

# Algorithms Transcending the SAT-Symmetry Interface

Markus Anders

TU Darmstadt, Germany

Pascal Schweitzer

TU Darmstadt, Germany

Mate Soos

National University of Singapore, Singapore

---

## Abstract

Dedicated treatment of symmetries in satisfiability problems (SAT) is indispensable for solving various classes of instances arising in practice. However, the exploitation of symmetries usually takes a black box approach. Typically, off-the-shelf external, general-purpose symmetry detection tools are invoked to compute symmetry groups of a formula. The groups thus generated are a set of permutations passed to a separate tool to perform further analyzes to understand the structure of the groups. The result of this second computation is in turn used for tasks such as static symmetry breaking or dynamic pruning of the search space. Within this pipeline of tools, the detection and analysis of symmetries typically incurs the majority of the time overhead for symmetry exploitation.

In this paper we advocate for a more holistic view of what we call the *SAT-symmetry interface*. We formulate a computational setting, centered around a new concept of joint graph/group pairs, to analyze and improve the detection and analysis of symmetries. Using our methods, no information is lost performing computational tasks lying on the SAT-symmetry interface. Having access to the entire input allows for simpler, yet efficient algorithms.

Specifically, we devise algorithms and heuristics for computing finest direct disjoint decompositions, finding equivalent orbits, and finding natural symmetric group actions. Our algorithms run in what we call instance-quasi-linear time, i.e., almost linear time in terms of the input size of the original formula and the description length of the symmetry group returned by symmetry detection tools. Our algorithms improve over both heuristics used in state-of-the-art symmetry exploitation tools, as well as theoretical general-purpose algorithms.

**2012 ACM Subject Classification** Theory of computation → Design and analysis of algorithms

**Keywords and phrases** boolean satisfiability, symmetry exploitation, computational group theory

**Digital Object Identifier** 10.4230/LIPIcs.SAT.2023.1

**Related Version** *Full Version*: <https://arxiv.org/abs/2306.00613> [5]

**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).

## 1 Introduction

Many SAT instances, especially of the hard combinatorial type, exhibit symmetries. When symmetries exhibited by these instances are not handled adequately, SAT solvers may repeatedly explore symmetric parts of the search space. This can dramatically increase runtime, sometimes making it impossible for the solver to finish within reasonable time [7].

One common method to handle the symmetries is to add symmetry breaking formulas to the problem specification [10, 2]. This approach is called static symmetry breaking. Another, competing, approach is to handle symmetries dynamically during the running of the SAT solver. There are a variety of such dynamic strategies, exploiting symmetry information



© Markus Anders, Pascal Schweitzer, and Mate Soos;  
licensed under Creative Commons License CC-BY 4.0

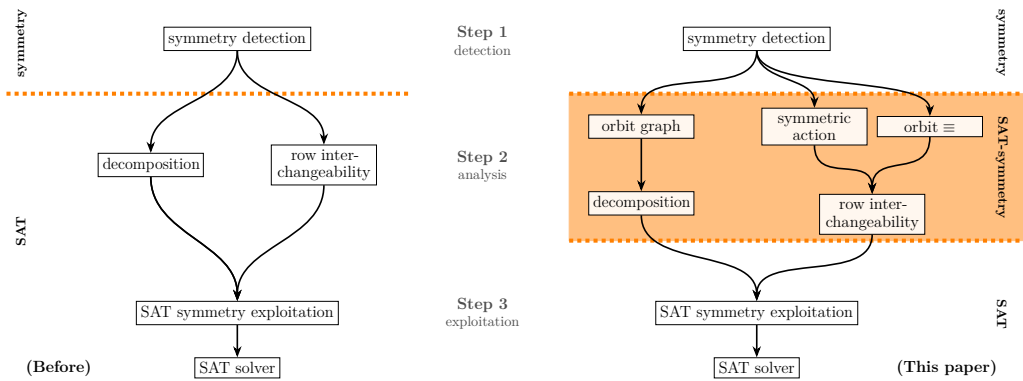
26th International Conference on Theory and Applications of Satisfiability Testing (SAT 2023).

Editors: Meena Mahajan and Friedrich Slivovsky; Article No. 1; pp. 1:1–1:21

Leibniz International Proceedings in Informatics



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



■ **Figure 1** Blurring the lines of the SAT-symmetry interface. We analyze existing practical routines (left), draw connections to existing concepts in computational group theory, and describe improved algorithms in our new SAT-symmetry context (right).

during variable branching [21] and learning [28, 12]. For SAT, the tools SHATTER [1] and BREAKID [13, 8] take the static symmetry breaking approach, while SYMCHAFF [28] and SMS [21] take the dynamic symmetry exploitation approach.

While there is a growing number of competing approaches of how best to handle symmetries, there are also a number of common obstacles: symmetries of the underlying formula have to be detected first, and the structure of symmetries has to be understood, at least to some degree. Approaches that handle symmetries can be typically divided into three distinct steps: (Step 1) symmetry detection, (Step 2) symmetry analysis, and (Step 3) symmetry breaking, or other ways of exploiting symmetry. In the following, we discuss each of these steps, also illustrated on the left side of Figure 1.

*Step 1.* In practice, symmetries are detected by modeling a given SAT formula as a graph, and then applying an off-the-shelf symmetry detection tool, such as SAUCY [11], to the resulting graph. Since symmetries form a permutation group under composition, a symmetry detection tool does not return all the symmetries. Instead, it only returns a small set of *generators*, which, by composition, lead to all the symmetries of the formula. Indeed, returning only a small set of generators is crucial for efficiency, since the number of symmetries is often exponential in the size of the formula.

*Step 2.* Symmetry exploitation algorithms apply heuristics to analyze the structure of the group described by the generators. This is necessary to enable the best possible use of the symmetries to improve SAT solver performance. We mention three examples for structural analyzes. Firstly, the *disjoint direct decomposition* splits a group into independent parts that can be handled separately. Secondly, so-called *row interchangeability* subgroups of the group [13, 14, 25] are of particular interest since they form a class of groups for which linear-sized, complete symmetry breaking constraints are known. Thirdly, *stabilizers* are commonly used for various purposes among both static and dynamic approaches [27].

*Step 3.* Lastly, the symmetries and structural insights are used to reduce the search space in SAT using one of the various static and dynamic symmetry exploitation approaches.

Designing symmetry exploitation algorithms typically involves delicately balancing computational overhead versus how thoroughly symmetries are used. In this trade-off, symmetry detection (Step 1) and analysis (Step 2) typically induce the majority of the overhead [13]. The main focus of this paper is improving the analysis of symmetries, i.e. (Step 2).

Practical implementations in use today that perform such structural analyzes do so through heuristics. While using heuristics is not an issue per se, some heuristics currently in use strongly depend on properties that the generators returned by symmetry detection tools may or may not exhibit. For example, BREAKID and the MIP heuristic in [25] both rely on so-called *separability* of the generating set and a specific arrangement of *transpositions* being present. Neither of these properties are guaranteed by contemporary symmetry detection tools [9].

In fact, modern symmetry detection tools such as TRACES [23] and DEJAVU [3] return randomly selected symmetries, since the use of randomization provides an asymptotic advantage in the symmetry detection process itself [4]. However, generating sets consisting of randomly selected symmetries are in a sense the exact opposite of what is desired for the heuristics, since with high probability random symmetries satisfy neither of the required conditions. This is particularly unfortunate, as DEJAVU is currently the fastest symmetry detection tool available for graphs stemming from SAT instances [6].


Another downside of the use of practical heuristics for the structural analysis of the group is that they are often also computationally expensive and make up a large portion of the runtime of the overall symmetry exploitation process. For example, the row interchangeability algorithm of BREAKID performs multiple callbacks to the underlying symmetry detection tool, where each call can be expensive.

Altogether, heuristics in use today sometimes cause significant overhead, while also posing an obstacle to speeding up symmetry detection itself. This immediately poses the question: why is it that these heuristics are currently in place that cause such a loss of efficiency when it comes to computations within the SAT-symmetry interface?

We believe that the issue is that tools on either side of the interface treat each other as a black box. Indeed, when considered as an isolated task, algorithms for the analysis of permutation groups are well-researched in the area of computational group theory [30]. Not only is the theory well-understood, but there are also highly efficient implementations [15]. However, we can make two crucial observations regarding the available algorithms. First and foremost, for group theoretic algorithms from the literature that are deemed to have linear or nearly-linear runtime [30], the concrete runtime notions actually differ from the ones applicable in the overall context. In fact, the runtime is essentially measured in terms of a dense rather than a sparse input description. Therefore, in the context of SAT-solving or graph algorithms, the runtime of these algorithms should rather be considered quadratic. Secondly, in computational group theory, algorithms assume that only generators for an input group are available. However, in the context of the SAT-symmetry interface, not only a group but also a graph (computed from the original formula) is available. It turns out as a key insight of our paper that lacking access to the graphs crucially limits the design space for efficient algorithms.

**Contributions.** Advocating a holistic view of the SAT-symmetry interface, we develop algorithms that transcend both into the SAT domain and the symmetry domain at the same time. This is illustrated in Figure 1 on the right side.

Firstly, we provide a definition for the computational setting such as input, output, and runtime, under which these algorithms should operate (Section 3). We then extract precise formal problem definitions from heuristics implemented in state-of-the-art tools (Section 4). Lastly, we demonstrate the efficacy of our new approach by providing faster theoretical algorithms for commonly used heuristics, as is described below.



■ **Figure 2** An illustration of a color refinement process.

