



A computational comparison of symmetry handling methods for mixed integer programs

Marc E. Pfetsch¹ · Thomas Rehn²

Received: 21 November 2015 / Accepted: 7 June 2018 / Published online: 4 July 2018
© Springer-Verlag GmbH Germany, part of Springer Nature and The Mathematical Programming Society 2018

Abstract

The handling of symmetries in mixed integer programs in order to speed up the solution process of branch-and-cut solvers has recently received significant attention, both in theory and practice. This paper compares different methods for handling symmetries using a common implementation framework. We start by investigating the computation of symmetries and analyze the symmetries present in the MIPLIB 2010 instances. It turns out that many instances are affected by symmetry and most symmetry groups contain full symmetric groups as factors. We then present (variants of) six symmetry handling methods from the literature. Their implementation is tested on several testsets. On very symmetric instances used previously in the literature, it is essential to use methods like isomorphism pruning, orbital fixing, or orbital branching. Moreover, tests on the MIPLIB instances show that isomorphism pruning, orbital fixing, or adding symmetry breaking inequalities allow to speed-up the solution process by about 15% and more instances can be solved within the time limit.

Keywords Symmetry · Mixed integer program · Branch-and-cut · Isomorphism pruning · Orbital branching

Mathematics Subject Classification 90C11 · 90C57

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s12532-018-0140-y>) contains supplementary material, which is available to authorized users.

✉ Marc E. Pfetsch
pfetsch@mathematik.tu-darmstadt.de
Thomas Rehn
research@carmen76.de

¹ Fachbereich Mathematik, Technische Universität Darmstadt, Dolivostrasse 15, 64293 Darmstadt, Germany

² initOS GmbH, Hegelstrasse 28, 39104 Magdeburg, Germany

Symmetries in mixed integer programs (MIPs) have been known to slow down branch-and-cut algorithms for a long time. The reason is that possibly many symmetric solutions appear in the branch-and-bound tree, although a single representative contains all essential information. Depending on the symmetry group, this significantly blows up the size of the branch-and-bound tree. During recent years, several approaches to deal with this problem have been developed. Historically the first approach is to add symmetry breaking inequalities or to perturb the objective function, which are folklore knowledge, see Sherahli and Smith [49] and Ghoniem and Sherahli [16] for recent examples. The handling of symmetries in the tree received more attention in mathematical programming beginning with “isomorphism pruning” by Margot [29–32]. Subsequent developments are, for example, “orbital branching” by Ostrowski et al. [36,38] and generalized investigations by Liberti [27]. We refer to Margot [33] for a nice comprehensive overview.

Moreover, symmetry handling techniques have found their way into commercial solvers, like CPLEX, GUROBI, and XPRESS. Although the approaches used in these solvers are not documented, it seems likely that they are based on techniques similar to orbital fixing, often combined with specialized methods to handle the symmetry groups.

Nevertheless, as far as we know, there is no computational comparison of the different techniques and their strengths and weaknesses. One goal of this paper is to close this gap. We present implementations of several techniques available in the literature using the branch-and-cut framework SCIP [2,47]. The implementation is publicly available through the web page of the first author and rests on `PermLib` [43], a C++-library of the second author, for the necessary computations with permutation groups. (We note that the implementation for isomorphism pruning of Margot was available on his web page, as well.) Additionally, we present several modifications and some new contributions for existing symmetry handling methods.

Our paper is structured as follows. In Sect. 1, we define symmetries of MIPs and introduce necessary notation. Section 1.1 deals with the computation of symmetry groups. It seems to be folklore knowledge that this can be reduced to the computation of graph automorphisms, and we discuss two ways to construct the corresponding graphs. In Sect. 2, we introduce the implemented symmetry handling methods, most of which are available in the literature:

- Orbital Fixing (Sect. 2.1),
- Isomorphism Pruning (Sect. 2.2),
- Orbital Branching (Sect. 2.3),
- Orbit Probing (Sect. 2.5),
- Symmetry Breaking Inequalities (Sect. 2.6),
- Projection onto the Fixed Space (Sect. 2.7).

Note that this list contains large parts of the proposed methods from the literature. However, it cannot be comprehensive. For example, methods to improve primal heuristics as in Christophel et al. [9] have not been implemented.

The heart of this paper are the extensive computational experiments in Sect. 5. We compare the presented approaches on large testsets, including the instances of Margot, MIPLIB 2003, and MIPLIB 2010.

These experiments have two goals: The first goal is to evaluate the potential of symmetry handling methods. To this end, we analyze the symmetry groups of the MIPLIB 2010 testset, see Sect. 4 and the “Appendix”. It turns out that symmetry appears in (surprisingly) many instances. Of course, the sizes of the symmetry groups and their degree vary from instance to instance. In any case, there are instances in this testset with a very high degree of symmetry, i.e., the sizes of the symmetry groups are huge and they affect many variables of the instances. Moreover, it turns out that most factors of the symmetry groups are symmetric groups. We conclude that for instances of the MIPLIB 2010 (or similar), existing algorithms that are simplified and tuned for the special case of symmetric groups will most probably be very useful for an efficient solving. On the other hand, there are special (combinatorial) instances like the Margot instances with large and very complicated symmetries. For such instances, any solver that does not handle symmetries will not be successful and symmetry handling will require more advanced techniques. It therefore might be that specialized techniques for such specially structured groups will appear in the future.

The second goal is to compare the performance of the different symmetry handling approaches on different testsets. It turns out that on the instances of Margot, with their relatively complicated symmetry groups, isomorphism pruning, orbital fixing, and orbital branching are the fastest; without these methods the performance decreases dramatically. For instances from the MIPLIB, isomorphism pruning and orbital fixing allow a speed-up of around 14% and solve more instances. Moreover, the addition of symmetry breaking inequalities performs very well with a speed-up of about 15%. This suggests that more research is needed on such inequalities and their ability to improve the dual bounds. For more conclusions from these experiments, see Sect. 6.

1 Symmetries of mixed integer optimization problems

Consider a *mixed integer program* (MIP) in the following form:

$$\begin{aligned} \max \quad & c^\top x \\ & Ax \leq b \\ & x \in \mathbb{Z}^p \times \mathbb{R}^{n-p}, \end{aligned} \tag{1}$$

where $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $c \in \mathbb{R}^n$, and $0 \leq p \leq n$ denotes the number of integer variables. We write $P := \{x \in \mathbb{R}^n : Ax \leq b\}$ for the polyhedron and $X := P \cap (\mathbb{Z}^p \times \mathbb{R}^{n-p})$ for the feasible region corresponding to the MIP (1).

A *symmetry* of (1) is a bijection $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that $f(X) = X$ and $c^\top f(x) = c^\top x$ for every $x \in X$. Thus, f maps feasible solutions of (1) to feasible solutions with the same objective. The symmetries of (1) form the so-called *symmetry group*. However, the definition of a symmetry explicitly is based on the feasible region X , which is in general not efficiently handable (it is NP-hard to decide whether $X = \emptyset$).

In practice, one therefore often only considers permutations of variables that leave the description $Ax \leq b$ and the objective function c invariant, see, e.g., Margot [33]. Here, the *symmetric group*, i.e., the set of all permutations on $\{1, \dots, n\}$, is denoted

by \mathcal{S}_n , and a permutation $\pi \in \mathcal{S}_n$ acts on \mathbb{R}^n by permuting the components, i.e., $\pi(x) := (x_{\pi^{-1}(1)}, \dots, x_{\pi^{-1}(n)})^\top$ for $x \in \mathbb{R}^n$.

Definition 1 A permutation $\pi \in \mathcal{S}_n$ of variables is a *formulation symmetry* of (1) if there exists a permutation $\sigma \in \mathcal{S}_m$ such that

- $\pi(\{1, \dots, p\}) = \{1, \dots, p\}$ (i.e., π preserves integer variables),
- $\pi(c) = c$,
- $\sigma(b) = b$, and
- $A_{\sigma(i), \pi(j)} = A_{ij}$.

Thus, the rows of A are permuted by σ and the columns by π .

Remark 2 The above definition could be generalized in several different ways. First, one can consider permutations that leave P (instead of its formulation), the integer variables, and the objective invariant, see, e.g., Bödi et al. [7]. Such symmetries can in principle be computed by normalization of the description of P and using graph automorphisms, see Herr [17]. This would avoid the effect that redundant inequalities might change the symmetries of Definition 1. To some extent, one can take care of this aspect by computing the symmetries after preprocessing, which performs a normalization and removes some redundant constraints.

A further generalization can be obtained by considering general linear bijection f with $f(P) = P$ that also preserve integrality of integer coordinates and the objective function. Thus, one is interested in maps from $\text{GL}_p(\mathbb{Z}) \times \text{GL}_{n-p}(\mathbb{R})$ that preserve P and do not change the value of the objective function, see Bremner et al. [8], Bödi et al. [7], Herr [17], and Padberg [39]. If all variables are integral, one extension is to consider subgroups of $O_n(\mathbb{Z})$, so-called signed permutations, see [7]. However, systematically computing $\text{GL}_p(\mathbb{Z})$ -symmetries of polyhedra is difficult, see [8]. Moreover, two different polyhedra yield the same set X , but have different symmetries because the feasible region is only a subset of the polyhedra.

As it is customary in many articles in mixed integer programming, we will exclusively deal with formulation symmetries as in Definition 1 and therefore simply refer to them as *symmetries* of (1) from now on. The group of all symmetries is the *symmetry group* G of (1) and is a subgroup of \mathcal{S}_n , which we denote by $G \leq \mathcal{S}_n$. The symmetry group can be (efficiently) computed in practice and can be represented via generators. We will discuss this in the next section.

1.1 Computing symmetry groups

Computing the symmetry group of a MIP is usually reduced to the determination of graph automorphisms, see Margot [33]. While the first (correct) approach seems to be due to Salvagnin [44], many authors have (re)discovered this fact, see, e.g., [6], Liberti [27], and Bödi et al. [7]. In the constraint programming literature, a similar idea already appears in Puget [40].

The computational complexity of the graph automorphisms problem is still unknown (it is neither known to be NP-hard, nor known to be solvable in polynomial time), see, e.g., Read and Corneil [41] and Johnson [20]. There are, however,

several software tools which compute graph automorphisms efficiently in practice even for large graphs, e.g., `nauty` [34], `saucy` [10], and `bliss` [21].

A natural way to model MIP symmetries as graph automorphisms is via the following bipartite graph $(V \dot{\cup} V', E)$ with vertex and edge colors. The set $V = \{v_1, \dots, v_n\}$ contains a vertex v_j for each variable x_j of the problem; v_j is colored according to the objective coefficient c_j of variable x_j . The second set $V' = \{v'_1, \dots, v'_m\}$ contains a vertex for each linear inequality in $Ax \leq b$. Each vertex v'_i is colored with respect to the coefficient b_i of the right-hand side. There is an edge $\{v'_i, v_j\} \in E$ if $A_{ij} \neq 0$, and it is colored by the coefficient A_{ij} of variable x_j in the i -th constraint. Moreover, vertices v_1, \dots, v_p , corresponding to integer variables, receive distinct colors from v_{p+1}, \dots, v_n . A graph automorphism for this bipartite graph is now given by two bijections $\pi : V \rightarrow V$, $\sigma : V' \rightarrow V'$ such that $\{\sigma(v'), \pi(v)\} \in E$ if and only if $\{v', v\} \in E$. It is easy to see that a graph automorphism that respects the vertex and edge colors corresponds to a symmetry according to Definition 1 and conversely.

