphi) = \{l \in \text{Lit}(F) : l^\varphi \neq l\}$, i.e., the set of all literals moved by φ . The set of all symmetries of F is $\text{Aut}(F)$. We can efficiently *test* if a permutation φ is an automorphism of F : for each clause C , we check whether $C^\varphi \in F$ holds.

An *assignment* of F is a function $\theta: \text{Var}(F) \rightarrow \{0, 1\}$. We define the evaluation of F under θ in the usual way, i.e., either $F[\theta] = 1$ or $F[\theta] = 0$ holds. A formula F is satisfiable if there exists an assignment θ with $F[\theta] = 1$, and unsatisfiable otherwise. Given an assignment θ of F and an automorphism $\varphi \in \text{Aut}(F)$, we define $\theta^\varphi(v) := \theta(v')$ if $\varphi(v) = v'$ for $v' \in \text{Var}(F)$ and $\theta^\varphi(v) := \neg\theta(v')$ if $\varphi(v) = \neg v'$ for $v' \in \text{Var}(F)$, where naturally $\neg 0 = 1$ and $\neg 1 = 0$. It follows readily that for $\varphi \in \text{Aut}(F)$, we have $F[\theta] = F^\varphi[\theta^\varphi] = F[\theta^\varphi]$.

Symmetry Breaking Constraints. All symmetry breaking constraints in this paper are *lex-leader* constraints. Let \prec denote a total order of $\text{Var}(F)$. We order an assignment θ according to \prec , yielding a $\{0, 1\}$ -string. We can then order assignments θ, θ' of F lexicographically by comparing their corresponding strings, denoted by \prec_{lex} . Given an automorphism φ of F , it suffices to evaluate F on those assignments θ for which $\theta^\varphi \preceq_{\text{lex}} \theta$ holds, since $F[\theta^\varphi] = F[\theta]$. In particular, we may add a *lex-leader constraint* LL_φ^\prec to F , which ensures that $\theta^\varphi \preceq_{\text{lex}} \theta$ holds. It is easy to see that F is satisfiable, if and only if $F \wedge_{\varphi \in \text{Aut}(F)} \text{LL}_\varphi^\prec$ is satisfiable [37]. Lex-leader constraints can be efficiently encoded as a CNF formula, and different encodings have been studied in detail [1, 21]. The practical encoding we use is reverse-engineered from BREAKID, and is described in [21]. Having detected structures of symmetries, SATSUMA attempts to determine a favorable variable order and a set of automorphisms for which lex-leader constraints are constructed (see Section 4).

2.2 Graphs and Symmetry

Graphs. An undirected graph $G = (V, E)$ consists of a vertex set V and an edge relation $E \subseteq \binom{V}{2}$. We refer to the set of vertices of G as $V(G)$, and to the set of edges as $E(G)$. A *vertex coloring* of G is a mapping $\pi : V(G) \rightarrow [k]$ to colors in $[k]$ for some $k \in \mathbb{N}$. We call (G, π) a vertex-colored graph. The *color class* of a color c consists of all vertices of G with color c . The color classes form a partition of $V(G)$, the *color partition* corresponding to π .

A bijection $\varphi : V(G) \rightarrow V(G)$ is called an *automorphism* of (G, π) , whenever $(G, \pi)^\varphi = (G^\varphi, \pi^\varphi) = (G, \pi)$ holds. Here, G^φ denotes the graph with vertex set $V(G)$ and edges $\{u^\varphi, v^\varphi\}$ whenever $\{u, v\}$ is an edge of G (where v^φ simply denotes the image of v under φ). The coloring π^φ is given by $\pi^\varphi(v) = \pi(v^\varphi)$ for every $v \in V(G)$. The set of all automorphisms of (G, π) is denoted by $\text{Aut}(G, \pi)$.

For a given CNF formula F , we define the *model graph* $G(F) = (G, \pi)$ as follows. The vertex set consists of the literals and clauses of F . There are edges connecting the literals of a common variable to each other. Clauses are connected to the literals they contain. Formally, let $E := \{\{v, \neg v\} : v \in \text{Var}(F)\} \cup \{\{C, l\} : l \in C, C \in F\}$. Define a coloring π by setting $\pi(l) := 0$ for all literals $l \in \text{Lit}(F)$ and $\pi(C) := 1$ for all clauses $C \in F$. It is well-known that the automorphisms of $G(F)$ restricted to $\text{Lit}(F)$ are precisely the automorphisms of F [37].

Permutation Groups. We recall some notions of permutation group theory. A detailed account can be found in [38]. Let Ω be a nonempty finite set. Let $\text{Sym}(\Omega)$ denote the *symmetric group* on Ω , i.e., the group of permutations of Ω , and set $\text{Sym}(n) := \text{Sym}([n])$. A *permutation group* is a subgroup Γ of $\text{Sym}(\Omega)$, denoted by $\Gamma \leq \text{Sym}(\Omega)$. We also say that Γ *acts on* Ω . For $g \in \Gamma$ and $\omega \in \Omega$, we write ω^g for the image of ω under g and $\omega^\Gamma = \{\omega^g : g \in \Gamma\}$ for the *orbit* of ω under Γ . In other words, ω^Γ consist of all points in Ω that can be reached from ω by applying elements of Γ . The partition of Ω into the orbits of Γ is called the *orbit partition*. For $\omega \in \Omega$, let $\Gamma_\omega := \{g \in \Gamma : \omega^g = \omega\}$ denote the *stabilizer* of ω in Γ . In other words, Γ_ω consists of those elements in Γ that map ω to itself. The *direct product* of permutation groups Γ_1 and Γ_2 which act on domains Ω_1 and Ω_2 , respectively, is the Cartesian product $\Gamma_1 \times \Gamma_2$, endowed with a component-wise multiplication. It naturally acts component-wise on $\Omega_1 \times \Omega_2$.

Individualization-Refinement. A central ingredient in our algorithms is the so-called *individualization-refinement* (IR) paradigm. The IR paradigm is the central technique in all state-of-the-art symmetry detection algorithms [33, 18, 27, 4], and highly engineered



■ **Figure 1** An illustration of the IR process. Individualization steps break symmetries or similarities (nodes marked with a cross are individualized). Refinement steps propagate this information.

implementations are available. The paradigm mainly consists of the *individualization* technique, paired with the so-called *color refinement algorithm*. In this paragraph, we focus on a high-level explanation of the routine. A detailed account can be found in [33].

The central idea of IR (see Figure 1 for an illustration) is the following: given a vertex-colored graph (G, π) and a vertex $v \in V(G)$, the vertex v is *individualized*. Basically this means that it obtains a new color. The routine then proceeds with a so-called *color refinement*: in each step, every vertex of G obtains a new color, based on its former color together with the colors of its neighbors in G . This recoloring procedure is repeated until the corresponding color partition stabilizes. The final coloring π' is then returned. We use the notation $\pi' = \text{IR}((G, \pi), v)$ to denote this process. The call $\text{IR}((G, \pi), v)$ can be computed in time $\mathcal{O}(|E(G)| \log |V(G)|)$ (see [11]).

The coloring π' is a *refinement* of π in the sense that vertices with the same color in π' already had the same color in π . In other words, a color c of π is either preserved in π' , or partitioned into several other colors c_1, \dots, c_n . For $i \in [n]$, we call the sets $\{u \in V(G) : \pi(u) = c, \pi'(u) = c_i\}$ the *fragments* of c in π' . The second crucial observation is that vertices in the same orbit under the stabilizer $\text{Aut}(G, \pi)_v$ obtain the same color in π' . However, it is possible that the color partition of π' is *coarser* than the orbit partition in the sense that the vertices of multiple orbits might obtain the same color in π' .

Clearly, this process can be applied inductively to individualize multiple vertices. It is also possible to pass the empty sequence ε to IR, i.e., to run only the color refinement procedure. Arguing as above, the resulting color partition is guaranteed to be at least as coarse as the orbit partition of $\text{Aut}(G, \pi)$ (i.e., the stabilizer of the empty sequence).

The next lemma summarizes the properties of IR to which we refer throughout the paper:

► **Lemma 1.** *Given a vertex-colored graph (G, π) and a vertex $v \in V(G)$, the refined coloring $\pi' = \text{IR}((G, \pi), v)$ has the following properties.*

1. *The coloring π' is a refinement of π : for $u, w \in V(G)$ with $\pi'(u) = \pi'(w)$, we have $\pi(u) = \pi(w)$.*
2. *The color partition of π' is at least as coarse as the orbit partition of $\Gamma = \text{Aut}(G, \pi)_v$: vertices $u, w \in V(G)$ with $\pi'(u) \neq \pi'(w)$ lie in different orbits of Γ , i.e., we have $w \notin u^\Gamma$.*
3. *The colors of π' are isomorphism-invariant: for every $\varphi \in \text{Sym}(V(G))$, it holds that $\text{IR}((G^\varphi, \pi^\varphi), v^\varphi) = \text{IR}((G, \pi), v)^\varphi$. In particular, if $\varphi \in \text{Aut}(G, \pi)$, then $\text{IR}((G, \pi), v)^\varphi = \text{IR}((G^\varphi, \pi^\varphi), v^\varphi) = \text{IR}((G, \pi), v)^\varphi$ holds.*

These properties follow almost immediately from the definition of IR, and we refer to [33] for a treatment of the topic. We also mention that usually, as opposed to the description above, IR is defined for *sequences* of vertices instead of single vertices.