**Computational Setting.** In our new computational setting, algorithms take as input a *joint graph/group pair*, meaning a group  $S$  and corresponding graph  $G$ , whose symmetry group is precisely  $\langle S \rangle$ . We define a precise notion of *instance-linear time*, meaning it is linear in the encoding size of the SAT formula, graph, and group.

**New Algorithms.** Given a joint graph/group pair, we develop and analyze the following algorithms:

- A1** An instance-linear algorithm for computing the *finest direct disjoint decomposition* of the symmetry group of a graph (Section 5). We also give a heuristic specific to SAT formulas, decomposing the symmetry group on the literals.
- A2** An algorithm to simultaneously detect *natural symmetric group actions* on all the orbits of a group (Section 6). Here we exploit randomized techniques from computational group theory for the detection of “giant” permutation groups. We give instance-linear heuristics which are able to exploit properties of the SAT-symmetry interface.
- A3** An instance-quasi-linear algorithm to compute *equivalent symmetric orbits*, under some mild assumptions about the generating set (Section 7). In conjunction with (A2), this enables us to detect *all* elementary row interchangeability subgroups.

Both (A1) and (A3) improve the (at least) quadratic runtime of previous, general-purpose permutation group algorithms of [9] and [30], respectively.

## 2 Preliminaries and Related Work

**Graphs and Symmetries.** A colored graph  $G = (V, E, \pi)$  consists of a set of vertices  $V$ , edges  $E \subseteq V \times V$ , and a vertex coloring  $\pi: V \rightarrow C$  which maps  $V$  to some set of colors  $C$ . We use  $V(G)$ ,  $E(G)$ , and  $\pi(G)$  to refer to the vertices, edges, and coloring of  $G$ , respectively.

A symmetry, or *automorphism*, of a colored graph  $G = (V, E, \pi)$  is a bijection  $\varphi: V \rightarrow V$  such that  $\varphi(E) = E$  as well as  $\pi(v) = \pi(\varphi(v))$  for all  $v \in V$ . In other words, symmetries preserve the neighborhood relation of the graph, as well as the coloring of vertices. The colors of vertices in the graph are solely used to ensure that distinctly colored vertices are not mapped onto each other using symmetries. Together, all symmetries of a graph form a permutation group under composition, which we call the *automorphism group*  $\text{Aut}(G)$ .

In this paper, we call software tools computing the automorphism group of a graph *symmetry detection tools* [23, 11, 18, 23, 3]. In the literature, these tools are also often called *practical graph isomorphism solvers*. In this paper, we avoid the use of this term in order not to confuse them with SAT *solvers*.

**Color refinement.** A common algorithm applied when computing the symmetries of a graph is *color refinement*. Given a colored graph  $G = (V, E, \pi)$ , color refinement *refines* the coloring  $\pi$  of  $G$  into  $G' = (V, E, \pi')$ . Crucially, the automorphism group remains invariant under color refinement, i.e.,  $\text{Aut}(G) = \text{Aut}(G')$ .

We now describe the algorithm. If two vertices in some color  $X = \pi^{-1}(c)$  have a different number of neighbors in another color  $Y = \pi^{-1}(c')$ , then  $X$  can be split by partitioning it according to neighbor counts in  $Y$ . After the split, two vertices have the same color precisely if they had the same color before the split, and they have the same number of neighbors in  $Y$ .

We repeatedly split classes with respect to other classes until no further splits are possible. Figure 2 shows an illustration of the color refinement procedure. A coloring which does not admit further splits is called *equitable*. For a graph  $G$ , color refinement can be computed in time  $\mathcal{O}(|E(G)| \log |V(G)|)$  [22, 26].

Let us also recall a different definition for equitable colorings: A coloring  $\pi$  of a graph is equitable if for all pairs of (not necessarily distinct) color classes  $C_1, C_2$ , all vertices in  $C_1$  have the same number of neighbors in  $C_2$  (i.e.,  $|N(v) \cap C_2| = |N(v') \cap C_2|$  for all  $v, v' \in C_1$ ). Given a coloring  $\pi$ , color refinement computes an equitable refinement  $\pi'$ , i.e., an equitable coloring  $\pi'$  for which  $\pi'(v) = \pi'(v')$  implies  $\pi(v) = \pi(v')$ . In fact, it computes the *coarsest* equitable refinement.

**Permutation Groups.** The *symmetric group*  $\text{Sym}(\Omega)$  is the permutation group consisting of all permutations of the set  $\Omega$ . A *permutation group* on domain  $\Omega$  is a group  $\Gamma$  that is a subgroup of  $\text{Sym}(\Omega)$ , denoted  $\Gamma \leq \text{Sym}(\Omega)$ . For a subset of the domain  $\Omega' \subseteq \Omega$ , the *restriction* of  $\Gamma$  to  $\Omega'$  is  $\Gamma|_{\Omega'} := \{\varphi|_{\Omega'} \mid \varphi \in \Gamma\}$  (where  $\varphi|_{\Omega}$  denotes restricting the domain of  $\varphi$  to  $\Omega$ ). The restriction is not necessarily a group since the images need not be in  $\Omega'$ . The *pointwise stabilizer* is the group  $\Gamma_{(\Omega')} := \{\varphi \in \Gamma \mid \forall p \in \Omega' : \varphi(p) = p\}$ , obtained by fixing all points of  $\Omega'$  individually.

Whenever we are dealing with groups, we use a specific, succinct encoding. Instead of explicitly representing each element of the group, we only store a subset that is sufficient to obtain any other element through composition. Formally, let  $S$  be a subset of the group  $\Gamma$ , i.e.,  $S \subseteq \Gamma$ . We call  $S$  a *generating set* of  $\Gamma$  whenever we obtain precisely  $\Gamma$  when exhaustively composing elements of  $S$ . We write  $\langle S \rangle = \Gamma$ . Moreover, each individual element  $\varphi \in S$  can be referred to as a *generator* of  $\Gamma$ .


We write  $\text{supp}(\varphi) := \{\omega \mid \omega \in \Omega \wedge \varphi(\omega) \neq \omega\}$  for the *support of a map*, meaning points of  $\Omega$  not fixed by  $\varphi$ . The *support of a group*  $\Gamma \in \text{Sym}(\Omega)$  is the union of all supports of elements of  $\Gamma$ , i.e.,  $\text{supp}(\Gamma) := \{\omega \mid \omega \in \Omega \wedge \exists \varphi \in \Gamma : \varphi(\omega) \neq \omega\}$ .

We use the *cycle notation* for permutations  $\varphi: \Omega \rightarrow \Omega$ . The permutation of  $\{1, \dots, 5\}$  given by  $1 \mapsto 2, 2 \mapsto 3, 3 \mapsto 1, 4 \mapsto 5, 5 \mapsto 4$  we write as  $(1, 2, 3)(4, 5)$ . Note that, for example  $(1, 2, 3)(5, 4)$  and  $(3, 1, 2)(4, 5)$  denote the same permutation. Algorithmically the cycle notation enables us to read and store a permutation  $\varphi$  in time  $\text{supp}(\varphi)$ .

When considering two permutation groups  $\Gamma$  and  $\Gamma'$  it is possible that the groups are isomorphic as abstract groups but not as permutation groups. For example, if we let the symmetric group  $\text{Sym}(\Omega)$  act component-wise on pairs of elements of  $\Omega$ , we obtain a permutation group with domain  $\Omega^2$  that also has  $|\Omega|!$  many elements. In fact this group is isomorphic to  $\text{Sym}(\Omega)$  as an abstract group. We say a group  $\Gamma$  is a *symmetric group in natural action* if the group is  $\text{Sym}(\Omega)$ , where  $\Omega$  is the domain of  $\Gamma$ .

**SAT and Symmetries.** A Boolean satisfiability (SAT) instance  $F$  is commonly given in *conjunctive normal form* (CNF), which we denote with  $F = \{(l_{1,1} \vee \dots \vee l_{1,k_1}), \dots, (l_{m,1} \vee \dots \vee l_{m,k_m})\}$ , where each element of  $F$  is called a clause. A clause itself consists of a set of *literals*. A literal is either a variable or its negation. We use  $\text{Var}(F) := \{v_1, \dots, v_n\}$  for the set of *variables* of  $F$  and we use  $\text{Lit}(F)$  for its literals.

A symmetry, or *automorphism*, of  $F$  is a permutation of the literals  $\varphi: \text{Lit}(F) \rightarrow \text{Lit}(F)$  satisfying the following two properties. First, it maps  $F$  back to itself, i.e.,  $\varphi(F) \equiv F$ , where  $\varphi(F)$  is applied element-wise to the literals in each clause. Here clauses are equivalent, if they are the same when treated as unordered sets of literals, for example  $(x \vee y) \equiv (y \vee x)$ . Then  $F' \equiv F$  if  $F'$  is obtained from  $F$  by reordering the literals of  $F'$  within the clauses. Second, for all  $l \in \text{Lit}(F)$  it must hold that  $\overline{\varphi(l)} = \varphi(\bar{l})$ , i.e.,  $\varphi$  induces a permutation of the



■ **Figure 3** Example model graph  $\mathcal{M}(F_E) = \mathcal{M}(\{(x_1 \vee \bar{y}_1), (x_2 \vee \bar{y}_2), (x_3 \vee \bar{y}_3), (x_1 \vee x_2 \vee x_3 \vee z_1 \vee z_2)\})$ . The automorphisms of the graph correspond to the automorphisms of the formula. The coloring on the right side shows the orbit partition.

variables. For example the permutation mapping  $x_i$  to  $\neg x_{i+1}$  and  $x_i$  to  $x_{i+1}$ , with indices taken modulo 4, is a symmetry of  $(x_1 \vee \neg x_2 \vee x_3 \vee \neg x_4) \wedge (x_4 \vee \neg x_1 \vee x_2 \vee \neg x_3)$ . The permutation group of all symmetries of  $F$  is  $\text{Aut}(F) \leq \text{Sym}(\text{Lit}(F))$ .