Note that 0-coefficients play a special role in this construction. Of course, we can replace 0 by any other number. However, the matrices appearing in MIPs are often sparse. In this case, the choice of 0 as special coefficient reduces the number of edges in the graph and speeds up the symmetry computation.

The above mentioned graph automorphism software packages can only handle vertex colors, but not edge colors. In fact, edge colors are not needed if A contains only one coefficient different from 0, e.g., if A is a 0/1-matrix. In the other cases, one can reduce the problem to a purely vertex-colored instance by applying two techniques that we describe in the following.

Salvagnin [44] discusses a transformation in which every edge $\{v', v\} \in E$ is replaced by two edges $\{v', w\}$, $\{w, v\}$, using an intermediate vertex w that is colored with the original edge color.

Additionally, the number of newly introduced vertices can often be substantially reduced by an idea of Puget [40]: Instead of adding new intermediate vertices for all edges, vertices with the same color can be combined. For each $v'_i \in V'$ let $V_{i,c} \subseteq V$ be the set of vertices which are incident to v'_i with an edge of color c . Then it is enough to introduce one intermediate c -colored vertex w with edges to v'_i and to all elements of $V_{i,c}$. We call this construction *grouping by variables*. In many MIP-instances, each constraint contains only few distinct variable coefficients. In this case, the sets $V_{i,c}$ are large, and the number of vertices in the graph is significantly reduced.

Depending on the distribution of coefficients, it may be beneficial to swap the roles of constraints and variables to add as few intermediate vertices as possible. For instance, if there are much more constraints than variables, it may help to add one intermediate vertex between each v_i and the set $V'_{i,c} \subseteq V'$ of all “constraint” vertices connected to v_i by an edge of color c . We call this construction *grouping by constraints*. Because our original graph is bipartite, both groupings are possible. In the following we will refer to either of these constructions as *matrix graph*; see Fig. 1 for an example.

The second, fundamentally different transformation, is described in the manual of the software `nauty` (version 2.4). Since it does not depend on bipartiteness, we describe it for a general edge-colored graph (V, E) . Let C be the total number of distinct edge colors in this graph, and let $L = \lceil \log_2(C + 1) \rceil$ be the number of bits needed to represent C . We introduce new vertex colors C_1, \dots, C_L and replace each




Fig. 1 Example for the reduction of symmetry computation with respect to (A, b, c) to graph automorphisms

vertex $v \in V$ by vertices v_1, \dots, v_L that are colored with C_1, \dots, C_L , respectively. Additionally, we add edges $\{v_1, v_2\}, \{v_2, v_3\}, \dots, \{v_{L-1}, v_L\}$. For each edge $\{v, w\} \in E$ with color $c \in \{1, \dots, C\}$ of the original graph, we add edges between v_i and w_i for each i -th bit that is 1 in the binary representation of c . Thus, we emulate the edge colors by vertex bit colors. Note that this transformation is described for a complete input graph. If the input graph is not complete, we can augment it to a complete graph without loss of generality by adding the missing edges with a new color which is not used in the input graph. We will refer to this construction as a whole as *layered graph*; see Fig. 1 for an example.

The matrix graph has $m + n + O(N)$ vertices, where N is the number of nonzeros in A . Depending on the distribution of different coefficients in the constraint matrix, the last part may be much smaller than N . The layered graph results in about $(n + m) \log_2(C)$ vertices, where C is the total number of distinct coefficients in the constraint matrix. Depending on the instance, either transformation might lead to a smaller graph. Thus, no construction dominates the other, in general. We report on experience with graph sizes for practical instances in Sect. 4.

In any symmetry computation it is important to take numerical issues into account. For a very simple example, consider three matrix coefficients α, β, γ , such that α and β as well as β and γ are numerically equal (e.g., $|\alpha - \beta| \leq \varepsilon$ and $|\beta - \gamma| \leq \varepsilon$, for some tolerance ε), but α and γ are not equal. Thus, the loss of transitivity has to be taken into account in order to avoid wrong symmetry computations. One method, also used by Salvagnin (2012, private communication), is to first sort all coefficients, say of A , non-decreasingly. Then passing through the sorted list, a new color for A_{ij} is used whenever $|A_{ij} - \delta| > \varepsilon$, where δ is the minimal coefficient belonging to the last used color. In this way, we have a stable behavior, but might consider two coefficients as different that are still numerically equal. Thus, we might compute a subgroup of the “real” symmetry group.

1.2 Permutation groups

As mentioned earlier, the symmetry group of a MIP is a permutation group, i.e., a subgroup of the (full) symmetric group S_n . In this section, we introduce basic notions and concepts of permutation groups, which help to classify different symmetry types (see the list of symmetry groups in Table 19).

The *degree* of a permutation group is the number of moved points. If a permutation group G is a direct product $G = H_1 \times \dots \times H_k$ of some permutation groups H_1, \dots, H_k , it suffices to consider each factor H_i independently, since the groups H_1, \dots, H_k act on disjoint sets of elements. In the following, we consider single factors only.

As we will see later, for some symmetry handling techniques it is not enough to know the abstract isomorphism type of a group, the action on the elements is important as well. The most important case is that of groups isomorphic to symmetric groups, which we will describe next.

Let $G \leq S_n$ be a permutation group of degree n , and let $G \cong S_m$ be isomorphic to a symmetric group of degree $m \leq n$. If $m = n$, we call the action of G a *coordinate action* (that is, G is a natural representation of S_n). In the case $m < n$, there may exist a group $H \leq S_m$ such that G is isomorphic to the particular form $\{(\pi, \pi) : \pi \in H\}$; this of course means that $H \cong G \cong S_m$. We say that such a group G has a *matrix action* (or is a *diagonal group*). Similar definitions for coordinate and matrix actions can also be made for cyclic groups. Coordinate and matrix actions of symmetric groups are especially nice from a computational perspective. Their corresponding orbits, stabilizers, and minimal orbit elements can be computed very fast in these representations.

Example 3 Some computations show that the groups

$$\begin{aligned} G_1 &= \langle (1\ 2\ 3\ 4\ 5\ 6), (1\ 2) \rangle \leq S_6, \\ G_2 &= \langle (1\ 2\ 3\ 4\ 5\ 6)(7\ 8\ 9\ 10\ 11\ 12), (1\ 2)(7\ 8) \rangle \leq S_{12}, \quad \text{and} \\ G_3 &= \langle (1\ 4\ 8\ 6\ 3\ 10)(2\ 7\ 9), (1\ 5\ 3\ 4\ 7)(2\ 10\ 6\ 8\ 9) \rangle \leq S_{10} \end{aligned}$$

are all isomorphic to the symmetric group S_6 on six elements. Here, we use the standard cycle notation (i_1, i_2, \dots, i_k) to denote a permutation $\pi \in S_n$ with $\pi(i_1) = i_2, \pi(i_2) = i_3, \dots, \pi(i_k) = i_1$, and $\pi(i) = i$ for all other indices $i \in \{1, \dots, n\} \setminus \{i_1, \dots, i_k\}$. Moreover, we denote by $\langle \pi_1, \dots, \pi_\ell \rangle$ the (permutation) group generated by permutations π_1, \dots, π_ℓ .

Nevertheless, as symmetry groups of integer programs, G_1, G_2 , and G_3 behave differently. The group G_1 is a coordinate action. Group G_2 , which has degree 12, is a matrix action and corresponds to permutations of the columns of a 2×6 matrix, where the elements of the ground set correspond to the entries of the matrix as follows:

$$A = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 7 & 8 & 9 & 10 & 11 & 12 \end{pmatrix}.$$

For the action of groups like G_3 there is in general no obvious and nice interpretation. Such groups appear, for instance, as symmetry groups of many instances from the testset `Margot1`, see Sect. 5.

2 Exploiting symmetries in MIPs

In this section, we review some well-known and some less-known techniques to handle symmetries in MIPs, as well as new variants of these. Almost all approaches are presented for the mixed 0/1 case, i.e., all integral variables x_1, \dots, x_p are restricted to have values 0 or 1. The few exceptions that work for general MIPs are mentioned explicitly.

One important issue that we also treat in this section is the compatibility of the symmetry breaking methods with other solver components. Clearly, using two symmetry breaking methods that are not based on the same principle can lead to cutting off all optimal solutions and thus to wrong results. We will therefore be careful not to mix different symmetry breaking methods.

An important method in MIP-solvers is *bound propagation*, i.e., based on the current bounds of variables, further bounds of variables can be strengthened. In the 0/1-case, bound propagations yield *fixings* of variables to 0 or 1 in the current node (see below for examples). We distinguish this from *branching* on variables, i.e., new nodes in the tree have been created by changing the bounds of a single variable. Note that this notation deviates from parts of the literature, e.g., the articles by Margot [30,33] and Ostrowski et al. [36,38], where fixings are called settings and branchings are called fixings.

It is important to note that the dynamic techniques explained in the following are not correct anymore if branching is not performed on variables, e.g., by branching on constraints. Moreover, components performing bound propagations/fixings have to be *symmetry independent* (this is called “strict setting algorithm working under symmetry” by Margot [30]), i.e., they have to respect symmetry in the following sense: if the method is able to fix variable x_i , it is (in principle) also able to fix variable $x_{\pi(i)}$ for all symmetries π . Of course, it may happen in practice that not all fixings have been carried out, because of performance (e.g., iteration) limits or numerical considerations in the implementation.

Most components of current MIP-solvers are symmetry independent, in particular, bound strengthening of linear constraints, dual reductions, reduced cost strengthening etc. One symmetry independent operation that is typically subject to performance limits is probing (see, e.g. Savelsbergh [46] and Achterberg [1]). Furthermore, cutting planes that are valid for the convex hull of all feasible solutions are symmetry independent, i.e., can be applied without conflict with symmetry handling methods. This includes all basic cutting planes like Gomory (mixed integer) cuts or knapsack based cuts. We refer the reader to Margot [33, Section 9.1.1] for a more elaborate discussion.

In the following, let $V := \{1, \dots, n\}$ be the indices of all variables and $B := \{1, \dots, p\} \subseteq V$ be the indices of all binary variables. Then $\text{Stab}(G, S)$ denotes the (set-wise) *stabilizer* of the symmetry group G with respect to a subset $S \subseteq V$ of variable indices, i.e.,

Algorithm 1: Orbital Fixing

Input: G : symmetry group;
 B_0, B_1 : binary variables branched on 0/1 resp.;
 F_0, F_1 : binary variables fixed to 0/1 resp.
Output: L_0, L_1 : binary variables that can be fixed to 0/1 resp.