We now recall the notion of Tinhofer graphs [7]. In view of the second part of Lemma 1, these are precisely the graphs for which the two partitions coincide.

► **Definition 2** (Tinhofer Graph [7, 40]). *A graph G is called Tinhofer if for all $v \in V(G)$, the orbit partition of $\Gamma := \text{Aut}(G, \pi)_v$ coincides with the color partition of $\pi' := \text{IR}((G, \pi), v)$ and the same applies recursively to the colored graph (G, π') (this corresponds to individualizing multiple vertices of G). Formally, the first property means that for all $u, w \in V(G)$, we have $w \in u^\Gamma$ if and only if $\pi'(u) = \pi'(w)$.*

In particular, IR works well on Tinhofer graphs: practical graph isomorphism solvers are guaranteed to not require any backtracking.

2.3 Symmetry Structures in SAT

The idea of our tool is to detect certain symmetry structures that are subsequently exploited. In this section, we describe the main structures detected by the tool. The description of the detection algorithms is the subject of Section 3.

Throughout, let F be a SAT formula. As a first step, consider the *disjoint direct decomposition* of the symmetries $\text{Aut}(F)$: this is a partition $\text{Lit}(F) = L_1 \dot{\cup} \dots \dot{\cup} L_k$ of $\text{Lit}(F)$ for which there exists a decomposition $\text{Aut}(F) = A_1 \times \dots \times A_k$ into a direct product of subgroups such that, for every $i \in [k]$, the automorphisms in A_i only move the literals in L_i . A disjoint direct decomposition naturally decomposes the symmetry breaking problem, and it suffices to treat each factor separately. In the following, we always refer to the finest such decomposition, which is clearly unique. We call its parts L_1, \dots, L_k the *disjoint direct factors* of F . Note that every disjoint direct factor is a union of orbits of $\text{Aut}(F)$.

As factors in the disjoint direct decomposition, we detect several variants of three main kinds of symmetries, namely row symmetry, row-column symmetry, and Johnson symmetry. Let us now define these notions in the special context of CNF formulas.

Row Symmetry. *Row interchangeability*, or *row symmetry*, naturally occurs in the context of matrix modeling [24] and is already successfully exploited in automated symmetry breaking. We say that a SAT formula F *exhibits row symmetry* if there exists a disjoint direct factor $L \subseteq \text{Lit}(F)$ which can be arranged in a matrix M such that $\text{Aut}(F)|_L$ acts by permuting the rows of M . In addition, we require that every column of M is an orbit of $\text{Aut}(F)$. See Figure 2a for an illustration. The colored boxes illustrate orbits, whereas dashed lines indicate vertices in the same row. The rows can be permuted using symmetry.

We should address a technical difference between the definition above and how BREAKID handles row symmetry: in our definition, a disjoint direct factor should *only* admit the action of the row symmetry group, or a particularly defined extension (see Section 3). BREAKID on the other hand would accept any row symmetry *subgroup* that it detects (see [5] for further discussion). Hence, in practice, it may happen that BREAKID reports row symmetry, when SATSUMA does not. However, SATSUMA may instead identify a larger, more expressive group, such as row-column symmetry, as explained below.

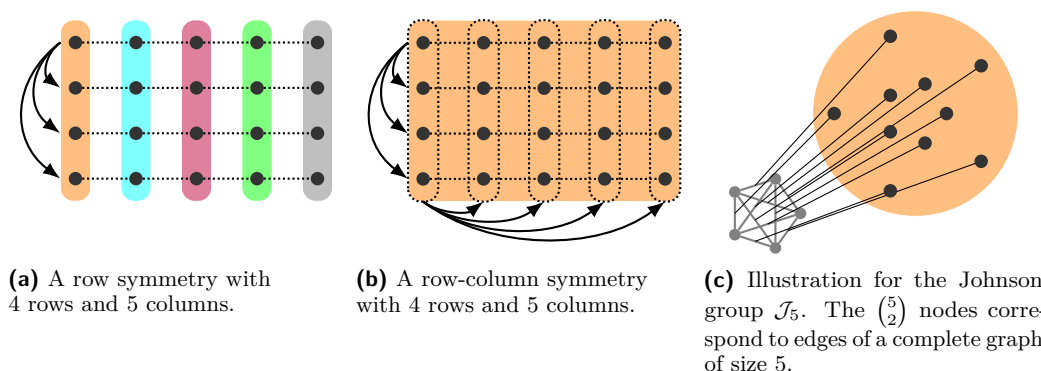
Let us make a general observation regarding negation symmetry.

► **Remark 3.** For an orbit σ of literals under $\text{Aut}(F)$, also the set $\neg\sigma := \{\neg v : v \in \sigma\}$ is an orbit of literals. Hence two cases can occur: either we have $\sigma = \neg\sigma$, or the orbits σ and $\neg\sigma$ are disjoint.


In order to simplify the exposition, we only consider the second scenario in the following.

Row-column symmetry. Row-column symmetries are an extension of row interchangeability. Row-column symmetry naturally occurs whenever both the rows and columns of a matrix of variables are interchangeable. Examples can be found in scheduling, design, and combinatorial problems (see [24]).

For $m, n \in \mathbb{N}$, the *row-column symmetry* group is $\Gamma := \text{Sym}(n) \times \text{Sym}(m)$, acting componentwise on $[n] \times [m]$. We think of $[n] \times [m]$ as an $n \times m$ matrix M , on which $(\sigma_1, \sigma_2) \in \Gamma$ acts by permuting the n rows according to σ_1 and the m columns according to σ_2 .



■ **Figure 2** Various group structures used throughout the paper. Colors indicate orbits of the group.




■ **Figure 3** An illustration of the Johnson group \mathcal{J}_5 .

A SAT formula F exhibits row-column symmetry if there exists a disjoint direct factor $L \subseteq \text{Lit}(F)$ consisting of an orbit σ of $\text{Aut}(F)$ and its negation $\neg\sigma$ such that the following holds: the literals in σ can be arranged in an $n \times m$ -matrix M such that $\text{Aut}(F)|_\sigma$ acts as a row-column symmetry group on M . See Figure 2b for an illustration. Note that the action of $\text{Aut}(F)$ on σ naturally extends to a row-column symmetry action on $\neg\sigma$. For this reason, our algorithm generates the matrix M of the literals in σ and extends this to $\neg\sigma$, see Section 3.2 for details.

Johnson symmetry. Johnson groups are naturally tied to the graph isomorphism problem. Whenever a problem asks for the existence of an undirected graph with a certain property, typically, the underlying symmetries form a Johnson group.

Observe that $\pi \in \text{Sym}(n)$ induces a permutation on the domain $\binom{[n]}{2}$ of 2-subsets of $[n]$, mapping $\{a_1, a_2\}$ to $\{a_1^\pi, a_2^\pi\}$. This way, $\text{Sym}(n)$ becomes a permutation group on a domain of size $\binom{[n]}{2}$, the Johnson group \mathcal{J}_n . Technically, these groups are specifically Johnson groups of arity 2. The corresponding action is called a Johnson action.

We now define Johnson symmetries for SAT formulas. Intuitively, the variables the formula correspond to the “edges” (i.e., sets of two vertices) of a complete graph. There is a symmetric action on the “vertices” of this underlying graph and the variables of the formula (“edges”) are permuted accordingly. See Figure 2c and Figure 3 for an illustration. Formally, a SAT formula F exhibits a Johnson symmetry if the following holds: there exists



(a) A row symmetry with 4 rows and 5 columns.

(b) Individualizing a vertex of a row identifies the entire row.

(c) A row symmetry on blocks of size 2.

■ **Figure 4** Illustrations of different aspects of row symmetry.

a disjoint direct factor $L \subseteq \text{Lit}(F)$ consisting of an orbit σ of $\text{Aut}(F)$ and its negation $\neg\sigma$ such that the literals in σ can be relabeled as $x_{\{i,j\}}$ for all $\{i,j\} \in \binom{[n]}{2}$ and $\text{Aut}(F)|_\sigma$ acts as the Johnson group \mathcal{J}_n (by permuting the index sets). Again, the action of $\text{Aut}(F)$ naturally extends to $\neg\sigma$.

3 Detection Algorithms

We now present our detection algorithms. All algorithms are centered around detecting structure on the model graph $G(F)$ of a given CNF formula F . Recall that $G(F)$ contains a vertex for each literal, so we may use these terms interchangeably. The major design principles of our algorithms are described in the following.

Colors are Orbits. Our algorithms work on the assumption that the model graph $G(F)$ is Tinhofer (see Definition 2). Then we can compute orbits of stabilizers using IR. In particular, the color classes of $\pi = \text{IR}(G(F), \varepsilon)$ are then the orbits of $\text{Aut}(G(F))$.