It is well understood that the symmetries of a SAT formula  $F$  can be captured by a graph. We call this the *model graph* and denote it with  $\mathcal{M}(F)$ . While there exists various constructions for the model graphs, we use the following common construction. Each literal  $l \in \text{Lit}(F)$  is associated with a vertex  $l$ . Each clause  $C \in F$  is associated with a vertex  $C$ . All pairs of literals  $l$  and  $\bar{l}$  are connected by an edge. For all literals  $l \in C$  of a clause  $C$ , we connect vertices  $l$  and  $C$ . Lastly, to distinguish clause vertices from literal vertices, we color all clauses with color 0 and all literals with color 1. As desired, for this graph,  $\text{Aut}(F) = \text{Aut}(\mathcal{M}(F))|_{\text{Lit}(F)}$  holds [29].

Consider the formula  $F_E := \{(x_1 \vee \bar{y}_1), (x_2 \vee \bar{y}_2), (x_3 \vee \bar{y}_3), (x_1 \vee x_2 \vee x_3 \vee z_1 \vee z_2)\}$ . Throughout the paper, we use  $F_E$  as our running example. Figure 3 shows its model graph. Regarding the symmetries of  $F_E$ , note that there are symmetries interchanging all of  $x_1, x_2, x_3$ , all of  $y_1, y_2, y_3$  and of the  $z_1, z_2$ . A generating set  $S_E$  for  $\text{Aut}(F_E)$  is  $S_E = \{(x_1, x_2, x_3)(\bar{x}_1, \bar{x}_2, \bar{x}_3)(y_1, y_2, y_3)(\bar{y}_1, \bar{y}_2, \bar{y}_3), (x_1, x_2)(\bar{x}_1, \bar{x}_2)(y_1, y_2)(\bar{y}_1, \bar{y}_2), (z_1, z_2)\}$ .

**Orbits.** Given a permutation group  $\Gamma \leq \text{Sym}(\Omega)$ , we denote with  $\omega^\Gamma \subseteq \Omega$  the *orbit* of a point  $\omega \in \Omega$ . That is, an element  $\omega' \in \Omega$  is in  $\omega^\Gamma$  whenever there is a  $\varphi \in \Gamma$  with  $\varphi(\omega) = \omega'$ .

The orbits of our example  $\text{Aut}(F_E)$  are shown in Figure 3, e.g., the orbit  $\{\bar{z}_1, \bar{z}_2\}$  is green.

### 3 SAT-Symmetry Computational Setting

Let us describe the computational setting in which our new algorithms operate. Since we want the theoretical runtimes to reflect more closely the runtimes in practice, there are two important differences compared to the traditional computational group theory setting. These differences are in the measure of runtime as well as in the format of the input.

**Joint Graph/Group Pairs.** Typically, algorithms in computational group theory dealing with permutation groups assume as their input a generating set of permutations  $S$  of a group  $\Gamma = \langle S \rangle$ . While this is certainly a natural setting when discussing algorithms for groups in general, in our setting this input format disregards further information that is readily available. Therefore, we require that algorithms in the SAT-symmetry interface have access to more information about the input group. Specifically, we may require that the input consists of *both* a generating set  $S$  and a graph  $G$  with  $\langle S \rangle = \text{Aut}(G)$ . We call this a *joint graph/group pair*  $(G, S)$ . For our SAT context, we may moreover assume that the SAT formula  $F$  with  $\mathcal{M}(F) = G$  is available, whenever necessary.

**Instance-Linear Runtime.** In computational group theory, given a generating set  $S$  for a permutation group  $\langle S \rangle \leq \text{Sym}(\Omega)$ , a runtime of  $\mathcal{O}(|S||\Omega|)$  is typically considered linear time [30]. This is however only a very crude upper bound when seen in terms of the actual encoding size of a given generating set. In particular, when generators are sparse, as is common in SAT [29, 11], linear time in this sense is *not* necessarily linear time in the encoding size, which is what we would use in a graph-theoretic or SAT context.

Specifically, we are interested in measuring the runtime of algorithms relative to the encoding size of a generating set given in a sparse format. Therefore, we define the encoding size of a generating set  $S$  as  $\text{enc}(S) := \sum_{p \in S} |\text{supp}(p)|$ .

In particular, given a SAT formula  $F$ , graph  $G = (V, E)$ , and generating set  $S$ , the goal is to have algorithms that (ideally) run in time linear in  $|F| + |V| + |E| + \text{enc}(S)$ . In order to not confuse the “types of linear time”, we refer to such algorithms as *instance-linear*. Analogously, an algorithm has *instance-quasi-linear* time if it runs in time  $\mathcal{O}((|F| + |V| + |E| + \text{enc}(S)) \cdot (\log(|F| + |V| + |E| + \text{enc}(S))))^c$  for some constant  $c$ .

**Illustrative Examples.** The task of computing the orbits is an excellent example demonstrating the usefulness instance-quasi-linear time. As transitive closure, we can find the orbit  $\Delta$  of an element in time  $\mathcal{O}(|\Delta||S|)$  [30]. However, with instance-quasi-linear time in mind, we quickly arrive at an algorithm to compute the entire orbit partition in time  $\mathcal{O}(\text{enc}(S) \cdot \alpha(\text{enc}(S)))$  using a union-find data structure, where  $\alpha$  is the inverse Ackermann function. The inverse Ackermann function exhibits substantially slower growth than  $\log(n)$ .

Furthermore, having access to the graph of a graph/group pair  $(G, S)$  gives a significant advantage in what is algorithmically possible. A good example of the difference is that testing membership  $\varphi \in \langle S \rangle$  is much easier for the graph/group pair: testing  $\varphi(G) = G$  (which is true if and only if  $\varphi \in \langle S \rangle$ ) can be done in instance-linear time. However, testing  $\varphi \in \langle S \rangle$  without access to the graph is much more involved. The best known method for the latter involves computing a strong generating set, corresponding base and Schreier table [30], followed by an application of the fundamental sifting algorithm [30]. Even performing only the last step of this process (sifting) is not guaranteed to be in instance-linear time.

## 4 Favorable Group Structures in SAT

We now propose problems which should be solved within the SAT-symmetry computational setting. We analyze heuristics used in advanced symmetry exploitation algorithms [13, 25, 17], extracting precise formal definitions.

**Disjoint Direct Decomposition.** Following [9], we say a direct product of a permutation group  $\Gamma = \Gamma_1 \times \Gamma_2 \times \cdots \times \Gamma_r$  is a *disjoint direct decomposition* of  $\Gamma$ , whenever all  $\Gamma_i$  have pairwise disjoint supports. We call  $\Gamma_i$  a *factor* of the disjoint direct decomposition of  $\Gamma$ . A disjoint direct decomposition is *finest*, if we cannot decompose any factor further into a non-trivial disjoint direct decomposition. A recent algorithm solves the problem of computing the finest disjoint direct decomposition for permutation groups in polynomial-time [9].

In our running example, the finest disjoint direct decomposition of  $\text{Aut}(F_E)$  splits the group into a subgroup  $H_1$  permuting only the  $x_i$  and  $y_i$  variables, and a subgroup  $H_2$  permuting  $z_1$  and  $z_2$ . Indeed, setting  $H_1 = \langle \{(x_1, x_2, x_3)(\bar{x}_1, \bar{x}_2, \bar{x}_3)(y_1, y_2, y_3)(\bar{y}_1, \bar{y}_2, \bar{y}_3), (x_1, x_2)(\bar{x}_1, \bar{x}_2)(y_1, y_2)(\bar{y}_1, \bar{y}_2)\} \rangle$ , and  $H_2 = \langle \{(z_1, z_2)\} \rangle$ , we have  $\text{Aut}(F_E) = H_1 \times H_2$ .

Computing a disjoint direct decomposition is a typical routine in symmetry exploitation tools [13, 25, 17]. It allows for separate treatment of each factor of the decomposition. The heuristics in use today do not guarantee that the decomposition is the finest disjoint direct

decomposition: indeed, the heuristics of the tools mentioned above assume that the given generating sets are already *separable* [9]. This means it is assumed, that every generator  $\varphi \in S$  only operates on one factor of the disjoint direct decomposition  $\Gamma_1 \times \Gamma_2 \times \dots \times \Gamma_r$ . Formally, this means for each  $\varphi \in S$  there is one  $i \in \{1, \dots, r\}$  for which  $\text{supp}(\varphi) \cap \text{supp}(\Gamma_i) \neq \emptyset$  holds, and for all  $j \in \{1, \dots, r\}, j \neq i$  it holds that  $\text{supp}(\varphi) \cap \text{supp}(\Gamma_i) = \emptyset$ . For example, the generating set  $S_E$  we gave for  $\text{Aut}(F_E)$  is separable.

It is not known how often generating sets given for graphs of SAT formulas are separable for a given symmetry detection tool, or in particular after reducing the domain to the literals of the SAT formula. It is however obvious that for the most advanced general-purpose symmetry detection tools, TRACES and DEJAVU, generators are not separable with very high probability due to the use of randomly selected generators [9].