- 1 compute $\text{Stab}(G, B_1)$ and its orbits \mathcal{O} ;
- 2 remove $O \in \mathcal{O}$ if $O \cap B_1 \neq \emptyset$ or $O \cap (V \setminus B) \neq \emptyset$;
- 3 $L_0 \leftarrow \emptyset, L_1 \leftarrow \emptyset$;
- 4 **foreach** orbit $O \in \mathcal{O}$ with $|O| \geq 2$ **do**
- 5 **if** $O \cap (B_0 \cup F_0) \neq \emptyset$ **then**
- 6 $L_0 \leftarrow L_0 \cup (O \setminus (B_0 \cup F_0))$;
- 7 **if** $O \cap F_1 \neq \emptyset$ **then**
- 8 $L_1 \leftarrow L_1 \cup (O \setminus F_1)$;

$$\text{Stab}(G, S) := \{\pi \in G : \pi(S) = S\}.$$

Moreover, the orbit of variable index $i \in V$ with respect to G is $O(i) := \{\pi(i) : \pi \in G\}$. It is well known that the orbits partition the ground set V . For more information about group theory, see, e.g., Lang [25].

Moreover, for the current node in a branch-and-bound tree, let $B_0, B_1 \subseteq B$ be the indices of binary variables that have been branched on 0 and 1, respectively, and let $F_0, F_1 \subseteq B$ be the indices of binary variables fixed to 0 and 1, respectively, by some symmetry independent method or a single symmetry handling method. We will assume that the sets B_0, B_1, F_0 , and F_1 are pairwise distinct; in particular, the current node is not trivially infeasible.

2.1 Orbital fixing

A basic method to exploit symmetries in order to fix variables is *orbital fixing*, which was introduced by Margot [30, Sect. 4] and is described in Algorithm 1.

In order to argue the correctness of orbital fixing, we first note that Step 1 of Algorithm 1 computes the orbital partition with respect to $\text{Stab}(G, B_1)$. Thus, every such orbit that contains a variable in B_1 only contains variables in B_1 . Moreover, the variable types in each orbit are the same. Thus, orbits containing a non-binary variable can be ignored for orbital fixing, which explains Step 2.

We cite the following result collected from the literature:

Lemma 4 *Let F_0 and F_1 be obtained by symmetry independent methods, and let O be an orbit of $\text{Stab}(G, B_1)$.*

- *If $O \cap (B_0 \cup F_0) \neq \emptyset$ then all variables in $O \setminus (B_0 \cup F_0)$ can be fixed to 0.*
- *If $O \cap F_1 \neq \emptyset$ then all variables in $O \setminus F_1$ can be fixed to 1.*

This result was first obtained by Margot [30, Corollary 2] for a ranked branching rule. The restriction on the particular branching rule can be removed, as observed by Ostrowski [37]. For the first part, see also Ostrowski et al. [38, Theorem 3]. By this lemma, Steps 6 and 8 are correct.

Algorithm 2: Isomorphism Pruning

Input: G : symmetry group;
 B_0, B_1 : binary variables branched on 0/1 resp.;
 F_0, F_1 : binary variables fixed to 0/1 resp.
Output: L_0, L_1 : binary variables that can be fixed to 0/1 resp. or “cut off”

- 1 **if** $B_1 = \emptyset$ **then**
- 2 | compute sets L_0 and L_1 by orbital fixing w.r.t. B_0, B_1, F_0, F_1 ;
- 3 **else**
- 4 | **if** B_1 is not a lexicographic representative **then**
- 5 | | cut off current node;
- 6 | compute sets L_0 and L_1 by orbital fixing w.r.t. B_0, B_1, F_0, F_1 ;

Note that since $O \cap B_1 = \emptyset$ after Step 2, the variables in B_1 are not relevant in the second case. Moreover, by skipping orbits O that contain continuous variables, orbital fixing can be applied to mixed 0/1 problems (see Remark 5 below).

2.2 Isomorphism pruning

Isomorphism pruning was introduced by Margot [29]. It modifies the search tree such that it contains exactly one element from every orbit of optimal solutions of the original problem. In the following, we consider a generalization for binary MIPs as described by Ostrowski [37], which allows to use isomorphism pruning without restrictions on the branching order.

Consider a node in the tree other than the root node, and let $D := (i_1, \dots, i_d)$ be the sequence of binary variable indices that have been branched on, from the root to the current node; thus, $d \geq 1$ is the depth of the node. Let $B_1 \subseteq D$ be the subset of variable indices that have been branched to 1. Let $\pi \in \mathcal{S}_n$ be an arbitrary permutation with $\pi(k) = i_k$ for $k = 1, \dots, d$. Then we can prune the current node, if the set $\{k \in \{1, \dots, d\} : i_k \in B_1\}$ is not lexicographically minimal in its orbit under the conjugated symmetry group $G^\pi := \{\pi^{-1}g\pi : g \in G\}$. The argument is that in this case a branching set symmetric to B_1 is used in some other node, which will lead to symmetric solutions. It thus suffices to consider only one such node, see Margot [33] and Ostrowski [37] for details. The combination of isomorphism pruning with orbital fixing yields Algorithm 2.

Note that we use the conjugated group G^π to connect the order of branching choices (which can be represented as a word in $\{0, 1\}^d$ and for which we want to determine whether it is canonical) to the domain of the symmetry group (the corresponding variable indexes in $\{1, \dots, d\}$). If we use the first index branching rule, meaning $i_j = j$ for all branching variables (see [30]), conjugation is not necessary because there is no difference between branching choices and variables.

Remark 5 Note that the presence of continuous variables does not influence the above arguments: If we set the indices of continuous variables to be last in the ordering on which the lexicographic comparison is based, they do not change the behavior of isomorphism pruning. Because continuous variables might be moved by the symmetries used for checking the lexicographical order, we have to guarantee, however, that

Algorithm 3: Orbital Branching

Input: G : symmetry group;

B_0, B_1 : binary variables branched on 0/1 resp.;

F_0, F_1 : binary variables fixed to 0/1 resp.

Output: two new branching nodes if possible

- 1 compute $\text{Stab}(G, B_1)$ and its orbits \mathcal{O} ;
 - 2 compute sets L_0 and L_1 by orbital fixing w.r.t. B_0, B_1, F_0, F_1 ;
 - 3 choose $O \in \mathcal{O}$ with $|O| \geq 2, O \cap (V \setminus B) = \emptyset, O \cap L_0 = \emptyset, O \cap L_1 = \emptyset$;
 - 4 return if no such O exists;
 - 5 create two branching nodes: one with variables $B_1 \cup \{\min O\}$ branched to 1 and one with variables $B_0 \cup O$ branched to 0;
 - 6 in both branches fix variables in L_0 and L_1 to 0 and 1, respectively;
-

no solutions are cut off by using symmetries on continuous variables, i.e., all other solver components have to be symmetry independent. This allows to use isomorphism pruning for mixed 0/1-problems.

Remark 6 General integer variables can be handled in several ways. First, Margot [32] generalized isomorphism pruning to general integer programs. Second, if we make sure that the integer variables are branched on after all binary variables are branched on, the above argument for continuous variables holds for general integral variables as well. This would, however, restrict the applicable branching rules. Third, one can use the subgroup of G that only moves binary and continuous variables, which is done in our implementation.

The key for a successful implementation is a representation of the symmetry group that allows to efficiently decide the lexicographic test above. Margot [29] describes how this can be done using bases and strong generating sets.

2.3 Orbital branching

Orbital branching was introduced by Ostrowski et al. [36,38]. The basic idea is the following: For each orbit O of the symmetry group containing only binary variables, the following disjunction is valid:

$$\sum_{j \in O} x_j \geq 1 \quad \vee \quad \sum_{j \in O} x_j \leq 0.$$

Because of symmetry this is equivalent to:

$$x_j = 1 \quad \text{for } j = \min O \quad \vee \quad x_j = 0 \quad \forall j \in O.$$

In other words, we can branch on an orbit by setting all variables to 0 in one child node and setting the variable with smallest index in O to 1 in the second child node. Ostrowski et al. proved that this results in a valid branching rule, see [38, Thm. 1 and 2]. Moreover, orbital fixing can also be included, see [38, Thm. 3].

Let G be the “global” symmetry group of the MIP (corresponding to the root node). Then the stabilizer $\text{Stab}(G, B_1)$ is a subgroup of the “local” symmetry group H_1 , obtained by fixing all variables in B_1 to 1 and recomputing the symmetry group, see [38, Thm. 4]. Moreover, H_1 is a subgroup of the local symmetry group H obtained by fixing all variables in B_0 and B_1 to 0 and 1, respectively, and recomputing the symmetry group, see [38, Thm. 5]. Furthermore, the latter result also holds if orbital fixing is performed. Following [38], in our implementation we perform the fixings already when creating the nodes. These arguments show that Algorithm 3 is correct.

There are two degrees of freedom in the algorithm: The choice of the orbit and of the symmetry group. For the symmetry group, we may either use the “global” symmetry group $\text{Stab}(G, B_1)$ or the “local” symmetry group H , which may be larger than the global symmetry group. Following the experimental results of [38], we use $\text{Stab}(G, B_1)$.

For the choice of the orbit to branch on, several rules were described and tested in [38]; we only describe four rules here. For this we denote by $G(O) := \text{Stab}(G, B_1 \cup \{\min O\})$ the symmetry group after 1-branching on O .

- *Branch largest*: choose orbit O for which $|O|$ is maximal (Rule 1);
- *Branch break symmetry*: choose orbit O for which $|G(O)|$ is minimal (Rule 4);
- *Branch keep symmetry*: choose orbit O for which $|G(O)|$ is maximal (Rule 5);
- *Branch max product*: choose orbit O for which

$$|O| \cdot \max\{|O'| : O' \text{ orbit of } G(O)\}$$

is maximal (Rule 6).

We evaluate these different settings in our computational experiments.

Remark 7 Note that again the presence of continuous variables does not influence the validity of orbital branching as long as no solutions are cut off by using symmetries on continuous variables. Moreover, orbital branching can in principle be extended to general bounded MIPs by an encoding with binary variables, but it is then questionable to yield good performance.

2.4 Truncation techniques

Because group computations can be expensive, we might save time by skipping dynamic symmetry exploiting techniques in the deeper parts of the branch-and-bound tree. Besides using a static limit on the tree depth, another possibility is to stop symmetry handling for all children of a node in which the computed symmetry group is trivial. As the following example shows, this is only a heuristic, because symmetries might reappear in deeper levels of the tree.