Certified Correctness. The input model graph might not be Tinhofer. However, each algorithm constructs a carefully chosen set of candidate permutations, which suffices to prove the existence of a certain group action. It is then verified that these permutations are automorphisms of the formula F , which ensures correctness. In our implementation, we produce lex-leader constraints only for automorphisms verified on the original formula.

Color-by-color. All of our detection algorithms proceed color-by-color, or orbit-by-orbit: given an orbit, the algorithms stabilize a specific set of points, observing the effect on the given orbit as well as other orbits. If an orbit exhibits a specific group action, then this effect is clearly defined, and a model of the purported structure is made.

3.1 Row Symmetry

We describe an algorithm for row symmetry. First, we define an auxiliary function that transposes two pair-wise disjoint lists of literals of equal length: For $l \in \text{Lit}(F)$, let

$$\text{transpose}_F((l_1 \dots l_k), (l'_1 \dots l'_k))_F(l) := \begin{cases} l'_i & \text{if } l = l_i \text{ with } i \in [k] \\ l_i & \text{if } l = l'_i \text{ with } i \in [k] \\ l & \text{otherwise.} \end{cases}$$

■ **Algorithm 1** Detection algorithm for row symmetry.

```

1 function DetectRowSymmetry
  Input : > formula  $F$ 
         > set  $\sigma \subseteq \text{Lit}(F)$  with  $|\sigma| \geq 3$ 
  Output : < matrix with row symmetry including  $\sigma$ , or  $\perp$  if check fails
2  $(G, \pi) := G(F)$ ,  $\pi' := \text{IR}((G, \pi), \varepsilon)$ ;
3 // construct a candidate row for each  $v \in \sigma$ 
4 foreach  $v \in \sigma$  do
5    $\pi_v := \text{IR}((G, \pi'), v)$ ;
6   let  $\tau$  be a list of literals that are singletons in  $\pi_v$  but not in  $\pi'$ ;
7   sort literals in  $\tau$  according to their color in  $\pi_v$ ;
8    $\text{row}[v] := \tau$ ;
9 check that rows are pair-wise disjoint;
10 // verify that  $M$  exhibits row symmetry
11 foreach  $i \in \{1 \dots |\sigma| - 1\}$  do
12    $v := \sigma[i - 1]$ ;  $v' := \sigma[i]$ ;
13   check that  $\text{transpose}_F(\text{row}[v], \text{row}[v'])$  is a symmetry of  $F$ ;
14 return matrix  $M$  constructed from row

```

(Description of Algorithm 1). For an illustration, see Figure 4a. The algorithm applies IR for each $v \in \sigma$ (see Figure 4b). All vertices v' in other orbits which are individualized in this process, i.e., which are fixed once v is fixed, are added to the purported “row” of v . We then verify that every row transposition of the resulting matrix is indeed a symmetry of F .

(Correctness of Algorithm 1). We first make the following observation for orbits of stabilizers in row interchangeability groups.

► **Lemma 4.** *Let $\Gamma = \text{Sym}(n)$ be a row interchangeability group acting on $[n] \times [m]$. For every $(i, j) \in [n] \times [m]$, the orbit of $(k, l) \in [n] \times [m]$ under the stabilizer $\Gamma_{(i,j)}$ of (i, j) is given by*

$$(k, l)^{\Gamma_{(i,j)}} = \begin{cases} \{(i, l)\} & \text{if } k = i \\ ([n] \setminus \{i\}) \times \{l\} & \text{otherwise.} \end{cases}$$

Proof. Interpreting $[n] \times [m]$ as $n \times m$ -matrix M , recall that Γ acts by permuting the rows of M . In other words, the stabilizer $\Gamma_{(i,j)}$ consisting of all row permutations that fix the i -th row and permute the other rows arbitrarily. Now consider the orbit of $(k, l) \in [n] \times [m]$ under the stabilizer $\Gamma_{(i,j)}$. If $k = i$, then (k, l) can only be mapped to elements in the same row as $\Gamma_{(i,j)}$ fixes the i -th row of M . On the other hand, since Γ acts by permuting the rows, every element of M can only be mapped to elements in the same column, that is, (k, l) must be fixed. Similarly, for $k \neq i$, the element (k, l) can be mapped to all elements in the l -th column except for (i, l) . ◀

Next, we prove that the algorithm always returns correct symmetries of F and that in case the model graph is Tinhofer, the algorithm is guaranteed to detect row interchangeability groups.

► **Theorem 5.** *Let F be a SAT formula.*

1. *If Algorithm 1 returns a matrix M , every row permutation of M is a symmetry of F .*
2. *If F exhibits row interchangeability with at least three rows including the input set σ and $G(F)$ is a Tinhofer graph, Algorithm 1 detects this structure and returns a corresponding matrix of literals.*

Proof. The first claim is guaranteed by the last part of Algorithm 1 which ensures that transpositions of the rows of the returned matrix M are indeed symmetries of F (Line 11). This implies that arbitrary row permutations are symmetries of F .


Now assume that F exhibits a row symmetry with at least three rows including σ and $G(F)$ is Tinhofer. We argue that the algorithm successfully detects this symmetry. We remark that the orbits of $\text{Aut}(G(F))$ restricted to the literals are precisely orbits of $\text{Aut}(F)$. Let L be the disjoint direct factor of F containing σ and assume that the literals in L can be partitioned into a matrix M that exhibits row symmetry (see Figure 4a). Due to the assumption that $G(F)$ is Tinhofer, if the vertex v corresponding to a literal l of F is individualized, the resulting refined coloring consists of the orbits of $\text{Aut}(G(F))_v$. In particular, due to Lemma 4, the vertices in the row of M are fixed and all other vertices are contained in orbits of size at least two since we have at least three rows (see Figure 4b). Note that since we have at least three rows, $\neg l$ must be in the row of l . Hence after executing the loop for v , $\text{row}[v]$ contains precisely the vertices in the row of v . Isomorphism-invariance of the IR routine (see Lemma 1) ensures that for each row, the order in which symmetrical singletons are colored will be consistent in each row (see Line 7). This ensures that the rows we construct can indeed be transposed (see Line 11 onwards), and the algorithm correctly returns a corresponding matrix. ◀

Recursive Row Symmetry. In practice, orbits often do not *just* exhibit a row symmetry. In particular, we consider the case that an orbit of size k , with a natural symmetric action, is connected to another orbit of size ck , where the symmetric action acts on blocks of size c (see Figure 4c). We extend our algorithm to detect this particular case as follows: in Line 6, we add fragments of other colors instead of vertices in singletons to the row. Let c be a color of π with a fragment c' in π' . We add the vertices $\pi'^{-1}(c')$ to the row, whenever $|\pi'^{-1}(c')| |\sigma| = |\pi^{-1}(c)|$. This means we consider vertices of c' , whenever there is the possibility that the color c is split into $|\sigma|$ parts of size $|\pi'^{-1}(c')|$. We call $\pi'^{-1}(c')$ a *block* of its orbit. On these blocks, we call our algorithm for row symmetry recursively. Essentially, this enables us to detect recursive structures of row symmetry.

Row Symmetry in Stabilizer. A slight extension is that if the test for row symmetry fails, we recurse on the largest fragment from the first IR call and check whether it exhibits row symmetry. This extension is used for the other detection algorithms as well.

3.2 Row-Column Symmetry

Next, we describe a detection algorithm for row-column symmetry. As discussed in Section 2.3, a disjoint direct factor exhibiting row-column symmetry consists of an orbit of literals and its negation, which is also an orbit of literals. We detect row-column symmetry only on one of these orbits, and expand the resulting automorphisms to the other one: For a permutation φ of $\text{Lit}(F)$ and all $l \in \text{Lit}(F)$, let



(a) A row-column symmetry with 4 rows and 5 columns. (b) Individualizing a vertex identifies its row and column.

■ **Figure 5** Illustrations of different aspects of row-column symmetry.

$$\text{expand}_F(\varphi)(l) := \begin{cases} \varphi(l) & \text{if } l \in \text{supp}(\varphi) \\ \neg\varphi(\neg l) & \text{if } \neg l \in \text{supp}(\varphi) \\ l & \text{otherwise.} \end{cases}$$

(Description of Algorithm 2). For an illustration, see Figure 5a. Given a set $\sigma \subseteq \text{Lit}(F)$, we apply IR to a fixed vertex $v \in \sigma$ (see Figure 5b). Assuming that a row-column symmetry is present, this determines a purported “row” $\text{row}[v]$ and “column” $\text{col}[v]$ of v . The algorithm now successively individualizes the vertices in $\text{row}[v]$ and $\text{col}[v]$. This way, every vertex in σ is assigned a reference vertex in each of $\text{row}[v]$ and $\text{col}[v]$, determining its position in the purported matrix. We then verify that the matrix is well-defined and that every row and column transposition, expanded to $\neg\sigma$, is indeed a symmetry of F .