**Row Interchangeability.** We now discuss the concept of row interchangeability [13, 14, 25]. Let  $F$  be a SAT formula. Let  $M$  be a variable matrix  $M: \{1, \dots, r\} \times \{1, \dots, c\} \rightarrow \text{Var}(F)$ . We denote the entries of  $M$  with  $x_{i,j}$  where  $i \in \{1, \dots, r\}$  and  $j \in \{1, \dots, c\}$ . We define the shorthand  $\text{supp}(M) := \{x_{i,j} \mid x_{i,j} \in M\} \cup \{\bar{x}_{i,j} \mid x_{i,j} \in M\}$ . The set  $\text{supp}(M)$  denotes all the literals involved with the matrix  $M$ . We say  $F$  exhibits *row interchangeability* if there exists a matrix  $M$  such that for every permutation  $\varphi \in \text{Sym}(\{1, \dots, r\})$ , for the induced literal permutation  $\varphi': \text{Lit}(F)|_{\text{supp}(M)} \rightarrow \text{Lit}(F)|_{\text{supp}(M)}$  given by  $x_{i,j} \mapsto x_{\varphi(i),j}, \neg x_{i,j} \mapsto \neg x_{\varphi(i),j}$  it holds that  $\varphi' \in \text{Aut}(F)|_{\text{supp}(M)}$ . Indeed, if this is the case, we can observe that the matrix  $M$  describes a subgroup of  $\text{Aut}(F)$  consisting of  $\{\pi \in \text{Aut}(F) \mid \exists \varphi \in \text{Sym}(\{1, \dots, r\}) : \varphi' = \pi|_{\text{supp}(M)}\}$ . We denote this group by  $H_M \leq \text{Aut}(F)$ .

A crucial fact is that for  $H_M|_{\text{supp}(M)}$ , linear-sized complete symmetry breaking is available [13, 14]. As is also in part discussed in [13, 14], we observe that the complete symmetry breaking for  $H_M$  is most effective whenever  $H_M$  is the only action on  $\text{supp}(M)$  in  $\text{Aut}(F)$ , or more precisely,  $\text{Aut}(F)|_{\text{supp}(M)} = H_M|_{\text{supp}(M)}$ . In this case, we call  $H_M$  an *elementary* row interchangeability subgroup. Otherwise, there are non-trivial symmetries  $\varphi \in \text{Aut}(F)|_{\text{supp}(M)}$  with  $\varphi \notin H_M|_{\text{supp}(M)}$  or  $\text{supp}(M)$  is not a union of orbits. Indeed, in this case, the complete symmetry breaking of  $H_M$  might make it more difficult to break such overlapping symmetries  $\varphi$ : for example, if two row interchangeability subgroups  $H_M$  and  $H_{M'}$  overlap, i.e.,  $\text{supp}(M) \cap \text{supp}(M') \neq \emptyset$ , complete symmetry breaking can only be guaranteed for one of them using the technique of [13].

Whenever  $H_M$  is an elementary row interchangeability subgroup, the situation is much clearer: we can produce a linear-sized complete symmetry breaking formula and this covers at least all symmetries on the literals  $\text{supp}(M)$ . In this paper, we therefore focus on computing elementary row interchangeability groups.

Let us consider  $F_E$  again: for the matrix  $M := \begin{bmatrix} x_1 & x_2 & x_3 \\ y_1 & y_2 & y_3 \end{bmatrix}$  there is indeed a row interchangeability subgroup. (Recall that the group  $H_M$  permutes positive and negative literals of variables appearing in  $M$ .) For this example,  $H_M$  is both an elementary row interchangeability group and a factor in the finest direct disjoint decomposition of  $\text{Aut}(F_E)$ .

**Row Interchangeability and Equivalent Orbits.** We now describe the matrix of elementary row interchangeability groups in more group-theoretic terms. We first define the notion of equivalent orbits:

► **Definition 1** (Equivalent orbits; see [30, Subsection 6.1.2]). *Two orbits  $\Delta_1, \Delta_2$  are equivalent, if and only if there is a bijection  $b: \Delta_1 \rightarrow \Delta_2$  such that for all  $\varphi \in \Gamma$  and  $\delta \in \Delta_1$ ,  $\varphi(b(\delta)) = b(\varphi(\delta))$ .*

We write  $\Delta_1 \equiv \Delta_2$  to indicate orbits  $\Delta_1$  and  $\Delta_2$  are equivalent. It is easy to see this indeed defines an equivalence relation on the orbits [30].



We observe that if a row interchangeability subgroup  $H_M$  is elementary, each row of the matrix  $M$  is an orbit of  $\text{Aut}(F)$ . Since all rows are moved simultaneously in the same way, we remark that rows of  $M$  are precisely equivalent orbits with a natural symmetric action:

► **Lemma 2.** *Let  $H_M$  be a row interchangeability subgroup of  $\Gamma = \text{Aut}(\mathcal{M}(F))$ , and let  $\Delta_i = (x_1, \dots, x_c)$  denote a row of  $M$ .  $H_M$  is an elementary row interchangeability subgroup if and only if all of the following hold: (1)  $\Delta_i$  is an orbit with a natural symmetric action in  $\Gamma$ . (2) For every other row  $\Delta_j$  of  $M$ , it holds that  $\Delta_i \equiv \Delta_j$ . (3) For  $\overline{\Delta}_i = (\overline{x}_1, \dots, \overline{x}_c)$ ,  $\overline{\Delta}_i$  is also an orbit with  $\overline{\Delta}_i \equiv \Delta_i$ .*

There is an exact algorithm which computes equivalent orbits in essentially quadratic runtime [30]. Again, runtimes are difficult to compare due to different pre-conditions in [30]. In any case, the algorithm for equivalent orbits depends on computing a base and strong generating set, which is too slow from our perspective.

We may split detecting elementary row interchangeability groups into detecting *natural symmetric action* on the orbits, followed by computing *equivalent orbits*. We now turn to solving the problems defined above in the computational setting of the SAT-symmetry interface. Specifically, we propose algorithms for the finest disjoint direct decomposition (Section 5), natural symmetric action (Section 6), and equivalent orbits (Section 7).

## 5 Finest Disjoint Direct Decomposition

Having established the problems we want to address, we now turn to presenting suitable algorithms in the SAT-symmetry computational setting. In particular, recall that we want to make use of joint graph/group pairs in order to state algorithms that run in instance-quasi-linear time. We begin by computing the finest disjoint direct decomposition.

Specifically, given a joint graph/group pair  $(G, S)$ , our aim is to compute the finest disjoint direct decomposition of the group  $\langle S \rangle$ . Our proposed algorithm, given the orbits, can do so in instance-linear time. The disjoint direct decomposition of a group allows us to separately treat each factor of the decomposition in symmetry exploitation or other consecutive algorithms.

To simplify the discussion, we assume the graph  $G$  to be undirected. However, the procedure generalizes to both directed and even edge-colored graphs.

**Orbit Graph.** We describe the *orbit graph*, which can be constructed from  $(G, S)$ . We are particularly interested in the connected components of the orbit graph, which turn out to correspond exactly to the factors of the finest disjoint direct decomposition.

First, note that the orbit partition  $\pi$  of  $\langle S \rangle$  can be viewed as a vertex coloring of the graph  $G$ , assigning to every vertex its orbit. We consider the graph  $G' = (V(G), E(G), \pi)$ , i.e.,  $G$  colored with its orbit partition (see Figure 4, left).

We call two distinct orbits  $\Delta, \Delta'$  *homogeneously connected*, whenever either all vertices  $v \in \Delta$  are adjacent to all vertices of  $\Delta'$ , or there is no edge with endpoints both in  $\Delta$  and  $\Delta'$ . Indeed, we could “flip edges” between homogeneously connected orbits such that they all become disconnected, without changing the automorphism group (see Figure 4, middle).

We now give the formal definition of the orbit graph. The orbit graph is essentially an adapted version of the so-called flipped quotient graph (see [20] for a discussion). The vertex set of the orbit graph is the set of orbits of  $\langle S \rangle$ , i.e.,  $\{\pi^{-1}(v) \mid v \in V(G')\}$ . Two orbits  $\Delta, \Delta'$  are adjacent in the orbit graph if and only if the orbits are *not* homogeneously connected in the original graph  $G$  (see Figure 4, right).

■ **Algorithm 1** Compute the orbit graph.


---

```

1 function OrbitGraph( $G'$ )
  Input : graph  $G' = (V, E, \pi)$ , where  $\pi$  is the orbit partition of  $\text{Aut}(G')$ 
  Output : orbit graph  $G_O$ 
2 initialize integer array  $A$  of size  $|V|$  with all 0;
3 initialize empty list  $W$ ;
4 initialize empty graph  $G_O$ ;
5  $V(G_O) := \pi(V)$ ;
6 for ( $\Delta \in \pi(V)$ )
7   pick an arbitrary  $v \in \Delta$ ;
8   for ( $(v, v') \in E$ )
9     increment  $A[\pi(v')]$ ;
10    add  $\pi(v')$  to  $W$ ;
11   for ( $\Delta' \in W$ )
12     if  $A[\Delta'] > 0$  and  $A[\Delta] < |\pi^{-1}(\Delta')|$  then add edge  $(\Delta, \Delta')$  in  $G_O$  ;
13      $A[\Delta'] := 0$ ;
14 return  $G_O$ ;

```

---




■ **Figure 4** Model graph of  $F_E$  colored with its orbit partition on the left. The corresponding graph with flipped edges is in the middle, which disconnects parts of the graph. On the right the orbit graph is shown, whose 3 connected components correspond to the factors of the finest disjoint direct decomposition.