Example 8 Consider the following MIP:

$$\begin{aligned} x_1 + 2x_2 + 3x_3 + 4x_4 &\leq 5, \\ 2x_1 + 3x_2 + 4x_3 + x_4 &\leq 5, \end{aligned}$$

Algorithm 4: Orbit Probing

Input: G : symmetry group
Output: L_0, L_1 : binary variables fixed to 0/1 resp. or message that the instance is infeasible

```

1  $L_0 \leftarrow \emptyset, L_1 \leftarrow \emptyset;$ 
2 repeat
3   compute  $H = \text{Stab}(\text{Stab}(G, L_1), L_0);$ 
4   compute orbits  $\mathcal{O}$  of  $H$ ;
5   foreach orbit  $O$  with  $|O| \geq 2, O \cap (L_0 \cup L_1) = \emptyset$  do
6     let  $k = \min O$  be the first variable in  $O$ ;
7     fix all variables in  $O$  to 0 and perform constraint propagation;
8     if infeasible then
9       |  $L_1 \leftarrow L_1 \cup \{k\}$  (fix variable  $k$  to 1);
10    fix variable  $k$  to 1 and perform constraint propagation;
11    if infeasible then
12      |  $L_0 \leftarrow L_0 \cup O$  (fix all variables in  $O$  to 0);
13    if both problems are infeasible then
14      | exit: the instance is infeasible;
15    Let  $S_0$  and  $S_1$  be the variables outside of  $O$  fixed to 0 and 1, respectively, by both constraint
    propagations;
16     $L_0 \leftarrow L_0 \cup S_0, L_1 \leftarrow L_1 \cup S_1;$ 
17    if fixings found then
18      | break for loop;
19 until no fixings are found;
```

$$\begin{aligned}
 3x_1 + 4x_2 + x_3 + 2x_4 &\leq 5, \\
 4x_1 + x_2 + 2x_3 + 3x_4 &\leq 5, \\
 x_1, x_2, x_3, x_4 &\in \mathbb{Z}.
 \end{aligned}$$

The symmetry group is the cyclic group $C_4 = \langle (1234) \rangle$. If we branch on $x_1 = 1$, then the corresponding node has only trivial symmetries. Branching further on $x_3 = 1$ yields the problem

$$\begin{aligned}
 2x_2 + 4x_4 &\leq 1, \\
 3x_2 + x_4 &\leq -1, \\
 4x_2 + 2x_4 &\leq 1, \\
 x_2 + 3x_4 &\leq -1, \\
 x_2, x_4 &\in \mathbb{Z}.
 \end{aligned}$$

The symmetry group of this child problem is $\text{Stab}(C_4, \{1, 3\}) = \langle (13)(24) \rangle$. Note that, in general, $\text{Stab}(G, F_1)$ is a proper subgroup of the symmetry group.

2.5 Orbit probing

Probing for mixed 0/1-problems refers to the following method (see, e.g., Savelsbergh [46] and Atamtürk et al. [4]): Each binary variable is tentatively fixed to 0 and 1, respectively. In each case, a propagation round is performed, i.e., it is checked whether additional variables can be fixed. If infeasibility is detected in either case,

one can fix the variable to the opposite value. (If both cases are infeasible, the whole instance is infeasible.) If some other non-fixed variable is fixed to the same value in either branch, one can fix it to this value. Moreover, this method allows to detect implications between variables that can later be used to derive additional inequalities. Probing is often successful for 0/1-variables that are affected by symmetry, if symmetry handling is performed by any method that allows to fix variables.

Moreover, combining the ideas of orbital branching and probing, one obtains *orbit probing*, which seems to be a new idea: All variables in an orbit are tentatively fixed to 0. If the resulting problem is infeasible, one may fix the first variable in the orbit to 1. Otherwise, the first variable in an orbit is tentatively fixed to 1. If this is infeasible, one may fix all variables in the orbit to 0. If the resulting problem is infeasible in both cases, the whole instance is infeasible. After fixing variables to 1, the stabilizer of the variables fixed to 1 during this process (or the symmetry group) has to be recomputed. The resulting fixings can be seen as the outcome of performing orbital branching on a tree in which all nodes except one are cut off. Correctness thus follows from orbital branching.

An extension is to also fix variables outside the orbit, if they have been fixed to the same values in both cases. However, since in this case more than one node may be open, the future symmetry group depends on the previous fixings. Thus, one needs to recompute the stabilizer with respect to all variables fixed to 0 or 1. This leads to Algorithm 4.

2.6 Symmetry breaking inequalities

Another way to use symmetry in discrete optimization is to add inequalities to tighten the search space. Many examples appear in the literature, see, for example Sherali and Smith [49] and Margot [33] for an overview.

In the case of a general MIP with symmetry group G , the goal is to modify the feasible region in such a way that it is contained in a *fundamental domain* for G , see Friedman [13]. A fundamental domain is a polyhedron containing only one element from each G -orbit (this definition can be generalized to arbitrary regions, see Margot [33]). Therefore, the branch-and-bound tree does not encounter symmetric (optimal) solutions twice. In general, for any $d \in \mathbb{R}^n$, one can add inequalities

$$d^\top x \leq d^\top \pi(x) \quad (2)$$

for each element $\pi \in G$. For instance, taking $d = (2^{n-1}, 2^{n-2}, \dots, 2, 1)^\top$ leads to a lexicographic ordering if x is binary and produces a minimal fundamental domain, as proved by Friedman [13]. However, tight IP formulations for fundamental domains are in general hard to obtain or not practical, because of their size and of numerical problems in the description if the above lexicographic ordering is used. For the special cases of coordinate actions of the symmetric group and cyclic group, Friedman discusses efficient separation algorithms.

For a cyclic group C_k acting on variables x_1, \dots, x_k the following give valid inequalities, which seem to be part of the folklore (see, e.g., Liberti [26,27]):

$$x_1 \leq x_2, \quad x_1 \leq x_3, \dots, x_1 \leq x_k. \quad (3)$$

If a symmetric group S_k acts on the variables x_1, \dots, x_k , a tight IP formulation is given by

$$x_1 \leq x_2 \leq \dots \leq x_k. \quad (4)$$

In fact, if all variables x_1, \dots, x_k are integral, [18] showed that

$$x_k \leq x_1 + 1 \quad (5)$$

is a valid inequality for the symmetric group case. For matrix actions of cyclic or symmetric groups, Inequalities (3) and (4) remain valid when restricted to a single row. However, the additional inequality (5) cannot be generalized to matrix actions (see [42, Sect. 6.2.2]).

For general groups, one can use Inequalities (2). In order to obtain numerically more stable inequalities, we use $d = (1, 2, 3, \dots, n)^\top$, i.e.,

$$x_1 + 2x_2 + \dots + kx_k \leq x_{\pi(1)} + 2x_{\pi(2)} + \dots + kx_{\pi(k)}. \quad (6)$$

To control the size, we only add these inequalities for the generators of G .

Liberti [27] and Liberti and Ostrowski [28] discuss the possibility of adding inequalities that amount to selecting an arbitrary element of each orbit and adding inequalities like (3) to ensure that this variable is the smallest. Of course, such inequalities or Inequalities (2) and (3) could also be considered with respect to the local symmetry group of a node within the branch-and-bound tree. In our implementation, however, inequalities are only added in the beginning (as in [27,28]). Approximations of fundamental domains, i.e., symmetry breaking inequalities were for example described by Liberti [26].

For cyclic or symmetric groups acting on columns of 0/1-matrices with additional partitioning or packing conditions on the rows, so-called *orbitopes* can be used to completely handle symmetries. On the one hand, a complete linear description is known, see [22] and also Faenza and Kaibel [11] for a shorter proof. On the other hand, a method (*orbitopal fixing*) to further fix variables based on all currently branched or fixed variables has been introduced in [23] and was shown to be slightly faster than the separation of inequalities. Thus, we use orbitopal fixing by default.

We generally have to be careful not to add conflicting symmetry breaking inequalities. In our implementation we ensure this by adding only one type of inequality for each component of the direct product. See Liberti [27] and Liberti and Ostrowski [28] for alternative approaches.

Note that all techniques explained in this section (with the exception of orbitopes), i.e., inequalities (3), (4), and (6), work for general MIPs; Inequality (5) is valid for integral variables.

2.7 Projection onto the fixed space

While all previously mentioned methods need binary (or integral) variables, the following result allows to fix continuous variables to be equal within orbits. We call a symmetry *continuous variable restricted (CVR)* if it fixes all binary and integer variables pointwise.

Lemma 9 *Let the MIP (1) be feasible and let G' be the subgroup of CVR symmetries. Then there exists an optimal solution \bar{x} with $\bar{x}_i = \bar{x}_j$ for all $i, j \in O$ and for all orbits O of continuous variables with respect to G' .*

Proof Let x^* be an optimal solution of the MIP (1), and define

$$\bar{x} = \frac{1}{|G'|} \sum_{\pi \in G'} \pi(x^*).$$

By assumption, $\pi(i) = i$ for all $\pi \in G'$ and all integer variable indices $i \in \{1, \dots, p\}$. Thus, $\bar{x}_i = x_i^*$ for all $i = 1, \dots, p$. Since (1) is convex with respect to the continuous variables and each $\pi(x^*)$ is feasible, \bar{x} is feasible for (1). Moreover, we have

$$c^\top \bar{x} = \frac{1}{|G'|} \sum_{\pi \in G'} \underbrace{c^\top \pi(x^*)}_{=c^\top x^*} = c^\top x^*.$$

Thus, \bar{x} is optimal as well. We have for any $\sigma \in G'$:

$$\sigma(\bar{x}) = \frac{1}{|G'|} \sum_{\pi \in G'} \sigma(\pi(x^*)) = \frac{1}{|G'|} \sum_{\gamma \in G'} \gamma(x^*) = \bar{x},$$

since there is a bijection between each $\gamma \in G'$ and $\sigma \circ \pi$ for $\pi = \sigma^{-1} \circ \gamma \in G'$. Consequently, \bar{x} is constant along each orbit, which shows the claim. \square

The core idea of this lemma has been known for more general convex programs for some time; see, for instance, Gatermann and Parrilo [15, Thm. 3.3] in the context of semidefinite programs. An extension to integer programs is “orbital shrinking” by Fischetti and Liberti [12] (see also Salvagnin [45] and Mittelmann and Salvagnin [35] for applications).

Remark 10 The above proof implicitly refers to the fixed space of the group:

$$\text{Fix}(G) = \{x \in \mathbb{R}^n : x = \pi(x) \text{ for all } \pi \in G\}.$$

In fact, setting the variables to be equal within orbits corresponds to a projection on the fixed space. Using the fixed space for integer programming is discussed by Bödi et al. [7] and in [18]; see also [17,19,42]. For highly symmetric problems—the symmetry group is the alternating or symmetric group—this yields a polynomial-time algorithm to solve IPs [7]. More generally, knowledge of lattice-free orbit polytopes

may allow to restrict IPs to integer points close to the fixed space, which improves the performance of MIP solvers in some cases, see [18].