(Correctness of Algorithm 2). In order to prove the correctness of Algorithm 2, we first observe the following:

► **Lemma 6.** *Let $\Gamma = \text{Sym}(n) \times \text{Sym}(m)$ be a row-column symmetry group acting on $[n] \times [m]$. For every $(i, j) \in [n] \times [m]$, the orbit of $(k, l) \in [n] \times [m]$ under the action of $\Gamma_{(i,j)}$ is given by*

$$(k, l)^{\Gamma_{(i,j)}} = \begin{cases} \{(k, l)\} & \text{if } (k, l) = (i, j) \\ \{i\} \times ([m] \setminus \{j\}) & \text{if } k = i, l \neq j \\ ([n] \setminus \{i\}) \times \{j\} & \text{if } k \neq i, l = j \\ ([n] \setminus \{i\}) \times ([m] \setminus \{j\}) & \text{otherwise.} \end{cases}$$

Proof. We identify $[n] \times [m]$ with the entries of an $n \times m$ -matrix M . Then Γ acts on M by permuting the rows and the columns of M . Let $\pi \in \Gamma$ be a permutation that fixes the entry (i, j) . Write $\pi = (\pi_r, \pi_c)$, where π_r is a permutation of the rows and π_c a permutation of the columns of M . Then π_r fixes the i -th row and π_c fixes the j -th column of M . On the other hand, every such element of Γ fixes the entry (i, j) .

Now consider the orbit of $(k, l) \in [n] \times [m]$ under the stabilizer $\Gamma_{(i,j)}$. By definition, it consists of (k, l) for $(k, l) = (i, j)$. For $k = i$ and $l \neq j$, we can map $(k, l) = (i, l)$ to all elements in the i -row, except for (i, j) . Similarly, we argue if $k \neq i$ and $l = j$. Finally, if $k \neq i$ and $l \neq j$, we can map (k, l) to all vertices (k', l') with $k' \neq i$ and $l' \neq j$. This shows the claim. ◀

We prove that the algorithm always returns correct symmetries of F and that in case the model graph is Tinhofer, it is guaranteed to detect row-column symmetry groups.

■ **Algorithm 2** Detection algorithm for row-column symmetry.

```


1 function DetectRowColumnSymmetry
    Input : > formula  $F$ 
           > set  $\sigma \subseteq \text{Lit}(F)$ 
    Output : < candidate matrix  $M$ , or  $\perp$  if check fails
2    $(G, \pi) := G(F), \pi' := \text{IR}((G, \pi), \varepsilon);$ 
3   choose arbitrary  $v \in \sigma;$ 
4    $\pi_v := \text{IR}((G, \pi), v);$ 
5   check that  $\sigma$  has 4 fragments in  $\pi_v;$ 
6   label fragments of  $\sigma$  in  $\pi_v$  not containing  $v$  as  $\sigma_1, \sigma_2, \sigma_3$  in increasing size;
7   // we determine "coordinates" in matrix relative to  $v$ 
8    $\text{row}[v] := \text{col}[v] = v;$  //  $v$  defines a row and a column
9   foreach  $r \in \sigma_1$  do
10  |    $\text{row}[r] := v, \text{col}[r] := r;$  //  $r$  is in row of  $v$ , and defines a column
11  |    $\pi_r := \text{IR}((G, \pi'), r);$ 
12  |   let  $\tau$  be the fragment of  $\sigma$  in  $\pi_r$  of size  $|\sigma_2|$  not containing  $v$  if exists;
13  |   foreach  $t \in \tau$  do  $\text{col}[t] := r;$  //  $t$  is in column of  $r$ 
14  foreach  $c \in \sigma_2$  do
15  |    $\text{col}[c] := v, \text{row}[c] = c;$  //  $c$  is in column of  $v$ , and defines a row
16  |    $\pi_c := \text{IR}((G, \pi'), c);$ 
17  |   let  $\tau$  be the fragment of  $\sigma$  in  $\pi_c$  of size  $|\sigma_1|$  not containing  $v$  if exists;
18  |   foreach  $t \in \tau$  do  $\text{row}[t] := c;$  //  $t$  is in row of  $c$ 
19  construct matrix  $M$  where  $M[r, c] = v'$  with  $\text{row}[v'] = r$  and  $\text{col}[v'] = c;$ 
20  // verify that  $M$  exhibits row-column symmetry
21  check that every vertex in  $\sigma$  has a unique row and a unique column label;
22  check that distinct vertices are assigned distinct label pairs;
23  check that  $M$  has pairwise disjoint rows, and pairwise disjoint columns;
24  foreach  $r \in \sigma_1$  do check that  $\text{expand}_F(\text{transpose}_F(M[*], r), M[*], v))$  is a
    symmetry of  $F;$  //  $M[*], x]$  denotes column of  $x$ 
25  foreach  $c \in \sigma_2$  do check that  $\text{expand}_F(\text{transpose}_F(M[c], *), M[v], *)$  is a
    symmetry of  $F;$  //  $M[x], *]$  denotes row of  $x$ 
26  return  $M$ 

```

► **Theorem 7.** *Let F be a SAT formula.*

1. *If Algorithm 2 returns a matrix M of literals, every permutation of the rows or the columns of M , expanded to the negations of the literals in M , is a symmetry of F .*
2. *If F exhibits a row-column symmetry with at least three rows and at least three columns including σ and $G(F)$ is a Tinhofer graph, then Algorithm 2 detects this structure and returns a corresponding matrix representation of the literals in σ .*

Proof. The first claim is guaranteed by the last part of Algorithm 2 which ensures that transpositions of the rows (Line 24) and columns (Line 25) of the returned matrix M , expanded to the corresponding negated literals, are indeed symmetries of F . By suitably composing such transpositions, we obtain that every permutation of the rows or columns of M induces a symmetry of F in this way.



■ **Figure 6** Individualizing a variable v in a Johnson symmetry (represented by edges in the illustration), splits the set of edges into edges incident to v , and edges not incident to v .

Now assume that $G(F)$ is Tinhofer and that F exhibits row-column symmetry with at least three rows and columns on σ . In other words, the literals in σ can be arranged in a matrix M on which $\text{Aut}(F)$ acts by row and column permutations (see Figure 5a). Individualizing a fixed vertex $v \in \sigma$ causes σ to split into four fragments according to the orbits of the stabilizer $\text{Aut}(F)_v$: the singleton $\{v\}$, two fragments σ_1 and σ_2 corresponding to the remainders of the row and the column of M containing v , and a fragment σ_3 containing the remaining vertices (see Lemma 6 Figure 5b). Since we assume that M has at least three rows and columns, $\sigma_1, \sigma_2, \sigma_3$ are non-singletons and σ_3 is the largest fragment. Without loss of generality, let $\sigma_1 \cup \{v\}$ be the row and $\sigma_2 \cup \{v\}$ be the column of v in M . Every column of M is determined by the unique element of $\sigma_1 \cup \{v\}$ that it contains (similarly for the rows). Individualizing a vertex $r \in \sigma_1$ leads a similar split of σ into four fragments. The fragments corresponding to the row and column of r can be distinguished by observing that v lies in the same row, but not in the same column as r . For all vertices in the column of r , we store this information (Line 13). Similarly, we proceed for the columns (Line 18). After this procedure, every element of σ is assigned a row and column representative in σ_2 and σ_1 respectively, which, up to a permutation of the rows and columns, allows us to recover the matrix M . ◀

3.3 Johnson Symmetry

Finally, we describe a procedure to detect Johnson actions. We remark that there is a classic algorithm to detect Johnson groups [10]. A difference to our heuristic is that we do not know the generators of the group, and instead apply techniques directly on a given graph.

Our aim is to identify the variables in the input set σ with the 2-subsets of $[n]$, where $|\sigma| = \binom{n}{2}$. We thus search for a bijection $b: \sigma \rightarrow \binom{[n]}{2}$ such that $\text{Aut}(F)$ acts as the Johnson group \mathcal{J}_n on σ via this bijection (see Section 2.3, Figure 2, and Figure 3). To avoid confusion, we refer to the elements of $[n]$ as *labels* and to those of $\text{Lit}(F)$ as *literals* or *vertices* of $G(F)$.

(*Description of Algorithm 3.*) Suppose that F exhibits a Johnson symmetry on σ . As described above, there is an (unknown) bijection $b: \sigma \rightarrow \binom{[n]}{2}$ (see Figure 3). We maintain a list $\text{label}[v]$ for every $v \in \sigma$, to which we add $i \in [n]$ when we deduce that $i \in b(v)$. If the algorithm returns a list label, a possible bijection b is given by $b(v) = \text{label}[v]$ for all $v \in \sigma$. Note that b is only determined up to permutation of the labels, so our algorithm merely determines vertices obtaining the same label and assigns the labels consecutively.

The algorithm proceeds as follows: we apply IR to $v \in \sigma$, yielding a coloring π_v . Write $b(v) = \{i, j\}$ for some $i, j \in [n]$. The coloring π_v has three fragments: $\{v\}$, the fragment σ_v containing all $u \in \sigma$ with $|b(u) \cap \{i, j\}| = 1$, and the remaining elements (see Figure 6). We call the vertices in σ_v *adjacent* to v and collect them in $\text{ad}[v]$. Now choose $w \in \text{ad}[v]$. We can assume $b(w) = \{j, k\}$ for some $k \notin \{i, j\}$. As before, we find the vertices adjacent to w by applying IR to w . Individualizing both v and w , the resulting coloring $\pi_{v,w}$ contains exactly one further singleton consisting of $y \in \sigma$ with $b(y) = \{i, k\}$. Now $\text{ad}[v] \cap \text{ad}[w] = \{y\} \cup \{u \in \sigma: b(u) = \{j, r\} \text{ for some } r \notin \{i, j, k\}\}$. The vertices in $\text{ad}[v] \cap \text{ad}[w] \setminus \{y\}$ thus obtain the label j . Similarly, we determine the vertices obtaining the label i or k . After ensuring that the labels have not been considered previously, we add them to the list label for the respective vertices.