*Description of Algorithm 1.* Algorithm 1 describes how to compute the orbit graph from  $G'$ . The algorithm first initializes the graph  $G_O$  with a vertex set that contains exactly one vertex for each orbit of  $G'$ . It then counts for each orbit  $\Delta$ , how many neighbors a vertex  $v \in \Delta$  has in the other orbits. Since  $\Delta$  is an orbit, this number is the same for all vertices, so it suffices to compute this for one  $v \in \Delta$ . Finally, the algorithm checks to which other colors the vertex  $v$  and thus the orbit  $\Delta$  is *not* homogeneously connected (Line 12). If  $\Delta$  and  $\Delta'$  are not homogeneously connected, the edge  $(\Delta, \Delta')$  is added to the orbit graph.

*Remark on the runtime of Algorithm 1.* Using appropriate data structures for graphs (adjacency lists) and colorings (see [23], which in particular includes efficient ways to compute  $|\pi^{-1}(C')|$ ), the algorithm can be implemented in instance-linear time.

**Orbit Graph to Decomposition.** Indeed, the connected components of the orbit graph represent precisely the factors of the finest disjoint direct decomposition of the automorphism group of the graph:

► **Lemma 3.** *Let  $\Gamma = \text{Aut}(G)$ . The vertices represented by a connected component of the orbit graph of  $G$  are all in the same factor of the finest direct disjoint decomposition of  $\Gamma$  and vice versa.*



■ **Figure 5** Illustration of two orbits that are non-homogeneously connected on the left in blue and purple. On the right, fixing a vertex of one orbit indicated in red immediately partitions the other orbit into two orbits: the neighbors of the red vertex in green, and the non-neighbors in orange.

**Proof.** Consider two orbits  $\Delta_1, \Delta_2$  of  $\Gamma$  in different factors of any direct disjoint decomposition. Towards a contradiction, assume  $\Delta_1, \Delta_2$  are not homogeneously connected. Note that naturally, the orbit coloring is equitable. Since the orbit coloring is equitable, the connection must be regular, i.e., each vertex of  $\Delta_1$  has  $d_1$  neighbors in  $\Delta_2$ , and every vertex of  $\Delta_2$  has  $d_2$  neighbors in  $\Delta_1$  for some integers  $d_1, d_2$ . However,  $0 < d_1 < |\Delta_2|$  and  $0 < d_2 < |\Delta_1|$  hold.

Let us now fix a point  $\delta \in \Delta_1$ , i.e., consider the point stabilizer  $\Gamma_{(\delta)}$ . If two orbits are in different factors of a direct disjoint decomposition, fixing a point of  $\Delta_1$  must not change the group action on  $\Delta_2$ . In particular,  $\Delta_2$  must be an orbit of  $\Gamma_{(\delta)}$ . However,  $\delta$  is adjacent to some vertex  $\delta' \in \Delta_2$  and non-adjacent to some vertex  $\delta'' \in \Delta_2$  (see Figure 5 for an illustration). Having fixed  $\delta$ , we can therefore not map  $\delta'$  to  $\delta''$ . This contradicts the assumption that  $\Delta_1$  and  $\Delta_2$  are in different factors of a direct disjoint decomposition. Hence, orbits in different factors of a direct disjoint decomposition must be homogeneously connected in  $G$ , i.e., non-adjacent in the orbit graph.

Now assume  $\Delta_1$  and  $\Delta_k$  are in the same component in the orbit graph. Then, there must be a path of orbits  $\Delta_1, \Delta_2, \dots, \Delta_k$  where each  $\Delta_i, \Delta_{i+1}$  is *not* homogeneously connected. In this case, we know for each  $i \in \{1, \dots, k-1\}$  that  $\Delta_i$  and  $\Delta_{i+1}$  must be in the same factor of any disjoint direct decomposition. Therefore,  $\Delta_1$  and  $\Delta_k$  must be in the same factor of every disjoint direct decomposition.

On the other hand, if  $\Delta_1$  and  $\Delta_k$  are in different components in the orbit graph, they are in different factors of the finest disjoint direct decomposition. ◀

Since connected components can be computed in linear time in the size of a graph, and the size of the orbit graph is at most linear in the size of the original graph, we can therefore compute the finest direct disjoint decomposition in instance-linear time. In a consecutive step, the generators could be split according to factors, producing a separable generating set, again in instance-linear time. This is done by separating each generator into the different factors. Finally, given the finest direct disjoint decomposition, we can again produce a joint graph/group pair for each factor, by outputting for a factor  $H_i$  the induced subgraph  $G'[H_i]$ . We summarize the above in a theorem:

► **Theorem 4.** *Given a joint graph/group pair  $(G, S)$  and orbit partition of  $\langle S \rangle$ , there is an instance-linear algorithm which computes the following:*

1. *The finest disjoint direct decomposition  $\langle S \rangle = H_1 \times H_2 \times \dots \times H_m$ .*
2. *A separable generating set  $S'$  with  $\langle S' \rangle = G$ .*
3. *For all factors  $i \in \{1, \dots, m\}$  a joint graph/group pair  $(G_i, S_i)$  with  $\text{Aut}(G_i) = \langle S_i \rangle = H_i$  in instance-linear time.*

We recall that if the orbit partition of  $\langle S \rangle$  is not yet available, we can compute it in instance-quasi-linear time.

**Domain Reduction to SAT Literals.** For a SAT formula  $F$ , we can apply the above procedure to its model graph  $\mathcal{M}(F)$ . However, as mentioned above, in SAT we are typically only interested in symmetries for a subset of vertices, namely the vertices that represent

literals. Therefore, we are specifically interested in the finest direct disjoint decomposition of the automorphism group reduced to literal vertices  $\text{Aut}(\mathcal{M}(F))|_{\text{Lit}(F)}$ . The crucial point here is that when removing orbits that represent clauses, orbits of literal vertices can become independent and the disjoint direct decomposition can therefore become finer. We cannot simply apply our algorithm for the induced group  $\text{Aut}(\mathcal{M}(F))|_{\text{Lit}(F)}$  since this is not a joint graph/group pair. Of course we could apply the algorithm from [9] that computes finest disjoint direct decomposition for permutation groups in general. However, we can detect *some* forms of independence by simple means using the original joint graph/group pair. Indeed, we will describe an algorithm that checks in instance-linear time whether the parts in a given partition of the literals are independent. We can thus at least check whether a given partition induces a disjoint direct decomposition (the proof can be found in the full version [5]):

► **Theorem 5.** *Let  $F$  be a CNF-Formula and  $(\mathcal{M}(F), S)$  be a joint graph/group pair for the model graph of  $F$ . Given a partition  $\text{Lit}(F) = L_1 \cup L_2 \cup \dots \cup L_t$  of the literals of  $F$ , the pair  $(\mathcal{M}(F), S)$ , and its orbits, we can check in instance-linear time whether the partition induces a disjoint direct product (that is, whether  $\text{Aut}(F) = \text{Aut}(F)_{(L \setminus L_1)} \times \dots \times \text{Aut}(F)_{(L \setminus L_t)}$ ).*

## 6 Natural Symmetric Action

Before we can begin our discussion of the natural symmetric action, we need to discuss generating (nearly-)uniform random elements (see [30]) of a given permutation group  $\langle S \rangle$ . There is no known algorithm which produces uniform random elements of  $\langle S \rangle$  in quasi linear time, even in computational group theory terminology [30]. However, there are multiple ways to produce random elements, most of which are proven to work well in practice, and can be implemented fairly easily [30, 15]. In this paper, we attempt to only make use of random elements sparingly. Whenever we do, as is common in computational group theory, we do not consider the particular method used to generate them and simply denote the runtime of the generation with  $\mu$ . Moreover, we discuss potential synergies in the SAT-symmetry context which might help to avoid random elements in practice, whenever applicable.

We now explain how to efficiently test whether a permutation group is a symmetric group in natural action. Then, we describe more generally how to determine simultaneously for all orbits of a permutation group whether the induced action is symmetric in natural action.

Detecting symmetric permutation groups in their natural action is a well-researched problem in computational group theory. State-of-the-art practical implementations are available in modern computer algebra systems (such as in [24]). Typically, a natural symmetric action is detected using a so-called probabilistic *giant* test, followed by a test to ensure that the group is indeed symmetric. The tests work by computing (nearly) uniform random elements of the group and inspecting them for specific properties.

A permutation group is called a giant if it is the symmetric group or the alternating group in natural action. In many computational contexts, giants are by far the largest groups that appear, hence their name. Because of this, giants form bottleneck cases for various algorithms and therefore often need to be treated separately. To test whether a permutation group is a symmetric group in natural action, we first test whether the group is a giant.

We leverage the following facts:

► **Fact 6** (see [30, Corollary 10.2.2.]). *If a permutation group of degree  $n$  contains an element with a cycle of length  $p$  for some prime  $p$  with  $n/2 < p < n - 2$  then  $G$  is a giant.*

► **Fact 7** (see [30, Corollary 10.2.3.]). *The proportion of elements in  $S_n$  containing a cycle of length  $p$  for a prime  $p$  with  $n/2 < p < n - 2$  is asymptotically  $1/\log(n)$ .*

Collectively, these statements show that we only need to generate few random elements of a group and inspect their cycle lengths to detect a giant. To then distinguish between the alternating group and the symmetric group, we can check whether all generators belong to the alternating group. This can be attained using basic routines, such as examining the so-called parity of a generator (see [30] for more details). Algorithm 2 generalizes the probabilistic test for a transitive group [30] to a test which is performed simultaneously to check for a natural symmetric action on all the orbits of a group.

■ **Algorithm 2** Compute whether orbits induce a natural symmetric action.