Remark 11 If we also allow non-CVR symmetries, i.e., we have a group acting on both continuous and integer variables, optimal solutions can be cut off by projection onto the fixed space as the following example shows:

$$\begin{aligned} \max \quad & x_4 + x_5 + x_6 \\ & 2x_4 + 3x_5 + x_1 + x_2 \leq 3, \\ & 2x_5 + 3x_6 + x_2 + x_3 \leq 3, \\ & 2x_6 + 3x_4 + x_3 + x_1 \leq 3, \\ & x_1 + x_2 + x_3 \geq 1, \\ & x_1, x_2, x_3 \in \mathbb{Z}_+. \end{aligned}$$

The symmetry group is a cyclic group generated by $(1\ 2\ 3)(4\ 5\ 6)$. An optimal solution is given by $(1, 0, 0, \frac{8}{35}, \frac{18}{35}, \frac{23}{35})^\top$. However, no point with $x_4 = x_5 = x_6 = \frac{7}{15} = \frac{1}{3} \cdot \frac{49}{35}$ is feasible.

3 Implementation details

We implemented the methods described in Sect. 2 in C++ using SCIP [47] version 3.2 with CPLEX 12.6.1 as LP solver. The symmetry computations are performed through `PermLib`, a C++ library for symmetry computations with permutations written by the second author [43]. The graph automorphisms are computed using `bliss` 0.72 [21].

We provide some details on the implementation.

Symmetry detection. Symmetry can be determined before or after presolving. Presolving can on the one hand remove symmetry, e.g., if variables are aggregated, or it can produce symmetry, e.g., if variables are eliminated that previously were the source of non-symmetry. If symmetry is detected before presolving, one needs to be careful that no successive presolving step destroys the found symmetry. In this case, we turn off presolving steps that can break symmetries. In SCIP this refers to the gateextraction, dominated column, and component presolvers. In any case, we make sure that symmetry is computed only once and only if it is required. Moreover, we make sure that certain types of variables are fixed pointwise, if this is needed. Furthermore, in order to avoid overly large computation times in this component, we limit the number of produced generators to 1500.

Isomorphism pruning. The implementation of isomorphism pruning only deals with the mixed 0/1-case, since the general case is significantly more complicated to implement. Note, however, that an implementation (`isop-1.2`) of the general integer case is available on the web page of Margot. Isomorphism Pruning is always combined with orbital fixing in our implementation, and the truncation techniques described in Sect. 2.4 can be applied.

Orbital branching. In the implementation of orbital branching, we also integrate orbital fixing. The corresponding fixings are performed at the time of creating the branching nodes. We also allow for the truncation of orbital branching in subnodes, see Sect. 2.4. Note that our implementation tries to apply orbital branching as long as there are nontrivial orbits containing fractional variables. Otherwise, other branching rules are applied.

Orbital fixing. Orbital fixing is also available as a separate propagation method and can be truncated in subnodes, see Sect. 2.4.

Symmetry breaking inequalities. We first split the symmetry group into a direct product so that each factor cannot be further decomposed. We then apply the basic group recognition of `PermLib` and add individual symmetry-breaking inequalities for each factor separately. If we detect a coordinate action of a cyclic or symmetric group, we add the Inequalities (3) and (4) (as well as (5) for integer variables), respectively. For matrix actions of cyclic or symmetric groups, we add Inequalities (3) or (4) for the first orbit only. We also try to detect an orbitope structure and apply the corresponding orbital fixing algorithm, see Sect. 2.6.

If none of the previous symmetry-breaking inequalities could be used, we allow to add inequalities (6) for each generator of the group. Additionally probing can be performed for all variables that are moved by a cyclic or symmetric group and for the diagonal variables in possibly detected orbitope structures.

Orbit probing. Orbit probing (see Sect. 2.5) is implemented straightforwardly, where the most time consuming step is to recompute the stabilizer (H in Algorithm 4). In order to limit the time resources, the default is to consider each variable at most once (even if probing with a different stabilizer group H might lead to reductions later). After orbit probing, we replace the symmetry group by the stabilizer of the variables fixed to 0 and 1 by orbit probing.

We also implemented the variant in which no deductions on the variables outside the considered orbit are made. Here, we only need to compute the stabilizer with respect to the variables fixed to 1.

Fixed space projection. In the implementation, we take care to apply the projection only for the fixed space of the subgroup of CVR symmetries. The corresponding variables are aggregated to be equal, i.e., they are removed from the problem.

PermLib. For all operations with permutation groups we use the second author's C++ library `PermLib` [43]. It offers basic functionality like orbit and stabilizer computations similar to GAP [14]. It represents permutation groups by bases and strong generating sets and employs backtrack search to compute set stabilizers (see also [48]). In order to decide whether a set is lexicographically minimal we implemented a backtrack search as described by Margot [30, Sect. 6.2].

Heuristic for symmetric subgroups. We also experimented with a fast heuristic that recognizes symmetric groups from transpositions. The resulting subgroup of the symmetry group is either trivial or is a direct product of coordinate actions of symmetric groups. However, our experiments (not reported here) show that this does not lead to an improvement using our implementation. The reason is that only using coordinate

actions of symmetric groups does not suffice to speed up the overall running time. It seems that a more sophisticated implementation would be needed that mixes the fast handling of symmetric groups with the more elaborate handling of other groups; moreover, the handling of matrix actions would need more attention.

Computational experiments. We performed extensive computational experiments to evaluate the performance of the various symmetry handling methods. The goal of the different experiments is to answer the following questions:

- How symmetric are different instances used in the literature?
- What kind of symmetries do appear?
- What is the proportion of time needed for computing symmetries?
- Can we reproduce computational experiments published in the literature?
- Depending on the instances, does it pay off to use symmetry handling?
- Which symmetry handling method is the best?

All computations were performed on a Linux cluster with Intel i3 CPUs with 3.2 Hz, 4 MB cache, and 8 GB memory running Linux. Each computation was performed single-threaded and a single process running on each computer. The code was compiled with gcc 4.4.5 with `-O3` optimization.

All times that we report are in seconds. The time limit for all computations has been set to 1 h. The time for instances that run into the time limit is evaluated as 3600 s.

4 Computational results for computing symmetries

The power of using symmetry handling methods depends on the amount and type of symmetry present in the tackled instances. Thus, as a first step, we analyzed the symmetry groups of MIPLIB instances.

4.1 MIPLIB 2003

For the MIPLIB 2003 [3], Table 1 shows the list of instances that contain symmetry before and after presolving. Note that Liberti [27] determined the types of the symmetry groups for these instances before presolving.

For instances containing symmetry, some additional work is necessary to transform the computed graph automorphisms into the internal data structures of `PermLib`. This explains the difference between “graph time”, i.e., the time needed by `bliss`, and the “total time”. Clearly, for instances without symmetry no additional work is necessary.

The results show that 30 of the 60 MIPLIB 2003 instances contain symmetry before presolving. The number of generators varies between 1 and 26,454. Not surprisingly, a larger number of generators typically leads to a larger time to compute symmetry. The time needed to compute graph automorphisms is usually small, with some extreme exceptions (e.g., `t1717` before presolving). The additional time needed to set up the `PermLib` data structures is usually small as well, but sometimes noticeable.

Moreover, 18 instances contain symmetry after presolving. Thus, presolving generally reduces symmetry, as can also be seen by the often significantly reduced number

Table 1 MIEP1B 2003 instances with symmetry before and after presolving; listed is the number of variables and constraints, the time to detect automorphisms of the graph as well as the total time (including the preparation of data structures), and the number of generators

Name	Before presolving				After presolving				
	#vars	#conss	Graph time	Total time	#gen	#vars	#conss	Total time	#gen
ark001	1388	1049	0.03	0.03	37	961	762	0.02	0
atlanta-ip	48,738	21,733	115.26	187.98	11,685	17,270	19,114	0.26	0
ds	67,732	657	0.28	0.30	2	64,030	626	0.28	0
fast0507	63,009	508	1.13	3.69	828	20,676	441	0.04	0
fiber	1298	364	0.00	0.00	1	1043	290	0.01	0
glass4	322	397	0.00	0.00	1	317	393	0.00	1
liu	1156	2179	0.01	0.01	1	1154	2179	0.01	0
mas74	151	14	0.00	0.00	2	150	14	0.00	2
mas76	151	13	0.01	0.01	2	150	13	0.00	2
misc07	260	213	0.01	0.01	2	232	224	0.00	2
mkc	5325	3412	0.20	0.23	195	3273	1288	0.13	194
mod011	10,958	4481	0.61	3.52	829	6489	1951	0.02	6
momentum3	13,532	56,823	1.04	1.07	55	13,151	49,376	0.90	0
msc98-ip	21,143	15,851	14.01	144.99	5151	12,725	14,988	0.13	4
mzzv11	10,240	9500	0.12	0.16	155	6558	6410	0.04	1
mzzv42z	11,717	10,461	0.11	0.15	110	7480	7372	0.06	0
net12	14,115	14,022	0.09	0.09	0	12,523	12,768	0.08	4
noswot	128	183	0.01	0.01	1	120	172	0.01	3
nsrand-ipx	6621	736	10.27	11.70	876	3798	536	0.04	1

Table 1 continued

Name	Before presolving			After presolving					
	#vars	#conss	Graph time	Total time	#gen	#vars	#conss	Total time	#gen
nw04	87,482	37	27.00	110.48	4505	46,143	36	0.18	0
opt1217	769	65	0.00	0.00	1	759	65	0.00	0
p2756	2756	756	0.01	0.01	29	2065	1422	0.02	25
protfold	1835	2113	0.02	0.02	2	1835	2113	0.02	2
qiu	840	1193	0.01	0.01	4	840	1193	0.01	4
rout	556	292	0.00	0.00	4	555	291	0.00	4
seymour	1372	4945	0.03	0.05	216	912	4412	0.01	2
stp3d	204,880	159,489	6.69	10.15	594	137,633	97,980	1.57	0
swath	6805	885	0.24	0.72	461	6260	483	0.01	0
tl1717	73,885	552	2565.42	2566.29	26,454	16,102	552	0.04	0
timtab1	397	172	0.00	0.00	1	201	167	0.00	1
timtab2	675	295	0.00	0.00	1	341	290	0.00	1

Table 2 Times for MIPLIB 2010 (total #331) symmetry detection split into instances with and without symmetry: column “graph” indicates the type of graph used. Column “presol” shows whether presolving is turned on. Column “graph time” and “total time” show the geometric means of the time in seconds to compute graph automorphisms and to compute symmetry in total, respectively. Column “limits” gives the number of instances for which symmetry computation failed because of the memory or time limit

Parameters		With symmetry			Without symmetry		
Graph	Presol	#	Graph time	Total time	#	Time	#Limits
Layered	No	171	2.47	3.05	149	2.12	11
Matrix	No	171	1.97	2.54	149	1.85	11
Layered	Yes	153	1.61	1.67	178	1.33	0
Matrix	Yes	153	1.36	1.43	178	1.24	0

of generators. One particular presolving step that reduces symmetry is the elimination of parallel or dominated columns. However, note that sometimes the number of generators increases (e.g., `noswot`) and there are instances in which presolving introduces symmetry (e.g., `net12`). The time for symmetry computation after presolving is quite small.