■ **Algorithm 3** Detection algorithm for Johnson actions.

```

1 function DetectJohnson
    Input : > formula  $F$ 
           > set  $\sigma \subseteq \text{Lit}(F)$ 
    Output : < bijective labeling of  $\sigma$  by 2-subsets of  $[n]$ , or  $\perp$  if check fails
2  $(G, \pi) := G(F)$ ,  $\pi' := \text{IR}((G, \pi), \varepsilon)$ ;
3 check that  $|\sigma| \geq 28$  and  $|\sigma| = \binom{n}{2}$  for some  $n \in \mathbb{N}$ ;
4 foreach vertex  $v \in \sigma$  do set  $\text{label}[v] = []$ ;
5  $\text{vnr} = 1$ ;
6 while there are vertices  $v \in \sigma$  with  $|\text{label}[v]| \leq 1$  do
7    $E_i := E_j := E_k := \{\}$ ;
8   choose  $v \in \sigma$  with  $|\text{label}[v]| \leq 1$ ;
9    $\pi_v := \text{IR}((G, \pi'), v)$ ;
10  check that number of fragments of  $\sigma$  in  $\pi_v$  is 3;
11  let  $\sigma_v$  be the smaller non-singleton fragment;
12  foreach  $x \in \sigma_v$  do add  $x$  to  $\text{ad}[v]$ ;
13  choose arbitrary  $w \in \text{ad}[v]$ ;
14   $\pi_w := \text{IR}((G, \pi'), w)$ ;
15  check that number of fragments of  $\sigma$  in  $\pi_w$  is 3;
16  let  $\sigma_w$  be the smaller non-singleton fragment;
17  foreach  $x \in \sigma_w$  do add  $x$  to  $\text{ad}[w]$ ;
18   $\pi_{v,w} := \text{IR}((G, \pi_v), w)$ ;
19  let  $\{y\}$  be the unique singleton fragment of  $\sigma$  in  $\pi_{v,w}$  different from  $\{v\}$  and
     $\{w\}$  if existent, otherwise return  $\perp$ ;
20   $\pi_y := \text{IR}((G, \pi'), y)$ ;
21  check that number of fragments of  $\sigma$  in  $\pi_y$  is 3;
22  let  $\sigma_y$  be the smaller non-singleton fragment;
23  foreach  $x \in \sigma_y$  do add  $x$  to  $\text{ad}[y]$ ;
24  add  $v$  to  $E_i$  and  $E_j$ , add  $w$  to  $E_j$  and  $E_k$ , add  $y$  to  $E_i$  and  $E_k$ ;
25  foreach  $x \in \text{ad}[v] \cap \text{ad}[y]$  and  $x \neq w$  do add  $x$  to  $E_i$ ;
26  foreach  $x \in \text{ad}[v] \cap \text{ad}[w]$  and  $x \neq y$  do add  $x$  to  $E_j$ ;
27  foreach  $x \in \text{ad}[w] \cap \text{ad}[y]$  and  $x \neq v$  do add  $x$  to  $E_k$ ;
28  foreach  $E \in \{E_i, E_j, E_k\}$  do
29    if  $\bigcap_{v \in E} \text{label}[v] = \emptyset$  then
30      append  $\text{vnr}$  to  $\text{label}[v]$  for  $v \in E$ ;
31       $\text{vnr} += 1$ ;
32  check that new labels were added to label in this iteration;
33 // verify that  $F$  exhibits Johnson symmetry
34 verify that label induces a bijection between  $\sigma$  and  $\binom{[n]}{2}$ ;
35 foreach  $i \in [n - 1]$  do
36   let  $\beta$  denote the permutation of  $\sigma$  induced by the Johnson action induced by
     $(i, i + 1) \in \text{Sym}(n)$  using label;
37   check that  $\text{expand}_F(\beta)$  is a symmetry of  $F$ ;
38 return label

```

(Description of Algorithm 3). We again make some observations about stabilizers in Johnson groups:

► **Lemma 8.** Let $n \in \mathbb{N}$ and consider the Johnson group $\Gamma := \mathcal{J}_n$, acting on 2-subsets of $[n]$.

1. For $\{i, j\} \in \binom{[n]}{2}$, the orbit of $S \in \binom{[n]}{2}$ under the stabilizer $\Gamma_{\{i, j\}}$ of $\{i, j\} \in \binom{[n]}{2}$ is given by

$$S^{\Gamma_{\{i, j\}}} = \begin{cases} \{S\} & \text{if } S = \{i, j\} \\ \{T \in \binom{[n]}{2} : |T \cap \{i, j\}| = 1\} & \text{if } |S \cap \{i, j\}| = 1 \\ \{T \in \binom{[n]}{2} : T \cap \{i, j\} = \emptyset\} & \text{if } S \cap \{i, j\} = \emptyset. \end{cases}$$

2. For $k \neq i, j$, the orbit of $S \in \binom{[n]}{2}$ under $\Gamma_{\{i, j\}} \cap \Gamma_{\{i, k\}}$ is given by

$$S^{\Gamma_{\{i, j\}} \cap \Gamma_{\{i, k\}}} = \begin{cases} \{S\} & \text{if } S \in \{\{i, j\}, \{i, k\}, \{j, k\}\} \\ \{\{i, r\} : r \in [n] \setminus \{i, j, k\}\} & \text{if } S = \{i, s\} \text{ for some } s \in [n] \setminus \{i, j, k\} \\ \{\{j, r\} : r \in [n] \setminus \{i, j, k\}\} & \text{if } S = \{j, s\} \text{ for some } s \in [n] \setminus \{i, j, k\} \\ \{\{k, r\} : r \in [n] \setminus \{i, j, k\}\} & \text{if } S = \{k, s\} \text{ for some } s \in [n] \setminus \{i, j, k\} \\ \{S \in \binom{[n]}{2} : S \cap \{i, j, k\} = \emptyset\} & \text{if } S \cap \{i, j, k\} = \emptyset. \end{cases}$$

Proof.

1. Let $S \in \binom{[n]}{2}$. If $S = \{i, j\}$, the orbit $S^{\Gamma_{\{i, j\}}}$ consists only of S by definition of the stabilizer. Now suppose that $|S \cap \{i, j\}| = 1$ holds. Without loss of generality, let $S = \{i, r\}$ for some $r \in [n] \setminus \{i, j\}$. Let $\pi \in \Gamma_{\{i, j\}}$. Either π fixes i and j , in which case we have $S^\pi = \{i, r'\}$ for some $r' \in [n] \setminus \{i, j\}$, or π interchanges i and j , in which case we have $S^\pi = \{j, r'\}$ for some $r' \in [n] \setminus \{i, j\}$. In both cases, we have $|S^\pi \cap \{i, j\}| = 1$. On the other hand, it is easy to see that for every set $T \in \binom{[n]}{2}$ with $|T \cap \{i, j\}| = 1$, there exists $\pi \in \Gamma_{\{i, j\}}$ with $S^\pi = T$. The description of $S^{\Gamma_{\{i, j\}}}$ in the case $S \cap \{i, j\} = \emptyset$ can be derived analogously.
2. Note that an element in $\Gamma_{\{i, j\}} \cap \Gamma_{\{i, k\}}$ fixes or interchanges the labels i and j , and at the same time fixes or interchanges the labels i and k . This is only possible if it fixes all of i , j and k . The structure of the orbits then follows similarly to the first claim. ◀

We now prove that the algorithm always returns correct symmetries of F and that in case the model graph is Tinhofer, the algorithm is guaranteed to detect that F exhibits a Johnson symmetry on the input set σ .

► **Theorem 9.** Let F be a SAT formula.

1. If Algorithm 3 returns a list label of labels in $[n]$, then for every element in \mathcal{J}_n , the induced permutation of σ according to label, expanded to $\neg\sigma$, is a symmetry of F .
2. If F exhibits a Johnson symmetry with Johnson group \mathcal{J}_n with $n \geq 8$ on σ and $G(F)$ is a Tinhofer graph, then Algorithm 3 detects this structure and returns a corresponding labeling of the literals in σ by 2-subsets of $[n]$.

Proof. The last part of Algorithm 3 ensures that the Johnson action j_π induced by a transposition $\pi := (i, i+1) \in \mathcal{J}_n$ by permuting the elements in σ according to their labels in label is a symmetry of F when expanded to $\neg\sigma$. By suitably composing these transpositions, it follows that every element of \mathcal{J}_n induces a symmetry of F in this way.