---

```

1 function SymmetricAction( $S, O$ )
   Input : generators  $S$  with  $\langle S \rangle = \Gamma \leq \text{Sym}(\Omega)$ , orbits  $O := \{\Delta_1, \dots, \Delta_m\}$  of  $\Gamma$ 
   Output : set of orbits  $S_O \subseteq O$ , where  $\Delta \in S_O$  induces a natural symmetric action
2 // first, we filter orbits which can at most be alternating
3 for ( $\Delta_i \in O$ )
4    $t := \top$ ;
5   for ( $s \in S$ )
6     if  $s|_{\Delta_i} \notin \text{Alt}(\Delta_i)$  then  $t := \perp$ ; //  $\Delta_i$  cannot be alternating
7     if  $t = \top$  then  $O := O \setminus \{\Delta_i\}$ ; //  $\Delta_i$  cannot induce symmetric action
8 // second, we test whether orbits are giants
9  $S_O := \emptyset$ ;
10 for ( $\_ \in \{1, \dots, c \log(|\Omega|)^2\}$ ) // repeat  $c \log(|\Omega|)^2$  times
11    $\varphi :=$  uniform random element of  $\langle S|_{\cup_{\Delta_i \in S_O} \Delta_i} \rangle$ ;
12   for ( $\Delta_i \in O$ )
13      $p :=$  cycle length of longest cycle in  $\varphi$  on  $\Delta_i$ ;
14     if  $p > n/2 \wedge p < n - 2 \wedge p$  prime then
15        $S_O := S_O \cup \{\Delta_i\}$ ; //  $\Delta_i$  induces symmetric action
16        $O := O \setminus \{\Delta_i\}$ ;
17 return  $S_O$ 

```

---

*Description of Algorithm 2.* Overall, the algorithm samples uniform random elements of the group and checks whether the random elements exhibit long prime cycles (see Fact 6). More precisely, the algorithm first distinguishes between potential alternating and symmetric groups on each orbit. Then, it computes  $d = c \log(|\Omega|)^2$  random elements. For each random element and each orbit, we then apply the giant test (Fact 6) to check whether the element certifies that the orbit induces a natural symmetric action.

*Runtime of Algorithm 2.* Let us assume access to random elements of the joint graph/group pair  $(G, S)$  with  $\langle S \rangle \leq \text{Sym}(\Omega)$  in time  $\mu$ . Assuming a random element can be produced in time  $\mu$ , the algorithm runs in worst-case time  $\mathcal{O}(\log(|\Omega|)^2(\mu + |\Omega|))$ .

*Correctness of Algorithm 2.* Regarding the correctness of the algorithm, the interesting aspect is to discuss the error probability. We argue that the error probability is at most  $1/4$  if  $c$  is chosen to larger than  $2 \ln(2)$ . Practical implementations use  $c = 20$  in similar contexts [24].

If an orbit  $\Delta$  does not induce a symmetric action, no error can be made. If an orbit  $\Delta$  induces a symmetric action, by Fact 7, the probability that one iteration does not produce a long prime cycle for  $\Delta$  is at most  $(1 - 1/\log(n))$ . Thus, the probability that none of the  $c \log(|\Omega|)^2$  iterations produces a long prime cycle for  $\Delta$  is bounded

by  $(1 - 1/\log(n))^{c \log(|\Omega|)^2} \leq (1/e)^{c \log(|\Omega|)} \leq 1/(4|\Omega|)$  since  $c > 2 \ln(2)$ . Since there can be at most  $|\Omega|$  orbits, using the union bound, we get that the probability that the test fails for at least one of the orbits is at most  $|\Omega| \cdot 1/(4|\Omega|) = 1/4$ .

When trying to compute a natural symmetric action on a graph/group pair, the following heuristics can be implemented in instance-linear time.

The first and most straightforward heuristic is that most of the time, it is fairly clear that the generators describe a natural symmetric action. In particular, symmetry detection based on depth-first search seem to often produce generators that are transpositions. From these symmetric actions can be detected immediately. This fact is implicitly used by column interchangeability heuristics in use today. There are however many more ways to detect a natural symmetric action, many of which are implemented in modern computer algebra systems such as [15, 24].

Next, the symmetry detection preprocessor SASSY [6] as well as the preprocessing used by TRACES sometimes detect a natural symmetric action on an orbit, by detecting certain structures of a graph. In these cases, the result should be immediately communicated to consecutive algorithms. We can also use the graph structure to immediately discard orbits from the test of Algorithm 2. In particular, all orbits  $\Delta$  where  $G[\Delta]$  is neither the empty graph nor the complete graph cannot have a natural symmetric action.

Furthermore, the generators produced by DEJAVU and TRACES are fairly random (for some parts even uniformly random [3]). This means they should presumably work well with the probabilistic tests above. Lastly, internally, symmetry detection tools often produce so-called Schreier-Sims tables [30], which can be used to produce random elements effectively.

Indeed, for our running example  $F_E$ , the natural symmetric action can be detected quite easily: let us consider the generators  $S_E$  reduced to the orbit of  $\{x_1, x_2, x_3\}$ . We observe that there is a generator  $(x_1, x_2, x_3)$  and  $(x_1, x_2)$ . While this is not a set of generators detected by current symmetry exploitation algorithms [13, 25], this is indeed also an arguably obvious encoding of a natural symmetric action: for an orbit of size  $n$ , an  $n$ -cycle in conjunction with a transposition encodes a symmetric action.

## 7 Equivalent Orbits

Towards our overall goal to compute row interchangeability subgroups, we can now already determine which orbits induce a natural symmetric action. By Lemma 2, to detect elementary row interchangeability subgroups, we only miss a procedure for orbit equivalence.

We describe now how to compute equivalent orbits as the automorphism group of a special, purpose-built graph. Then, we give a faster algorithm computing equivalent orbits with natural symmetric action in instance-quasi-linear time, under mild assumptions. In particular, we can find all classes of equivalent orbits described by Lemma 2.

### 7.1 Cycle Type Graph

If two orbits are equivalent, they appear in every permutation in “the same manner”: for example, if orbits  $\Delta_1$  and  $\Delta_2$  are equivalent then for every generator  $g$ , the cycle types  $g$  induces on  $\Delta_1$  are the same as the cycle types  $g$  induces on  $\Delta_2$ . More generally, equivalent orbits must be equivalent with respect to *every* generating set of the group. We introduce the *cycle type graph* whose symmetries capture orbit equivalence. This means we can use a symmetry detection tool to detect equivalent orbits.

For a group  $\Gamma \leq \text{Sym}(\Omega)$  and generating set  $\langle S \rangle = \Gamma$ , we define the *cycle type graph*  $\mathcal{C}(S)$  as follows.

Firstly, the **vertex set** of  $\mathcal{C}(S)$  is the disjoint union  $V(\mathcal{C}(S)) := \Omega \dot{\cup} \bigcup_{g \in S} \text{supp}(g)$ . In other words, there is a vertex for each element of  $\Omega$  and there are separate elements for all the points moved by the generators. In particular if a point is moved by several generators there are several copies of the point.

Secondly, the **edges** of  $\mathcal{C}(S)$  are added as follows: each corresponding vertex for  $x \in \text{supp}(g)$  is adjacent to the corresponding vertex for element  $x \in \Omega$  via an undirected edge. Furthermore, there are directed edges  $\{(x, g(x)) \mid x \in \text{supp}(g)\}$ . In other words, for each generator  $s_i \in S$ , we add directed cycles for each cycle of the generator, as shown in Figure 6a. In the following, we refer to directed cycles added in this manner as *cycle gadgets*.

Thirdly, we define a **vertex coloring** for  $\mathcal{C}(S)$ . For this we enumerate the generators, i.e.,  $S = \{s_1, \dots, s_m\}$ . We then color the vertices in  $\Omega$  with color 0 and an  $x \in \text{supp}(g_i)$  is colored with  $(i, t)$ , where  $t$  is the length of the cycle in  $g_i$  containing  $x$ . With this, the cycle type graph is constructed in such a way that its automorphism structure captures equivalence of orbits, as is described in more detail below.

We record several observations on automorphisms of the cycle type graph (missing proofs can be found in the full version [5]).

► **Lemma 8.** *If  $\Delta_1, \Delta_2$  are orbits of  $\Gamma$  then there is some  $b \in \text{Aut}(\mathcal{C}(S))$  for which  $b(\Delta_1) = \Delta_2$ , if and only if  $\Delta_1 \equiv \Delta_2$ .*

We may formulate the observations in group theoretic terms, giving the following lemma.

► **Lemma 9.** *Given a group  $\Gamma \leq S(\Omega)$ , its centralizer in the symmetric group  $C_{\text{Sym}(\Omega)}(\Gamma)$  and the cycle type graph  $\mathcal{C}(\Gamma)$  are a joint graph/group pair, i.e.,  $\text{Aut}(\mathcal{C}(\Gamma)) = C_{S(\Omega)}(\Gamma)$ .*

Since the centralizer of  $\text{Sym}(\Omega)$  in  $\text{Sym}(\Omega)$  is trivial for  $|\Omega| > 2$ , we get the following corollary.


► **Corollary 10.** *If  $\Gamma = \text{Sym}(\Omega)$  with  $|\Omega| > 2$ , the cycle type graph of  $\Gamma$  is asymmetric.*

It follows from the corollary that for two equivalent orbits with a natural symmetric action the bijection  $b$  commuting with the generators and interchanging the orbits is in fact unique.