4.2 MIPLIB 2010

We next report on the symmetries present in the MIPLIB 2010 [24] instances. The testset for this section consists of all 361 problems in the MIPLIB 2010, excluding the 21 “unstable” and 11 “XXL” instances. This leaves 331 instances (two instances are contained in both subsets).

On these instances, we ran the two algorithms described in Sect. 1.1 to construct the vertex-colored graphs whose automorphism group corresponds to the symmetry group of the MIPs. As a heuristic for the decision of whether to group by variables or by constraints in the matrix graph, we group by variables whenever there are more variables than constraints.

The results are given in Table 2. The lines indicate whether the “matrix” or “layered” graph is used, and whether presolving is used or not. For some instances, we ran into the time limit or the memory limit of 8 GB, which is accounted for in the last column.

The results show that there are at least 171 instances initially containing symmetry, but only 153 if presolving is used. Thus, as for the MIPLIB 2003 instances, presolving already eliminates a certain amount of symmetry. Moreover, the version using the matrix approach for the graph is faster on average, both before and after presolving.

We then tried to analyze the symmetry group G using `PerMLib` on a computer with more memory (32 GB). The results are given in the Table 19 in the “Appendix”. There are 194 instances that might have a nontrivial symmetry group. The analysis ran into the memory limit or time limit of 10 h for 19 instances. Thus, there are 175 instances for which we actually tried to analyze the symmetry group.

The size of the symmetry groups range from 2 to about $10^{41.641.2}$. The percentage of variables that are moved by symmetry ranges from close to 0–100%. In total 35,042 factors were analyzed, and they very often consist of coordinate or matrix actions

of some \mathcal{S}_k : There are 29,962 and 4769 factors of coordinate and matrix actions of symmetric groups, respectively. However, the coordinate actions very often involve small symmetric groups. The type of about 294 factors could not be identified. No cyclic groups (other than \mathcal{S}_2) were detected.

Moreover, for 120 of the 331 examined MIPLIB2010 instances, in total 102,190 coordinate actions of full symmetric groups were found by a heuristic that recognizes symmetric groups from transpositions; the computation was stopped because of the memory or time limit for only two instances. In fact, this method finds significantly more symmetric group factors than the first analysis since it is able to complete the computations also for highly symmetric instances. These instances contribute the major share of symmetric group factors (see the table in the online supplement).

If we analyze the groups after presolving the picture changes as follows (see the table in the online supplement): 157 still contain symmetry, and we ran into the memory limit for two instances. In total 10,505 factors were analyzed, 9662 and 716 where coordinate and matrix actions of symmetric groups, respectively. The type of 125 factors could not be identified.

5 Computational results for symmetry handling methods

We use following testsets:

- Margot1: instances used in [30] (total: 16); we complemented the STS instances as described there.
- Margot2: additional highly symmetric instances by Margot [32,33], which used to be available on the web page of François Margot (total: 79);
- M2003-sym: all instances from MIPLIB 2003 [3] for which we found a non-trivial symmetry group after presolving (total: 18).
- M2010-sym: all instances from the MIPLIB 2010 [24] for which we found a non-trivial symmetry group after presolving (total: 154);
- M2010-bench: all instances from the MIPLIB 2010 benchmark suite (total: 87);

In this section, we report on computational experiments that investigate the different symmetry handling methods on these testsets. On one hand, we study different settings for the same method and on the other hand compare the different strategies. As a basis of comparison, we take SCIP with default settings, labeled as `default`.

For reporting aggregated results, we use the *shifted geometric mean* of values t_1, \dots, t_n :

$$\left(\prod (t_i + s) \right)^{1/n} - s$$

with shift s . We use a shift $s = 10$ for time and $s = 100$ for branch-and-bound nodes in order to decrease the strong influence of very easy instances in the mean values, see Achterberg [1] for more information. In all tables below that report aggregated results, `#nodes` refers to the shifted geometric mean of the number of nodes in the branch-and-bound tree, `time` refers to the shifted geometric mean of the CPU time

Table 3 Comparison of the number of nodes in the B&B-tree used by different implementations of isomorphism pruning and orbital branching on the `Margot1` testset: 1. isomorphism pruning implementation of Margot [30] (Table 2, version BC2), 2. our implementation, 3. implementation of orbital branching by Ostrowski et al. [38] (Table 3, Rule 5), and 4. our implementation

Problem	Isomorphism pruning		Orbital branching	
	#nodes [30]	#nodes ISP-first	#nodes [38]	#nodes OB-orbit3
cov954	655	364	249	79
cov1053	681	549	9775	1027
cov1054	447	1171	1249	15,884
cov1075	470	342	381	130
cov1076	22,454	23,822	31,943	26,372
cov1174	69,036	162,731	–	227,356
cod83	79	43	25	43
cod83r	121	69	–	68
cod93	653	400	1361	3030
cod93r	1301	2737	–	3100
cod105	19	7	11	7
cod105r	13	5	–	5
sts45	1571	1820	4709	2284
sts63	4499	4261	5533	5550
sts81	503	895	6293	1217

in seconds. Note that all times for the symmetry handling methods always include symmetry computation. Moreover, note that it depends on the particular symmetry group and the types of the variables it acts on whether a particular method allows to exploit symmetry.

The detailed results of all settings on all testsets are given in the extensive online supplement, containing about 150 tables.

5.1 Experiment 1: comparison to the literature

To begin our computational evaluation of the implemented symmetry handling methods, we compare the results of our implementation of isomorphism pruning and orbital branching with results published in the literature. Clearly, since the results have been obtained on different computers and with different branch-and-cut frameworks, it cannot be expected that the running times are the same. We would, however, expect the number of nodes to be roughly the same. Moreover, symmetry handling should clearly turn out to be effective, since this is one of the main results of the papers published in the literature.

Table 3 presents the best results from Margot [30] and Ostrowski et al. [38] on the `Margot1` testset and compares them to two variants of our implementation that were able to solve all instances (see Sects. 5.2 and 5.3, respectively). For isomorphism pruning, branching on the first index was used, as it was done in [30]. Let us mention that the results get worse, if we use a different branching rule—compare the experiments

in Sect. 5.2. Moreover, note that [38] does not present results of orbital branching for some of the instances in the testset.

The results show that our implementation produces about the same number of nodes as the implementations in the literature. Moreover, there seems to be a slight tendency that isomorphism pruning produces less nodes than orbital branching; but see Sect. 5.8 for the comparison of the timings.

5.2 Experiment 2: isomorphism pruning

We next compare different versions of isomorphism pruning (ISP) (see Sect. 2.2):

ISP :	isomorphism pruning with default SCIP branching rule;
ISP-NST :	ISP with “no subtree”: symmetry handling is turned off in the subtree of a node for which the stabilizer of the fixed variables is empty;
ISP-first :	isomorphism pruning with first index branching;
ISP-NST-first :	ISP-NST with first index branching.

Table 4 presents the results of these settings on the five testsets; we additionally present the results of the default settings (“default”) and default settings with a first index branching rule (“first”). Moreover, Table 5 presents a comparison of selected settings on all instances that could be solved to optimality by all of these settings; this, for instance, allows for a fair comparison of the number of nodes.

As reported in the literature, isomorphism pruning is very effective for the highly symmetric instances in `Margot1`: the running time is about two orders of magnitude smaller than the default and it can solve eleven more instances (for `ISP-first`). This tremendous improvement is due to the extraordinary reduction of the number of branch-and-bound nodes.

The time needed for symmetry computation is negligible, never exceeding 0.1 s. However, the additional time needed for isomorphism pruning is often quite significant, in extreme cases using more than 90% of the total time. However, for these instances, this effort is very well invested.

Using isomorphism pruning together with first-index branching (`ISP-first` and `ISP-NST-first`) about halves the number of nodes and roughly cuts the running time to one third compared to `ISP` (see Table 5). This is also statistically significant: a Wilcoxon signed rank test, see Berthold [5], confirmed a statistically significant reduction of the time with a p -value of less than 0.0005. Note that this reduction is only realized together with isomorphism pruning: running first index branching without isomorphism pruning (setting `first`) does not improve upon the default settings. Furthermore, the success of the first-index branching rule is quite specific to these instances: randomly permuting the variables results in `ISP-NST-first` being slower than `ISP-NST`; see also the comments in Sect. 5.8.

Finally, variant `ISP-NST` improves upon `ISP` in terms of running time; this is statistically significant with a p -value less than 0.01. However `ISP-NST` increases the number of nodes, as can be expected; see Table 5.

The results for `Margot2` turn out to be different: The default version performs quite well—it allows to solve 65 (of 79) instances. Obviously many instances are easily

Table 4 Comparison of different isomorphism pruning variants: given are the shifted geometric means of the number of branch-and-bound nodes (#nodes) and CPU time in seconds (time), the number of instances solved to optimality (#opt), the number of instances in which isomorphism pruning was active (#act), the geometric mean of the number of calls of isomorphism pruning (#calls), the geometric mean of the number of domain reductions performed (#red), the geometric mean of the number of node cutoffs detected (#cutoff), and the shifted geometric mean of the time used by isomorphism pruning including symmetry computation (ISP-time)

Setting	#nodes	Time	#opt	#act	#calls	#red	#cutoff	ISP-time
Margot1 (16)								
Default	253,743.3	1174.06	5	0	–	–	–	–
First	276,284.5	1210.15	5	0	–	–	–	–
ISP	1844.1	55.58	15	16	827.5	1603.4	60.4	18.50
ISP-NST	2208.4	48.31	15	16	985.7	1029.5	11.1	3.97
ISP-first	816.4	15.21	16	16	447.2	1772.9	35.9	6.29
ISP-NST-first	951.3	14.38	16	16	519.8	1631.2	18.2	4.76
Margot2 (79)								
Default	969.7	32.64	65	0	–	–	–	–
First	1117.8	35.94	64	0	–	–	–	–
ISP	517.4	21.75	69	28	11.3	8.2	3.0	5.03
ISP-NST	521.7	20.95	69	28	11.4	8.0	2.1	2.11
ISP-first	555.8	21.30	69	26	13.4	11.0	4.2	6.18
ISP-NST-first	589.8	21.00	69	26	14.1	10.5	3.3	4.50
M2003-sym (18)								
Default	81,955.4	568.43	13	0	–	–	–	–
ISP	53,800.6	526.18	13	8	2762.1	6.6	1.8	13.73
ISP-NST	69,514.8	524.47	13	7	3603.3	6.6	1.2	5.95
M2010-sym (154)								
Default	9061.6	1084.21	60	0	–	–	–	–
ISP	3334.9	1004.48	68	94	791.9	33.7	4.1	92.63
ISP-NST	5090.2	925.03	70	88	1241.6	34.1	2.3	29.17
M2010-bench (87)								
Default	16,375.3	580.00	62	0	–	–	–	–
ISP	10,589.4	495.80	66	24	22.4	4.4	1.5	9.72
ISP-NST	10,778.4	467.62	67	21	22.8	4.3	1.1	4.69

solved. Moreover, *ISP*, *ISP-first*, *ISP-NST*, and *ISP-NST-first* perform similar: they roughly halve the number of nodes and reduce time to 65% compared to *default* (see Table 5). The best variant seems to be *ISP-NST* by a slight margin, which solves 69 instances within the time limit.