Now suppose that F exhibits a Johnson symmetry with Johnson group \mathcal{J}_n with $n \geq 8$ (i.e., $|\sigma| \geq 28$). Furthermore, assume that $G(F)$ is Tinhofer. In particular, there is a bijection $b: \sigma \rightarrow \binom{[n]}{2}$ (see Figure 3). We claim that when the algorithm terminates, there is

a permutation $\tau \in \text{Sym}(n)$ of the label set $[n]$ such that we have $\text{label}[v] = \{\tau(i), \tau(j)\}$ if $b(v) = \{i, j\}$. Note that the bijection b itself is determined only up to permutation of the labels. Again, for the sake of clarity, we refer to the elements of $[n]$ as *labels* and reserve the term *vertices* for the vertices of the graph $G(F)$.

The individualization of a vertex v with $b(v) = \{i, j\}$ (Line 9) leads to a color partition with three fragments since $G(F)$ is Tinhofer (see Lemma 8 and Figure 6). The smaller non-singleton fragment is $\sigma_v = \{u \in \sigma : |b(u) \cap \{i, j\}| = 1\}$. For this, note that $|\sigma_v| = 2(n-2)$ holds and that we have $n \geq 8$ by assumption. The list $\text{ad}[v]$ (Line 12) then consists of all vertices $u \in \sigma$ with $b(u) = \{i, r\}$ or $b(u) = \{j, r\}$ with $r \in [n] \setminus \{i, j\}$.

Now let $w \in \text{ad}[v]$. Up to this point, the labels i and j are interchangeable, so we may assume $b(w) = \{j, k\}$ for some $k \in [n] \setminus \{i, j\}$. We repeat the above procedure with w in place of v . In particular, $\text{ad}[w]$ (Line 17) contains all vertices $u \in \sigma$ with $b(u) = \{j, r\}$ or $b(u) = \{k, r\}$ for $r \in [n] \setminus \{j, k\}$.

Finally we individualize both v and w to obtain the coloring $\pi_{v,w}$. The fragments are given by Lemma 8. In particular, we obtain $b(y) = \{i, k\}$. Apart from y , the intersection $\text{ad}[v] \cap \text{ad}[w]$ contains all vertices $u \in \sigma$ with $b(u) = \{j, r\}$ for $r \in [n] \setminus \{i, j, k\}$, and we add them to E_j (Line 26). Similarly, we construct the sets E_i and E_k (Lines 25 and 27).

From this explicit description, it is clear that $u \in \sigma$ is added to E_i precisely if $i \in b(u)$ (similarly for E_j and E_k). In particular, for distinct vertices $u_1, u_2 \in E_i$, we have $b(u) \cap b(v) = \{i\}$. Thus if the lists $\text{label}[u]$ for $u \in E_i$ have a common entry, the label i has been considered before (recall that $|E_i| > 1$ holds). Otherwise, we add the current vertex number vnr to $\text{label}[u]$ for all $u \in E_i$ (Line 28) and set $\tau(i) = \text{vnr}$. This way, $\text{label}[u]$ remains duplicate-free and only ever contains labels $\tau(l)$ for $l \in b(u)$. In particular, we always maintain the property $|\text{label}[u]| \leq 2$. In each iteration of the while loop, one of the labels i and j was not considered before (due to $|\text{label}[v]| \leq 1$). In particular, the loop is executed at most n times. When it stops, we have $|\text{label}[v]| = 2$ for all vertices v . ◀

Johnson Action on Row Symmetry. Quite commonly, SAT instances which search for a graph, will search for a graph with a certain *vertex property*. For example, when asking for a k -colorable graph, there will be (interchangeable) colors attached to each vertex of the graph. In order to detect a corresponding symmetry structure, we want to detect blocks which correspond to the labels in the Johnson domain. The detection works by stabilizing vertices in other orbits, and checking whether they split apart the Johnson orbit precisely into the vertices marked with a particular label, and a remainder. If so, these blocks are collected and considered in our overall Johnson action. Finally, we run row symmetry detection on the collected blocks.

4 Implementation

We now give an overview of our new symmetry breaking tool SATSUMA. The input of our algorithm is a CNF formula F . The output is a symmetry breaking constraint for F . We first discuss the breaking constraints produced for a given detected structure.

Breaking Constraints. We produce lex-leader constraints for each detected structure: we use precisely the automorphisms constructed in Algorithm 1, Algorithm 2, and Algorithm 3. Before we can produce lex-leader constraints, we must however fix an ordering on the variables. The ordering used for matrix models simply orders the matrix row-by-row. For Johnson groups, we begin with the vertices of the first label (see Algorithm 3), then the remaining vertices of the second label, and so forth.

family		CMS		BREAKID+CMS			SATSUMA+CMS		
name	size	solved	avg	prep	solved	avg	prep	solved	avg
channel	10	2	484.99	4.727	10	0.032	0.404	10	0.033
cliquecolor	20	2	574.734	0.129	13	228.998	0.058	20	0.845
coloring	55	21	377.338	42.12*	26	317.32	1.071	27	307.632
fpga	10	6	321.596	0.035	10	0.01	0.01	10	0.008
md5	11	5	358.616	0.635	5	359.382	0.548	6	349.171
php	10	3	423.266	6.337	10	0.043	0.128	10	0.036
ramsey	7	2	428.613	1.681	3	343.086	0.394	5	235.27
urquhart	6	6	0.768	0.14	6	0.008	0.032	6	0.066

■ **Figure 7** Benchmarks comparing BREAKID to SATSUMA, using the solver CRYPTOMINISAT (CMS). The timeout is 600 seconds, all times are given in seconds. The columns “prep” denote the average time used for symmetry breaking. Columns “solved” refer to the number of solved instances by CMS, and “avg” is the average time spent by CMS. *BREAKID could not compute the symmetry breaking constraints of two coloring instances within the timeout. We declared these as a timeout for the SAT benchmarks (but the other configurations also timed out on these instances).

High-level Algorithm. The high-level algorithm proceeds as follows:

(Step 1). Construct a model graph from the given CNF formula.

(Step 2). Run the algorithms described in the previous section in the following order: Johnson groups (Algorithm 3), row-column symmetry (Algorithm 2), row interchangeability (Algorithm 1). Whenever a structure is found, all orbits covered by the structure are marked. The subsequent analysis only considers *unmarked* orbits. For each structure, symmetry breaking constraints are constructed as described above. Lastly, we maintain a vertex coloring of the model graph, which we call the *remainder coloring*: this coloring restricts the symmetries of the model graph to symmetries not yet covered by detected structures.

(Step 3). Run symmetry detection on the graph colored with the remainder coloring. Then, the *binary clause* heuristic of BREAKID is applied for all variables not yet ordered by already produced lex-leader constraints: a stabilizer chain of the automorphism group is approximated, and for each stabilized variable x a short lex-leader constraint for each other literal y of its orbit is produced, i.e., essentially the binary constraint $x \leq y$ (see [21] for a detailed description). Lastly, a lex-leader constraint for each generator is produced.

Implementation. The tool is written in C++, and is freely available as open source software [2]. The tool DEJAVU [3, 4, 6] is used for providing general-purpose symmetry detection, the individualization-refinement framework, and data structures for symmetries. Significant parts of the implementation, in particular the generation of lex-leader constraints and binary clauses, are reverse-engineered from BREAKID. Our reimplementations of these routines differs in two crucial aspects from the original one: first, BREAKID uses the symmetry detection tool SAUCY [18] instead of DEJAVU. Second, we use different data structures and algorithms for the handling of symmetry.

5 Benchmarks

We compare the state-of-the-art static symmetry breaking tool BREAKID (version 2.6) to SATSUMA.

As SAT solvers, we use CRYPTOMINISAT [39] and CADICAL [12]. The benchmarks using CADICAL largely concur with the CRYPTOMINISAT benchmarks, and our descriptions will focus on the results using CRYPTOMINISAT. The timeout for all benchmarks is 600

family		CAD		BREAKID+CAD			SATSUMA+CAD		
name	size	solved	avg	prep	solved	avg	prep	solved	avg
channel	10	2	494.226	4.727	10	0.077	0.404	10	0.077
cliquecolor	20	9	442.373	0.129	13	216.999	0.058	20	0.2
coloring	55	20	393.864	42.12*	26	316.779	1.071	28	301.783
fpga	10	5	391.29	0.035	10	0.008	0.01	10	0.025
md5	11	6	339.378	0.635	6	343.324	0.548	6	324.716
php	10	3	422.976	6.337	10	0.085	0.128	10	0.1
ramsey	7	2	428.583	1.681	3	342.908	0.394	5	192.299
urquhart	6	2	449.622	0.14	6	0.005	0.032	6	0.052

■ **Figure 8** Benchmarks comparing BREAKID to SATSUMA. The SAT solver used is CADICAL (CAD). The timeout used is 600 seconds. The columns “prep” refer to the time in seconds used to compute the symmetry breaking constraint. Columns “solved” refer to the number of solved instances by CAD, and “avg” is the average time used by CAD (*excluding* the time used for symmetry breaking). *BREAKID could not compute the symmetry breaking constraints of two coloring instances within the timeout.


seconds. We separately measure the time spent on symmetry breaking itself, and SAT solving. All benchmarks ran sequentially on an Intel Core i7 9700K with 64GB of RAM on Ubuntu 20.04.