While the cycle type graph and the centralizer in the symmetric group  $(\mathcal{C}(\Gamma), C_{\text{Sym}(\Omega)}(\Gamma))$  is a joint graph/group pair, we still have to compute the group: so far, we only have access to  $\mathcal{C}(\Gamma)$ . One option is a symmetry detection tool. However, this goes against our goal of invoking symmetry detection tools unnecessarily often – and against our goal to find instance-linear algorithms. Hence, instead of computing the entire automorphism group, our approach is to make due with less: in the following, we enhance the cycle type graph in a way such that it becomes “easy” for color refinement. Color refinement is usually applied as a heuristic approximating the orbit partition of a graph. However, on the enhanced graphs, we prove that color refinement is guaranteed to compute the orbit partition. Then, we show that the orbit partition suffices to determine equivalent orbits. Overall, these methods are only guaranteed to work for orbits with a natural symmetric action, as is the case in elementary row interchangeability groups.

## 7.2 Symmetries of Cycle Type Graph with Unique Cycles

Our goal is now to enhance the cycle type graph such that color refinement is able to compute its orbit partition. This in turn enables us to detect equivalent orbits, and in turn elementary row interchangeability groups. Towards this goal, we first discuss an algorithm to compute



**Figure 6** The cycle type graph and enhanced cycle type graph. The figure shows a cycle type gadget and canonical cyclic order for the permutation  $(1234)(5678)$ . Automorphisms of this graph are elements of the centralizer, which indicate equivalent orbits.

unique cycles on orbits. Unique cycles are a key ingredient for our enhanced cycle type graph. These cycles should be invariant with respect to an ordered generating set, a concept we explain first.

Given an ordered generating set  $(s_1, \dots, s_m)$  for a permutation group  $\Gamma = \langle \{s_i \mid i \in \{1, \dots, m\}\} \rangle \leq \text{Sym}(\Omega)$ , a permutation  $\varphi \in \text{Sym}(\Omega)$  *fixes* the ordered generating set (pointwise under conjugation) if for all  $s_i$  we have  $\varphi \circ s_i \circ \varphi^{-1} = s_i$ .

A permutation  $\pi \in \text{Sym}(\Omega)$  is *invariant* with respect to the ordered generating set if all permutations  $\varphi$  that fix  $(s_1, \dots, s_m)$  under conjugation also fix  $(s_1, \dots, s_m, \pi)$  under conjugation<sup>1</sup>. Note that all group elements in  $\pi \in \Gamma$  are invariant. However, there can be further invariant permutations.

We say a permutation  $\pi$  has a *unique cycle* if for some length  $\ell > 1$ , the permutation  $\pi$  contains exactly one cycle of  $\ell$ . We now describe a two-step process. Step one is to compute an invariant permutation with a unique cycle. Step two is to use this to compute an invariant permutation with a cyclic order.

**Unique Cycle from Generators.** As a first step we now need an invariant *unique cycle* to proceed. We argue how to compute such a cycle for orbits on which our group induces a natural symmetric action.


We may use random elements to find a unique cycle. In fact, if we perform the giant test of Section 6, we get access to a unique cycle. However, in that section we needed a prime length cycle. If we are only interested in unique cycles, not necessarily of prime length, this process terminates much more quickly: *Golomb's constant* [16] measures, as  $n \rightarrow \infty$ , the probability that a random element  $\varphi \in S_n$  has a cycle of length greater than  $\frac{n}{2}$ . The limit is greater than  $\frac{4}{5}$ .

In practice the existence of a unique cycle is a mild assumption: on the one hand some practical heuristics only apply if specific combinations of transpositions are present in the generators [13, 25]. Each transposition is a unique cycle. On the other hand, randomly distributed automorphisms, as returned by TRACES and DEJAVU, satisfy having a unique cycle with high probability, as argued above.

**Unique Cycle To Cyclic Order.** We now assume we are given an invariant unique cycle  $C$ . The idea is now to extend  $C$  using the generators  $S$  to a cycle which encompasses the entire orbit. Crucially, the extension ensures that the result is still invariant (i.e. if we do this for all orbits simultaneously, the final permutation will be invariant).

<sup>1</sup> In group theoretic terms,  $\pi$  is in  $C_{\text{Sym}(\Omega)}(C_{\text{Sym}(\Omega)}(\langle s_1, \dots, s_m \rangle))$ .





■ **Figure 7** An illustration of the cycle overlap algorithm. Overlapping cycles  $C = (0, 1, 2, 3)$  and  $D = \{(0, a, 1, b, c), (2, d)\}$ .

■ **Algorithm 3** The cycle overlap algorithm.

---

```

1 function CycleOverlap( $C, D$ )
  Input : directed cycle  $C$ , overlapping pair-wise disjoint directed cycles
          $D = \{D_1, \dots, D_m\}$  (with  $\forall D_i \in D: D_i \cap C \neq \emptyset$ ,
          $\forall D_i, D_j \in D: i \neq j \implies D_i \cap D_j = \emptyset$ )
  Output: directed cycle containing all vertices  $C \cup D$ 
2  $C' := \bigcup_{D_i \in D} C \cap D_i$ ;
3 for ( $c \in C'$ )
4    $D' :=$  read  $D$  from  $c$  to next vertex of  $C'$ ;
5   insert  $D'$  after  $c$  in  $C$ ;
6 return  $C'$ ;

```

---

We now describe the *cycle overlap* algorithm, which gets as input a directed cycle  $C$ , as well as a collection of cycles  $D_1, \dots, D_m$  which must be pair-wise disjoint. Furthermore, each  $D_i$  must have one vertex in common with  $C$ . The result is a cycle  $C'$  that contains all vertices of all the cycles. A formal description can be found in Algorithm 3.

*Description of Algorithm 3.* The algorithm first checks which vertices of  $D = D_1 \cup \dots \cup D_m$  appear in  $C$ , and records them into the set  $C'$ . Then, for each  $c \in C'$  in the overlap of  $C$  and  $D$ , the algorithm walks along the respective cycle  $D_i$  containing  $c$ , and records all vertices it observes into  $D'$ . It walks along the cycle until another  $c' \in C'$  is reached (it may record the entire cycle  $D_i$ , i.e.,  $c' = c$  may hold). Finally,  $D'$  is inserted as a path into  $C$ .

The output of the process is invariant under the cyclic orders involved. This means no matter in which order the cycles from  $D$  are processed, the algorithm always returns the same cyclic order. Figure 7 illustrates the algorithm.

*Runtime of Algorithm 3.* We may use a doubly-linked list structure for directed cycles  $C$  and  $D$ , and an array  $A$  to link vertices  $V$  to their position in  $C$  in time  $\mathcal{O}(1)$ . Assuming these data structures, inserting a  $D'$  into  $C$  can be performed in time  $|D'|$ . Indeed, with these data structures, we can implement the entire algorithm in time  $\mathcal{O}(\sum_{i \in \{1, \dots, m\}} |D_i|)$ . We may also update the array  $A$  to include the new vertices of  $C$  added from  $D$ .

To get a unique cyclic order, we repeatedly combine  $C$  with cycles appearing in generators that intersect  $C$ . Every cycle in a generator only has to be processed once. Eventually  $C$  contains the entire orbit. With careful management of usage-lists of vertices in cycles of generators, the overall algorithm can be implemented in instance-linear time.

### 7.3 Cyclic Order to Equivalent Orbits

We finally describe how to find equivalent orbits, assuming invariant cyclic orders are given on the orbits. An invariant cyclic order for the vertices of each orbit moves us one step closer to the orbits of the cycle type graph. There are however still many potential bijections between orbits: indeed, we do not know how each cyclic order should be rotated. We therefore describe a procedure to refine the cyclic order further.

■ **Algorithm 4** High-level procedure to obtain equivalent orbits.

---

```

1 function EquivalentOrbits( $S, \Delta_1, \dots, \Delta_m, C_{\Delta_1}, \dots, C_{\Delta_m}$ )
   Input : group  $\langle S \rangle = \Gamma \leq S(\Omega)$ , orbit partition  $\Delta_1, \dots, \Delta_m$  of  $\Gamma$ , unique cycles
            $C_{\Delta_1}, \dots, C_{\Delta_m}$ 
   Output: partition of  $v \in \Omega$  into equivalent orbits
2   overlap each unique cycle  $C_{\Delta_i}$  with  $S$  to obtain canonical cycle  $C'_{\Delta_i}$ ;
3   construct enhanced cycle type graph  $\mathcal{C}'(S)$ ;
4    $\pi :=$  apply color refinement to  $\mathcal{C}'(S)$ ;
5   return  $\pi$ 

```

---

We introduce the *enhanced cycle type graph*  $\mathcal{C}'(S)$ . We are provided an invariant cyclic order for each orbit  $\Delta$  of  $\Gamma$ , which we denote by  $C_\Delta$ . First, we add to the cycle type graph (Subsection 7.1) a cycle gadget for each  $C_\Delta$ . As before, we color the cycle gadget  $C_\Delta$  according to its cycle length. Next, we enhance all other cycle gadgets using distance information of  $C_\Delta$ : in every cycle gadget we mark each directed edge  $v_1 \rightarrow v_2$  with the length of the (directed) path from  $v_1$  to  $v_2$  in  $C_\Delta$  (see Figure 6b). We write  $v_1 \xrightarrow{c} v_2$  whenever the path from  $v_1$  to  $v_2$  in  $C_\Delta$  has length  $c$ . Note that while we use edge-labels for clarity, these can be encoded back into vertex colors (see [19, Proof of Lemma 15]).