The picture again changes, when considering the results on the MIPLIB testsets *M2003-sym*, *M2010-sym*, and *M2010-bench*. As can be expected, the variants using first index branching perform significantly worse than all other variants, since it does not take the particular problem structure into account. The corresponding results are therefore not reported in the following.

Table 5 Comparison of different isomorphism pruning variants on instances solved to optimality for all selected settings

Setting	Margot1 (15)		Margot2 (65)		M2003-sym (13)		M2010-sym (58)		M2010-bench (61)	
	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time
Default	-	-	255.2	6.9	22,322.5	177.5	818.1	149.5	6053.9	272.0
ISP	1226.6	40.3	113.8	4.0	16,418.9	155.9	649.1	159.4	5391.9	272.3
ISP-NST	1456.9	34.4	113.2	3.7	16,418.9	155.0	648.9	138.0	5344.9	263.0
ISP-first	548.8	9.8	115.6	3.6	-	-	-	-	-	-
ISP-NST-first	658.3	9.4	120.1	3.6	-	-	-	-	-	-

The best variant is `ISP-NST`. In particular, `ISP-NST` improves upon the default settings by about 7% (`M2003-sym`), 15% (`M2010-sym`), and 19% (`M2010-bench`) in Table 4. Moreover, the improvements in Table 5 are about 13% (`M2003-sym`), 8% (`M2010-sym`), and 3% (`M2010-bench`). Although these are very good improvements, they are not statistically significant, except for `M2010-sym` with a p -value less than 0.0005. This seems to be due to the fact that there are certain instances where `ISP-NST` is significantly slower. However, since `ISP-NST` is also able to increase the number of solved instances (10 for `M2010-sym`, and 5 for `M2010-bench`), we conclude that this variant seems to be a good choice, even for `MIPLIB` instances and definitely for the Margot instances.

In order to reduce the effect of performance variability, we also ran `ISP-NST` and `default` on ten permuted instances for each instance in `M2010-bench`. On average, `ISP-NST` is 14% faster than `default` on all instances. On the instances solved to optimality by all settings and for all permutations, `ISP-NST` is 5% faster and uses 12% less nodes. Moreover, `ISP-NST` is able to solve two more instances for *each* permutation than `default`. These results support the above claims.

Note that isomorphism pruning is considered active if it allows to fix variables or cut off a node. Thus, these numbers may vary across different variants.

Finally, recall that the time needed for the computation of symmetries is included in the time for isomorphism pruning. In particular, this time also arises for the 53 instances in the `M2010-bench` testset that are not symmetric after presolving. Moreover, note that isomorphism pruning was not applied for some instances of the `MIPLIB` testsets, since no symmetry on 0/1 variables was present. Furthermore, the time needed for isomorphism pruning is notable in relation to the total time, but is still reasonable on average. However, there are extreme cases, where most of the total time is used in isomorphism pruning. This situation could be improved with a filtering of instances for which isomorphism pruning is reasonably fast and effective, and it could possibly be improved with a more refined implementation.

5.3 Experiment 3: orbital branching

We now turn to the investigation of the reimplementations of orbital branching (OB), see Sect. 2.3. We investigate the following orbital branching variants:

<code>OB-orbit0</code>	<code>OB</code> : orbital branching with choosing the largest orbit (default);
<code>OB-orbit1</code> :	orbital branching with choosing the orbit that locally tends to break the highest amount of symmetry;
<code>OB-orbit2</code> :	orbital branching with choosing the orbit that locally tends to preserve the highest amount of symmetry;
<code>OB-orbit3</code> :	orbital branching with choosing the orbit that maximizes the product of the size of the orbit with the maximal size of a child orbit;
<code>OB-min4</code> :	<code>OB</code> , only executed for orbits of size at least four;
<code>OB-NST</code> :	<code>OB</code> , turning off symmetry handling in the subtree of a node for which the stabilizer of the fixed variables is empty;
<code>OB-first</code> :	<code>OB-NST</code> with first index branching rule as a fall-back rule.

Table 6 Comparison of different orbital branching variants: Given are the shifted geometric means of the number of B&B-nodes (#nodes) and CPU time in seconds (time), the number of instances solved to optimality (#opt), the number of instances in which orbital branching was active (#act), the geometric mean of the number of children created by orbital branching (#children), the geometric mean of the number of variables fixed to 0 and 1, respectively (#fixed0, #fixed1), the geometric mean of the number of cutoffs (#cutoffs), and the shifted geometric mean of the time used by orbital branching including symmetry computation (OB-time)

Setting	#nodes	Time	#opt	#act	#children	#fixed0	#fixed1	#cutoffs	OB-time
Margot1 (16)									
Default	253, 743.3	1174.06	5	0	-	-	-	-	-
OB-orbit0	2034.1	40.90	15	16	193.2	151.8	1.7	1.5	11.19
OB-orbit1	3113.5	58.85	14	16	299.8	70.5	1.5	1.3	16.70
OB-orbit2	1186.7	36.20	15	16	453.2	175.5	1.4	1.5	22.44
OB-orbit3	976.9	33.75	16	16	146.9	95.8	1.2	1.4	19.04
OB-NST	2131.0	36.70	15	16	93.6	54.8	1.0	1.1	0.91
OB-min4	2710.0	49.99	15	16	52.5	69.5	1.3	1.2	12.92
OB-first	2685.5	42.07	14	16	271.4	160.6	2.5	1.8	13.09
Margot2 (79)									
Default	969.7	32.64	65	0	-	-	-	-	-
OB-orbit0	533.4	21.05	68	28	5.9	4.9	1.2	1.0	4.28
OB-orbit1	562.7	21.69	67	28	5.8	3.4	1.1	1.0	4.47
OB-orbit2	452.8	21.93	67	28	8.0	4.3	1.3	1.7	8.88
OB-orbit3	492.4	21.54	68	28	5.6	4.7	1.1	1.4	7.83
OB-NST	553.6	20.91	67	28	4.6	3.9	1.1	1.0	0.13
OB-min4	568.5	22.35	67	28	3.6	4.0	1.1	1.0	3.94
OB-first	624.3	21.31	68	28	6.8	5.2	1.4	1.1	5.20

Table 6 continued

Setting	#nodes	Time	#opt	#act	#children	#fixed0	#fixed1	#cutoffs	OB-time
M2003-sym (18)									
Default	81, 955.4	568.43	13	0	–	–	–	–	–
OB-orbit0	42, 699.9	553.40	12	9	11.6	1.5	1.0	1.0	47.98
OB-orbit1	36, 271.2	553.41	12	9	9.6	1.0	1.0	1.0	48.18
OB-orbit2	41, 718.4	578.59	12	7	3.4	1.2	1.0	1.1	51.11
OB-orbit3	42, 636.6	567.38	12	8	4.4	1.1	1.0	1.1	50.88
OB-NST	50, 106.6	510.20	13	5	2.9	1.4	1.0	1.0	4.23
OB-min4	55, 223.0	604.05	12	3	1.8	1.1	1.0	1.0	47.21
M2010-sym (154)									
Default	9061.6	1084.21	60	0	–	–	–	–	–
OB-orbit0	3493.9	1276.23	56	112	86.4	8.8	1.8	1.4	234.13
OB-orbit1	3107.9	6673.38	50	109	68.7	7.0	1.6	1.3	731.44
OB-orbit2	2482.3	1322.42	56	79	25.3	5.7	1.6	2.0	341.19
OB-orbit3	2289.4	1554.07	47	100	30.4	5.5	1.5	1.8	422.50
OB-NST	5168.0	1215.22	59	107	49.9	5.8	1.4	1.0	52.08
OB-min4	3551.1	1232.82	58	79	16.2	3.6	1.1	1.1	207.71
M2010-bench (87)									
Default	16, 375.3	580.00	62	0	–	–	–	–	–
OB-orbit0	12, 040.2	663.83	58	28	6.9	2.5	1.4	1.1	24.74
OB-orbit1	11, 443.3	700.40	57	28	6.3	2.2	1.3	1.1	28.60
OB-orbit2	10, 272.5	666.24	61	22	3.6	2.2	1.2	1.4	27.95
OB-orbit3	10, 914.7	702.22	57	26	4.0	2.1	1.2	1.4	29.52
OB-NST	13, 550.8	629.97	61	24	5.0	1.9	1.2	1.0	9.61
OB-min4	11, 832.2	629.75	60	17	2.6	1.4	1.1	1.0	20.60

Table 7 Comparison of different orbital branching variants on instances solved to optimality for all selected settings

Setting	Margot1 (13)		Margot2 (65)		M2003-sym (12)		M2010-sym (38)		M2010-bench (52)	
	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time
Default	-	-	255.2	6.9	24,462.4	149.7	1089.4	147.5	7806.4	285.3
OB-orbit0	557.4	10.6	113.6	3.5	14,408.0	129.9	1035.4	144.2	7280.6	278.3
OB-orbit1	915.3	17.9	121.7	3.8	14,301.6	129.9	1072.4	165.9	7529.3	288.1
OB-orbit2	629.8	18.4	108.9	4.0	18,504.6	140.9	989.2	166.5	7233.9	283.2
OB-orbit3	355.0	14.1	107.8	3.6	17,824.1	136.0	949.1	184.9	7148.1	286.2
OB-NST	581.5	8.7	116.8	3.5	14,439.5	122.6	1066.0	140.1	7411.2	276.5
OB-min4	799.6	14.5	129.0	4.1	22,729.5	152.7	1079.4	156.2	7631.3	290.2
OB-first	782.1	11.6	130.6	3.7	-	-	-	-	-	-

The results for orbital branching are given in Table 6. In general, the number of nodes created by orbital branching is quite small (see column “#children”): For *Margot1* the percentage is less than 30% and for all other testsets less than 1%. Thus, in most of the nodes, no orbits that can be used for branching were found, i.e., all orbits were trivial or they did not contain fractional variables or orbital fixing can be applied to fix all variables in the orbits. Note, however, that, as mentioned above, symmetry is not recomputed in our implementation; see also Ostrowski et al. [38] for a discussion.

For the *Margot* instances, the basic results are similar to isomorphism pruning: *default* is significantly worse than the other variants—the results being more pronounced for *Margot1* than for *Margot2*. Variant *OB-min4* is not able to improve on the basic version *OB=OB-orbit0*. Variant *OB-first*, which uses first index branching if no branching orbit has been chosen, performs quite well, but *OB* and *OB-NST* are still faster. We conclude that first index branching is not as important for orbital branching as it is for isomorphism pruning.