Benchmark Instances. We run benchmarks on a variety of well-established instance families exhibiting symmetry (see Figure 7). The sets coloring, urquhart, fpga, md5, and channel are part of the distribution of BREAKID [21]. We generate pigeonhole principle (php) instances, Ramsey instances, and clique coloring instances using the tool CNFGEN [31]. The set of parameters for clique coloring is similar to [26], but we added larger instances. All instances are unsatisfiable.

Regarding the detected symmetry structures of these instances, we detect Johnson symmetry on the ramsey and cliquecolor families. On php, channel, and fpga, SATSUMA detects row-column symmetry, and BREAKID corresponding row interchangeability (see also [36, 21]). The coloring instances exhibit a variety of different symmetries, but in particular also row symmetry [21]. In urquhart and md5, no structure is detected by either of the tools.

Regarding our choice of benchmark instances, we stress that our main goal is to observe whether detecting richer structures can improve performance compared to existing approaches.

SAT Benchmarks. An overview of the results can be found in Figure 7 (for CADICAL, see Figure 8). Considering the results, we observe that SATSUMA solves more instances, and solving times are considerably lower on average on the cliquecolor and ramsey instances. We recall that these instance families exhibit Johnson symmetry. On all sets with row and row-column symmetry, that is channel, coloring, fpga, and php, we observe that solved instances and average solving times are comparable. On coloring, we observe that SATSUMA solves one more instance than BREAKID (and two more using CADICAL). For urquhart, both SATSUMA and BREAKID rely on the binary clause strategy. The results indicate that BREAKID is more effective in breaking symmetry, which is however outweighed by the faster runtime of SATSUMA. The md5 instances only contain a single non-trivial symmetry. Here, SATSUMA produces more breaking clauses, and we observe a consistent albeit marginal speedup. It should be mentioned that it does however seem plausible that the observed speed-up may be due to shuffling of literals in clauses, or other factors.



■ **Figure 9** Benchmarks comparing the computational overhead of BREAKID to SATSUMA. The shown computation time is the time spent computing symmetry breaking constraints for an instance using the respective tool. The red bar indicates the timeout of 600 seconds.

In particular, we point out that SATSUMA compares favorably on instance families which exhibit Johnson symmetry. We believe this to be due to our detection of Johnson symmetry and the subsequent generation of more favorable constraints. Crucially, on all successfully solved instances of cliquecolor and ramsey, the *remainder contains no symmetry*: all symmetries are detected and in turn broken solely using the algorithms of this paper, and no general-purpose symmetry detection and breaking is applied.

We observe that the average time spent computing the symmetry breaking constraints is lower on all families for SATSUMA. A more in-depth analysis follows below.

Computational Overhead. We conduct further benchmarks to gauge the computational overhead incurred by BREAKID and SATSUMA. We test three different benchmark families: php, cliquecolor, and urquhart (generated using CNFGEN). For php, we increase the number of pigeons from 10 to 150 (with $n - 1$ holes, respectively). For cliquecolor, we increase the number of vertices of the prospective graph from 10 to 300 (the size of the clique is 3 and number of colors 2). In urquhart, we use random 5-regular graphs, increasing the number of vertices from 10 to 350. We chose these instance families such that they cover the different symmetry detection routines in SATSUMA: the family php essentially measures the runtime of our row-column routine, cliquecolor that of the Johnson routine, and urquhart uses general purpose symmetry detection, followed by the binary clause strategy.

Figure 9 summarizes the results. In all instance families, the data suggest that SATSUMA asymptotically scales better than BREAKID. These results match our observations regarding overhead from the first part of the benchmarks (see Figure 7).

We believe there are multiple reasons why SATSUMA runs faster than BREAKID. First, our new algorithms of Section 3 verify symmetries on the CNF formula instead of the model graph. This is advantageous because symmetries of the CNF only explicitly map literals, whereas symmetries of the model graph also explicitly map clauses. Second, most routines in our implementation run proportional in the size of the *support* of symmetries, as opposed to the number of literals of F . Third, for general-purpose symmetry detection, DEJAVU seems to be more efficient in computing automorphism groups of SAT instances than SAUCY [3].

We mention that in the urquhart instances, the outliers with high running time seem to be due to SAUCY taking a long time to compute symmetries for BREAKID. On the other hand, in these cases, we observe that the symmetries as returned by DEJAVU are less suitable for the binary clause heuristic, leading to fewer produced clauses. This could however be easily alleviated by a strengthening of the heuristic (e.g., by sometimes applying the Schreier-Sims algorithm for stabilizers as already pointed out in [21]).

6 Conclusions and Future Work

We described a new structure-based approach to symmetry breaking, and demonstrated the effectiveness of our implementation SATSUMA. There seem to be many promising directions in which the present work could be expanded:

- Detect more group structures: in particular, a more generic approach to detect aggregates of groups would be of great interest. Another interesting case might be the symmetries of the family `urquhart`, which are isomorphic to C_2^k and have been studied previously [32].
- Consider other breaking approaches for certain group structures. So far, we used the knowledge of group structures to pick out automorphisms, for which off-the-shelf lex-leader constraints are generated. Since optimal handling of row-column symmetry and Johnson symmetry seems infeasible with lex-leader constraints [32], other breaking constraints could lead to better results. Moreover, Johnson symmetry allows the use of symmetry reduction developed specifically for graph generation [15, 16, 30].
- Improved techniques for handling of the “remainder”. As already pointed out in [21], one potential direction would be to apply the random Schreier-Sims algorithm [38] to produce more small symmetry breaking clauses.
- An enticing feature is proof-logging, as was recently introduced to BREAKID [13].
- The new detection algorithms could be applied in other domains as well: for example, seeing as row interchangeability is successfully used in MIP, it seems only natural that MIP instances may also contain richer structures.
- Sometimes symmetries are not present in a compiled CNF of a given problem (as, e.g., analyzed in [26]). A possible remedy is to allow the user to provide an auxiliary graph that models the original symmetry (see [26]), and the methods proposed in this paper should generalize to this setting.

References

- 1 Fadi A. Aloul, Igor L. Markov, and Kareem A. Sakallah. Shatter: efficient symmetry-breaking for boolean satisfiability. In *Proceedings of the 40th Design Automation Conference, DAC 2003, Anaheim, CA, USA, June 2-6, 2003*, pages 836–839. ACM, 2003. doi:10.1145/775832.776042.
- 2 Markus Anders, Sofia Brenner, and Gaurav Rattan. `satsuma`. Software, version 1.0., swHId: `swh:1:dir:134ea7952a25d92ba2addf7ff25aa18183550735` (visited on 2024-08-05). URL: <https://github.com/markusa4/satsuma>.
- 3 Markus Anders and Pascal Schweitzer. `dejavu`. <https://automorphisms.org>.
- 4 Markus Anders and Pascal Schweitzer. Parallel computation of combinatorial symmetries. In *29th Annual European Symposium on Algorithms, ESA 2021, September 6-8, 2021, Lisbon, Portugal (Virtual Conference)*, volume 204 of *LIPICs*, pages 6:1–6:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.ESA.2021.6.
- 5 Markus Anders, Pascal Schweitzer, and Mate Soos. Algorithms transcending the SAT-symmetry interface. In *26th International Conference on Theory and Applications of Satisfiability Testing, SAT 2023, July 4-8, 2023, Alghero, Italy*, volume 271 of *LIPICs*, pages 1:1–1:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICs.SAT.2023.1.
- 6 Markus Anders, Pascal Schweitzer, and Julian Stieß. Engineering a preprocessor for symmetry detection. In *21st International Symposium on Experimental Algorithms, SEA 2023, July 24-26, 2023, Barcelona, Spain*, volume 265 of *LIPICs*, pages 1:1–1:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICs.SEA.2023.1.
- 7 Vikraman Arvind, Johannes Köbler, Gaurav Rattan, and Oleg Verbitsky. Graph isomorphism, color refinement, and compactness. *Comput. Complex.*, 26(3):627–685, 2017. doi:10.1007/s00037-016-0147-6.