Just like with the cycle type graph, the automorphism group of the enhanced cycle type graph  $\mathcal{C}'(S)$  is the centralizer of  $\Gamma$  and  $(C_{\text{Sym}(\Omega)}(\Gamma), \mathcal{C}'(\Gamma))$  is a joint graph/group pair (see Lemma 9). However, it is easier to compute the orbit partition of  $\mathcal{C}'(S)$ .

In fact, our method of obtaining the orbits of  $\mathcal{C}'(S)$  is rather straightforward: we apply the color refinement procedure to the enhanced cycle type graph  $\mathcal{C}'(S)$ . The technical proof that color refinement indeed computes the orbit partition can be found in the full version [5]:

► **Lemma 11.** *Color refinement computes the orbit partition of the enhanced cycle type graph.*

Given our high-level procedure in Algorithm 4, and given that color refinement can be computed in quasi-linear time as previously discussed, leads to the following theorem:

► **Theorem 12.** *Given access to a unique cycle per orbit, there is an instance-quasi-linear algorithm which computes for a joint graph/group pair  $(G, S)$  a partition  $\pi$  of equivalent orbits. Given two equivalent orbits  $\Delta_1 \equiv \Delta_2$ , there is an algorithm which computes from  $\pi$  a corresponding matching  $b: \Delta_1 \rightarrow \Delta_2$  such that for all  $\varphi \in \Gamma$  and  $\delta \in \Delta_1$ ,  $\varphi(b(\delta)) = b(\varphi(\delta))$  in time  $\mathcal{O}(|\Delta_1|)$ .*

## 8 Conclusion and Future Work

Exploiting our concept of joint graph/group pairs, we proposed new, asymptotically faster algorithms for the SAT-symmetry interface. However, most of the new concepts and approaches of this paper do not only apply to the domain of SAT, but also for example to MIP [25] and CSP [14]. More computational tasks should be considered in this context, the most prominent one arguably being pointwise stabilizers [30].

Our new algorithms exploit subroutines with highly efficient implementations available, but otherwise do not use any complicated data structures. We intend to implement the algorithms and integrate them into the symmetry detection preprocessor SASSY [6].

Finally, in some classes of SAT instances, more complex symmetry structures may arise. Analyzing and taking advantage of these structures is potential future work. For example,

in the pigeonhole principle, BREAKID finds overlapping row interchangeability groups and breaks these groups partially. By virtue of being overlapping, the symmetry breaking constraints produced are *not* guaranteed to be complete. Specifically, the pigeonhole principle is an example of instances whose symmetries form a *wreath product* of two symmetric groups, i.e.,  $S_n \text{ wr } S_m$  ( $n$  pigeons,  $m$  holes). Procedures to detect and exploit such groups (e.g., by first using blocks of imprimitivity [30] followed by the algorithms of this paper) could be of practical interest.

---

## 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*, pages 836–839. ACM, 2003. doi:10.1145/775832.776042.
- 2 Fadi A. Aloul, Arathi Ramani, Igor L. Markov, and Kareem A. Sakallah. Solving difficult SAT instances in the presence of symmetry. In *Proceedings of the 39th Design Automation Conference, DAC*, pages 731–736. ACM, 2002. doi:10.1145/513918.514102.
- 3 Markus Anders and Pascal Schweitzer. Parallel computation of combinatorial symmetries. In Petra Mutzel, Rasmus Pagh, and Grzegorz Herman, editors, *29th Annual European Symposium on Algorithms, ESA*, volume 204 of *LIPICs*, pages 6:1–6:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.ESA.2021.6.
- 4 Markus Anders and Pascal Schweitzer. Search problems in trees with symmetries: Near optimal traversal strategies for individualization-refinement algorithms. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP*, volume 198 of *LIPICs*, pages 16:1–16:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.ICALP.2021.16.
- 5 Markus Anders, Pascal Schweitzer, and Mate Soos. Algorithms transcending the sat-symmetry interface. *CoRR*, abs/2306.00613, 2023. doi:10.48550/arXiv.2306.00613.
- 6 Markus Anders, Pascal Schweitzer, and Julian Stieß. Engineering a preprocessor for symmetry detection. *CoRR*, abs/2302.06351, 2023. doi:10.48550/arXiv.2302.06351.
- 7 Paul Beame, Richard M. Karp, Toniann Pitassi, and Michael E. Saks. The efficiency of resolution and davis–putnam procedures. *SIAM J. Comput.*, 31(4):1048–1075, 2002. doi:10.1137/S0097539700369156.
- 8 Bart Bogaerts, Stephan Gocht, Ciaran McCreesh, and Jakob Nordström. Certified symmetry and dominance breaking for combinatorial optimisation. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI*, pages 3698–3707. AAAI Press, 2022. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/20283>.
- 9 Mun See Chang and Christopher Jefferson. Disjoint direct product decompositions of permutation groups. *J. Symb. Comput.*, 108:1–16, 2022. doi:10.1016/j.jsc.2021.04.003.
- 10 James M. Crawford, Matthew L. Ginsberg, Eugene M. Luks, and Amitabha Roy. Symmetry-breaking predicates for search problems. In Luigia Carlucci Aiello, Jon Doyle, and Stuart C. Shapiro, editors, *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning (KR’96)*, pages 148–159. Morgan Kaufmann, 1996.
- 11 Paul T. Darga, Mark H. Liffiton, Kareem A. Sakallah, and Igor L. Markov. Exploiting structure in symmetry detection for CNF. In Sharad Malik, Limor Fix, and Andrew B. Kahng, editors, *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.
- 12 Jo Devriendt, Bart Bogaerts, and Maurice Bruynooghe. Symmetric explanation learning: Effective dynamic symmetry handling for SAT. In Serge Gaspers and Toby Walsh, editors, *Theory and Applications of Satisfiability Testing – SAT*, volume 10491 of *LNCS*, pages 83–100. Springer, 2017. doi:10.1007/978-3-319-66263-3\_6.
- 13 Jo Devriendt, Bart Bogaerts, Maurice Bruynooghe, and Marc Denecker. Improved static symmetry breaking for SAT. In Nadia Creignou and Daniel Le Berre, editors, *Theory and*

- Applications of Satisfiability Testing – SAT*, volume 9710 of *LNCS*, pages 104–122. Springer, 2016. doi:10.1007/978-3-319-40970-2\_8.
- 14 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 Pascal Van Hentenryck, editor, *Principles and Practice of Constraint Programming – CP*, volume 2470 of *LNCS*, pages 462–476. Springer, 2002. doi:10.1007/3-540-46135-3\_31.
  - 15 The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.12.2*, 2022. URL: <https://www.gap-system.org>.
  - 16 Solomon W. Golomb and Peter Gaal. On the number of permutations on  $n$  objects with greatest cycle length  $k$ . *Advances in Applied Mathematics*, 20(1):98–107, 1998. doi:10.1006/aama.1997.0567.
  - 17 Andrew Grayland, Christopher Jefferson, Ian Miguel, and Colva M. Roney-Dougal. Minimal ordering constraints for some families of variable symmetries. *Annals of Mathematics and Artificial Intelligence*, 57:75–102, 2009.
  - 18 Tommi A. Junttila and Petteri Kaski. Conflict propagation and component recursion for canonical labeling. In Alberto Marchetti-Spaccamela and Michael Segal, editors, *Theory and Practice of Algorithms in (Computer) Systems – First International ICST Conference, TAPAS*, volume 6595 of *LNCS*, pages 151–162. Springer, 2011. doi:10.1007/978-3-642-19754-3\_16.
  - 19 Sandra Kiefer, Iliia Ponomarenko, and Pascal Schweitzer. The weisfeiler-leman dimension of planar graphs is at most 3. *J. ACM*, 66(6):44:1–44:31, 2019. doi:10.1145/3333003.
  - 20 Sandra Kiefer, Pascal Schweitzer, and Erkal Selman. Graphs identified by logics with counting. In Giuseppe F. Italiano, Giovanni Pighizzini, and Donald Sannella, editors, *Mathematical Foundations of Computer Science 2015 – 40th International Symposium, MFCS, Part I*, volume 9234 of *LNCS*, pages 319–330. Springer, 2015. doi:10.1007/978-3-662-48057-1\_25.
  - 21 Markus Kirchweger and Stefan Szeider. SAT modulo symmetries for graph generation. In Laurent D. Michel, editor, *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.
  - 22 Brendan D. McKay. Practical graph isomorphism. In *10th. Manitoba Conference on Numerical Mathematics and Computing (Winnipeg, 1980)*, pages 45–87, 1981.
  - 23 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.
  - 24 M. Neunhöffer, Á. Seress, and M. Horn. *recog*, a package for constructive recognition of permutation and matrix groups, Version 1.4.2. <https://gap-packages.github.io/recog>, September 2022. GAP package.
  - 25 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.
  - 26 Adolfo Piperno. Isomorphism test for digraphs with weighted edges. In Gianlorenzo D’Angelo, editor, *17th International Symposium on Experimental Algorithms, SEA*, volume 103 of *LIPICs*, pages 30:1–30:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPICs.SEA.2018.30.
  - 27 Jean-Francois Puget. Symmetry breaking using stabilizers. In Francesca Rossi, editor, *Principles and Practice of Constraint Programming – CP*, volume 2833 of *LNCS*, pages 585–599. Springer, 2003. doi:10.1007/978-3-540-45193-8\_40.
  - 28 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.
  - 29 Karem A. Sakallah. Symmetry and satisfiability. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability – Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, pages 509–570. IOS Press, 2021. doi:10.3233/FAIA200996.

- 30 Ákos Seress. *Permutation Group Algorithms*. Cambridge Tracts in Mathematics. Cambridge University Press, 2003. doi:10.1017/CB09780511546549.