- 8 Gilles Audemard, Saïd Jabbour, and Lakhdar Sais. Symmetry breaking in quantified boolean formulae. In *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, pages 2262–2267, 2007. URL: <http://ijcai.org/Proceedings/07/Papers/364.pdf>.
- 9 László Babai. Graph isomorphism in quasipolynomial time [extended abstract]. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 684–697. ACM, 2016. doi:10.1145/2897518.2897542.
- 10 László Babai, Eugene M. Luks, and Ákos Seress. Permutation groups in NC. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 409–420. ACM, 1987. doi:10.1145/28395.28439.
- 11 Christoph Berkholz, Paul S. Bonsma, and Martin Grohe. Tight lower and upper bounds for the complexity of canonical colour refinement. *Theory Comput. Syst.*, 60(4):581–614, 2017. doi:10.1007/s00224-016-9686-0.
- 12 Armin Biere, Katalin Fazekas, Mathias Fleury, and Maximillian Heisinger. CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling entering the SAT Competition 2020. In *Proc. of SAT Competition 2020 – Solver and Benchmark Descriptions*, volume B-2020-1 of *Department of Computer Science Report Series B*, pages 51–53. University of Helsinki, 2020.
- 13 Bart Bogaerts, Stephan Gocht, Ciaran McCreesh, and Jakob Nordström. Certified dominance and symmetry breaking for combinatorial optimisation. *J. Artif. Intell. Res.*, 77:1539–1589, 2023. doi:10.1613/jair.1.14296.
- 14 Bart Bogaerts, Jakob Nordström, Andy Oertel, and Çağrı Uluç Yıldırımoglu. BreakID-kissat in SAT competition 2023 (system description). In *Proceedings of SAT Competition 2023: Solver, Benchmark and Proof Checker Descriptions*, Department of Computer Science Series of Publications B, Finland, 2023. Department of Computer Science, University of Helsinki.
- 15 Michael Codish, Graeme Gange, Avraham Itzhakov, and Peter J. Stuckey. Breaking symmetries in graphs: The nauty way. In *Principles and Practice of Constraint Programming – 22nd International Conference, CP 2016, Toulouse, France, September 5-9, 2016, Proceedings*, volume 9892 of *Lecture Notes in Computer Science*, pages 157–172. Springer, 2016. doi:10.1007/978-3-319-44953-1_11.
- 16 Michael Codish, Alice Miller, Patrick Prosser, and Peter J. Stuckey. Constraints for symmetry breaking in graph representation. *Constraints An Int. J.*, 24(1):1–24, 2019. doi:10.1007/s10601-018-9294-5.
- 17 James M. Crawford, Matthew L. Ginsberg, Eugene M. Luks, and Amitabha Roy. 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*, pages 148–159. Morgan Kaufmann, 1996.
- 18 Paul T. Darga, Mark H. Liffiton, Karem A. Sakallah, and Igor L. Markov. Exploiting structure in symmetry detection for CNF. In *Proceedings of the 41th Design Automation Conference, DAC 2004, San Diego, CA, USA, June 7-11, 2004*, pages 530–534. ACM, 2004. doi:10.1145/996566.996712.
- 19 Jo Devriendt and Bart Bogaerts. BreakID: Static symmetry breaking for ASP (system description). *CoRR*, abs/1608.08447, 2016. doi:10.48550/arXiv.1608.08447.
- 20 Jo Devriendt, Bart Bogaerts, and Maurice Bruynooghe. Symmetric explanation learning: Effective dynamic symmetry handling for SAT. In *Theory and Applications of Satisfiability Testing – SAT 2017 – 20th International Conference, Melbourne, VIC, Australia, August 28 – September 1, 2017, Proceedings*, volume 10491 of *Lecture Notes in Computer Science*, pages 83–100. Springer, 2017. doi:10.1007/978-3-319-66263-3_6.
- 21 Jo Devriendt, Bart Bogaerts, Maurice Bruynooghe, and Marc Denecker. Improved static symmetry breaking for SAT. In *Theory and Applications of Satisfiability Testing – SAT 2016 – 19th International Conference, Bordeaux, France, July 5–8, 2016, Proceedings*, volume 9710 of *Lecture Notes in Computer Science*, pages 104–122. Springer, 2016. doi:10.1007/978-3-319-40970-2_8.

- 22 Jo Devriendt, Bart Bogaerts, Broes De Cat, Marc Denecker, and Christopher Mears. Symmetry propagation: Improved dynamic symmetry breaking in SAT. In *IEEE 24th International Conference on Tools with Artificial Intelligence, ICTAI 2012, Athens, Greece, November 7-9, 2012*, pages 49–56. IEEE Computer Society, 2012. doi:10.1109/ICTAI.2012.16.
- 23 Pierre Flener, Alan M. Frisch, Brahim Hnich, Zeynep Kiziltan, Ian Miguel, Justin Pearson, and Toby Walsh. Breaking row and column symmetries in matrix models. In *Principles and Practice of Constraint Programming – CP 2002, 8th International Conference, CP 2002, Ithaca, NY, USA, September 9-13, 2002, Proceedings*, volume 2470 of *Lecture Notes in Computer Science*, pages 462–476. Springer, 2002. doi:10.1007/3-540-46135-3_31.
- 24 Pierre Flener, Alan M. Frisch, Brahim Hnich, Zeynep Kiziltan, Ian Miguel, and Toby Walsh. Matrix modelling. Technical Report APES-36-2001, APES group (2001), 2001.
- 25 Ian P. Gent, Karen E. Petrie, and Jean-François Puget. Symmetry in constraint programming. In *Handbook of Constraint Programming*, volume 2 of *Foundations of Artificial Intelligence*, pages 329–376. Elsevier, 2006. doi:10.1016/S1574-6526(06)80014-3.
- 26 Tommi A. Junttila, Matti Karppa, Petteri Kaski, and Jukka Kohonen. An adaptive prefix-assignment technique for symmetry reduction. *J. Symb. Comput.*, 99:21–49, 2020. doi:10.1016/j.jsc.2019.03.002.
- 27 Tommi A. Junttila and Petteri Kaski. Conflict propagation and component recursion for canonical labeling. In *Theory and Practice of Algorithms in (Computer) Systems – First International ICST Conference, TAPAS 2011, Rome, Italy, April 18-20, 2011. Proceedings*, volume 6595 of *Lecture Notes in Computer Science*, pages 151–162. Springer, 2011. doi:10.1007/978-3-642-19754-3_16.
- 28 George Katsirelos, Nina Narodytska, and Toby Walsh. On the complexity and completeness of static constraints for breaking row and column symmetry. In *Principles and Practice of Constraint Programming – CP 2010 – 16th International Conference, CP 2010, St. Andrews, Scotland, UK, September 6-10, 2010. Proceedings*, volume 6308 of *Lecture Notes in Computer Science*, pages 305–320. Springer, 2010. doi:10.1007/978-3-642-15396-9_26.
- 29 Markus Kirchweger, Manfred Scheucher, and Stefan Szeider. A SAT attack on rota’s basis conjecture. In *25th International Conference on Theory and Applications of Satisfiability Testing, SAT 2022, August 2-5, 2022, Haifa, Israel*, volume 236 of *LIPICs*, pages 4:1–4:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.SAT.2022.4.
- 30 Markus Kirchweger and Stefan Szeider. SAT modulo symmetries for graph generation. In *27th International Conference on Principles and Practice of Constraint Programming, CP*, volume 210 of *LIPICs*, pages 34:1–34:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.CP.2021.34.
- 31 Massimo Lauria, Jan Elffers, Jakob Nordström, and Marc Vinyals. Cnfgcn: A generator of crafted benchmarks. In *Theory and Applications of Satisfiability Testing – SAT 2017 – 20th International Conference, Melbourne, VIC, Australia, August 28 – September 1, 2017, Proceedings*, volume 10491 of *Lecture Notes in Computer Science*, pages 464–473. Springer, 2017. doi:10.1007/978-3-319-66263-3_30.
- 32 Eugene M. Luks and Amitabha Roy. The complexity of symmetry-breaking formulas. *Ann. Math. Artif. Intell.*, 41(1):19–45, 2004. doi:10.1023/B:AMAI.0000018578.92398.10.
- 33 Brendan D. McKay and Adolfo Piperno. Practical graph isomorphism, II. *J. Symb. Comput.*, 60:94–112, 2014. doi:10.1016/j.jsc.2013.09.003.
- 34 Hakan Metin, Souheib Baarir, Maximilien Colange, and Fabrice Kordon. Cdclsym: Introducing effective symmetry breaking in SAT solving. In *Tools and Algorithms for the Construction and Analysis of Systems – 24th International Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings, Part I*, volume 10805 of *Lecture Notes in Computer Science*, pages 99–114. Springer, 2018. doi:10.1007/978-3-319-89960-2_6.

- 35 Marc E. Pfetsch and Thomas Rehn. A computational comparison of symmetry handling methods for mixed integer programs. *Math. Program. Comput.*, 11(1):37–93, 2019. doi:10.1007/s12532-018-0140-y.
- 36 Ashish Sabharwal. Symchaff: exploiting symmetry in a structure-aware satisfiability solver. *Constraints An Int. J.*, 14(4):478–505, 2009. doi:10.1007/s10601-008-9060-1.
- 37 Kareem A. Sakallah. Symmetry and satisfiability. In *Handbook of Satisfiability – Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, pages 509–570. IOS Press, 2021. doi:10.3233/FAIA200996.
- 38 Ákos Seress. *Permutation Group Algorithms*. Cambridge Tracts in Mathematics. Cambridge University Press, 2003. doi:10.1017/CB09780511546549.
- 39 Mate Soos, Karsten Nohl, and Claude Castelluccia. Extending SAT solvers to cryptographic problems. In *Theory and Applications of Satisfiability Testing – SAT 2009, 12th International Conference, SAT 2009, Swansea, UK, June 30 – July 3, 2009. Proceedings*, volume 5584 of *Lecture Notes in Computer Science*, pages 244–257. Springer, 2009. doi:10.1007/978-3-642-02777-2_24.
- 40 Gottfried Tinhofer. A note on compact graphs. *Discret. Appl. Math.*, 30(2-3):253–264, 1991. doi:10.1016/0166-218X(91)90049-3.