According to Table 7, the fastest among the four different variants (*OB-orbit0*, *OB-orbit1*, *OB-orbit2*, *OB-orbit3*) to choose the branching orbit is *OB-orbit0* (i.e., our default orbital branching rule). However, Table 6 suggests *OB-orbit3* to be faster. Moreover, the only variant that solves all instances in the *Margot1* testset is *OB-orbit3*, which also produces the smallest number of nodes in Table 7. Finally, in the results of [38, Table 3] the rule corresponding to *OB-orbit2* performed best for the case of using the global symmetry group, as we do in our implementation. There might be several reasons for these differences: The testset in [38] is slightly different and the implementation is based on different frameworks, the differences between our variants are not large, and the results depend on the particular instances on a small testset. In summary, *OB-orbit0* seems to be a solid default choice (especially considering Table 7 and the other testsets). Its offspring *OB-NST* is the fastest method in Tables 6 and 7.

Turning to the *MIPLIB* testsets, we expect orbital branching to perform less well, because branching on orbits tends to be less important compared to branching on variables not affected by symmetry. Indeed, the best variant based on orbital branching in Table 6 is *OB-NST*, which uses significantly more time and solves less instances than the default settings on *M2010-sym* and *M2010-bench*. Nevertheless, with respect to the instances solved to optimality in Table 7, *OB-NST* is faster than the default settings by about 18% for *M2003-sym*, 5% for *M2010-sym*, and 3% for *M2010-bench*. Thus, *OB-NST* performs quite well on these instances and the total number of solved instances is among the largest of the orbital branching variants, but still less than the default. As mentioned above, one reason for the worse performance of orbital branching on the *MIPLIB* instances is that branching on orbits tends to be less important. Another reason is that there are several instances for which orbital branching takes a very long time, compared to the default; for *M2010-bench*, these instances are *biella1*, *lectsched-4-obj*, *macrophage*, *mcsched*, *n3seq24*, *neos18*, and *tanglegram2*. For some instances (e.g., *biella1* and *tanglegram2*), the time needed to compute the stabilizers is too large, for the others the branching decisions made by orbital branching are obviously not helpful. In fact, the current implementation always uses orbital branching first, if this is applicable. Otherwise the default (or first index) branching rule takes over. It is likely that a prior-

itization of the different branching possibilities, for instance using strong branching, would help to improve the situation. However, there are still several exceptional instances for which the time needed for orbital branching is too large to be competitive; again, this might be remedied by improving the implementation or applying more strict working limits.

Note that the different variants have a different number of “active” instances; we count an instance as active if orbital branching created children, fixed variables, or cut off nodes. For example, `OB-orbit2` allows to cut off many nodes for the MIPLIB instances, but generates few children.

5.4 Experiment 4: orbital fixing

As mentioned in Sect. 2.1, orbital fixing (OF) can also be applied as a stand-alone component. We consider the following variants:

- OF: orbital fixing as described in Sect. 2.1;
- OF-NST: OF with “no subtree”: symmetry handling is turned off in the subtree of a node for which the stabilizer of the fixed variables is empty;
- OF-first: OF with first index branching;
- OF-NST-first: OF-NST with first index branching.

The corresponding results are given in Tables 8 and 9. The general behavior is quite similar to isomorphism pruning. Interestingly, the running times are similar as well. The results show that orbital fixing decreases the number of nodes in the tree (see Table 9), if it can be applied. The first index branching variants `OF-first` and `OF-NST-first` are the fastest for `Margot1` and are the only variants to solve all instances. Moreover, `OF-NST` improves upon the default: e.g., the time is reduced by about 19% for `M2010-bench` in Table 8 and by about 4% in Table 9. In this case, the result are statistically significant for `M2010-sym` with a p -value of less than 0.0005 and for `M2010-bench` with a p -value less than 0.005, but not for `M2003-sym`. The number of solved instances is increased by 10 for `M2010-sym` and 5 for `M2010-bench`. Moreover, a comparison of ten permuted runs of `M2010-bench` shows that `OF-NST` is about 14% faster than the default on all instances on average, 5% faster on the instances solved to optimality for all permutations, and it allows to solve two more instances for *all* permutations. This makes orbital fixing a competitive variant.

5.5 Experiment 5: orbit probing

We start the evaluation of orbit probing (OP) (see Sect. 2.5) using the following two variants:

- OP: orbit probing;
- OP-nod: orbit probing without deductions outside orbits.

Variant `OP-nod` does not need to compute the stabilizer with respect to the variables fixed to 0. Thus, the symmetry group considered tends to be larger, which might provide

Table 8 Comparison of orbital fixing variants: given are the shifted geometric means of the number of branch-and-bound nodes (#nodes) and CPU time in seconds (time), the number of instances solved to optimality (#opt), the number of instances in which orbital fixing was active (#act), the geometric mean of the number of calls of orbital fixing (#calls), the geometric mean of the number of domain reductions performed (#red), the geometric mean of the number of node cutoffs detected (#cutoff), and the shifted geometric mean of the time used by orbital fixing including symmetry computation (OF-time)

Setting	#nodes	Time	#opt	#act	#calls	#red	#cutoff	OF-time
Margot1 (16)								
Default	253, 743.3	1174.06	5	0	–	–	–	–
OF	2192.5	61.70	15	16	985.4	1659.8	1.4	16.28
OF-NST	2313.5	49.47	15	16	1038.0	1126.7	1.4	3.50
OF-first	1097.8	21.16	16	16	587.7	2190.8	2.0	8.40
OF-NST-first	1089.9	17.78	16	16	582.0	1955.2	1.9	6.03
Margot2 (79)								
Default	969.7	32.64	65	0	–	–	–	–
OF	520.1	21.56	68	28	11.3	8.3	1.1	4.24
OF-NST	518.0	21.03	68	28	11.3	8.1	1.2	1.88
OF-first	576.9	21.36	68	26	14.0	11.2	1.4	5.64
OF-NST-first	591.9	21.03	68	26	14.3	10.7	1.2	3.95
M2003-sym (18)								
Default	81, 955.4	568.43	13	0	–	–	–	–
OF	52, 754.7	524.34	13	7	2722.8	6.6	1.0	11.86
OF-NST	69, 554.7	523.52	13	7	3607.8	6.6	1.0	5.27
M2010-sym (154)								
Default	9061.6	1084.21	60	0	–	–	–	–
OF	3560.7	974.22	67	82	852.3	32.9	1.6	83.52
OF-NST	5301.0	917.86	70	82	1301.7	32.3	1.5	26.00
M2010-bench (87)								
Default	16, 375.3	580.00	62	0	–	–	–	–
OF	10, 216.5	489.68	66	21	21.6	4.5	1.1	9.09
OF-NST	10, 978.2	469.00	67	21	23.2	4.3	1.0	4.45

a higher potential for fixing variables directly from the tentative fixings. However, it also loses the possibility to exploit the fixings outside of the considered orbits.

The results of these two variants are displayed in Tables 10 and 11. These results show that orbit probing fixes relatively few variables, in general. In fact, there is a small speed-up of both variants on Margot1 and Margot2. The results on the MIPLIB instances are somewhat inconclusive: For M2010-sym and M2010-bench, the number of solved instances increases by two for OP, but decreases by one for M2003-sym. For M2010-sym the time of OP and OP-nod increases with respect to the default both in Tables 10 and 11, while on M2010-bench, the running time of OP decreases by about 5% in Table 10, but only very slightly in Table 11.

Table 9 Comparison of different orbital fixing variants on instances solved to optimality for all selected settings

Setting	Margot1 (15)		Margot2 (65)		M2003-sym (13)		M2010-sym (58)		M2010-bench (61)	
	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time
Default	-	-	255.2	6.9	22,322.5	177.5	818.1	149.5	6053.9	272.0
OF	1476.8	45.4	111.4	3.9	16,452.2	154.9	640.2	145.2	5329.4	267.3
OF-NST	1529.6	35.3	111.6	3.7	16,452.2	154.5	642.8	135.6	5347.3	262.2
OF-first	727.4	14.0	120.1	3.6	-	-	-	-	-	-
OF-NST-first	721.5	11.4	122.2	3.6	-	-	-	-	-	-

Table 10 Results of orbit probing variants: given are the shifted geometric means of the number of branch-and-bound nodes (#nodes) and CPU time in seconds (time), the number of instances solved to optimality (#opt), the number of instances in which orbit probing was active (#act), the number of fixed variables (#fixed), and the shifted geometric mean of the time used for probing including symmetry computation (OP-time)

Setting	#nodes	Time	#opt	#act	#fixed	OP-time
Margot1 (16)						
Default	253, 743.3	1174.06	5	0	–	0.00
OP	217, 600.5	1021.72	6	6	6	0.02
OP-nod	217, 398.5	1020.92	6	6	6	0.02
Margot2 (79)						
Default	969.7	32.64	65	0	–	0.00
OP	686.5	25.18	67	49	96	0.00
OP-nod	686.1	25.18	67	45	80	0.00
M2003-sym (18)						
Default	81, 955.4	568.43	13	0	–	0.00
OP	72, 801.1	579.37	12	5	32	0.18
OP-nod	73, 454.4	579.51	12	4	31	0.15
M2010-sym (154)						
Default	9061.6	1084.21	60	0	–	0.00
OP	7824.0	1120.56	62	43	4420	8.32
OP-nod	8157.5	1125.89	60	36	5608	7.22
M2010-bench (87)						
Default	16, 375.3	580.00	62	0	–	0.00
OP	14, 569.8	547.89	64	12	675	2.05
OP-nod	15, 732.3	568.57	63	9	1098	1.82

Comparing OP and OP-nod shows that the time used within OP-nod is less than the time spend within OP and that OP-nod finds more fixings. However, OP-nod results in a larger overall computation time in Table 10 (except for Margot1). Table 11 does not show a significant difference.

Consequently, orbit probing should be a candidate for the application together with other techniques, since it sometimes improves the performance and is relatively inexpensive. We therefore combined orbit probing with isomorphism pruning (OP-ISP-NST), orbital branching (OP-OB-NST), and orbital fixing (OP-OF-NST). However, the results show that these variants often solve less instances and are generally slower. We therefore only present the corresponding tables in the online supplement. The somewhat surprising outcome is that orbit probing does not seem to improve the performance of other methods.

5.6 Experiment 6: symmetry breaking inequalities

With respect to symmetry breaking constraints, we investigate the following variants as described in Sect. 2.6:

Table 11 Comparison of different orbit probing variants on instances solved to optimality for all selected settings

Setting	Margot1 (5)		Margot2 (65)		M2003-sym (12)		M2010-sym (58)		M2010-bench (61)	
	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time
Default	29,761.4	91.8	255.2	6.9	24,462.4	149.7	886.3	149.1	6059.9	267.7
OP	21,261.9	70.5	149.2	4.1	22,873.0	141.3	850.7	162.7	5900.4	266.5
OP-nod	21,261.9	70.4	149.2	4.1	22,873.0	141.5	908.9	159.5	6163.4	266.4

<code>S</code> :	add inequalities (3), (4), (5) or orbitopes if applicable;
<code>S-fund</code> :	variant <code>S</code> with fundamental domain inequalities (6);
<code>S-fundnolp</code> :	variant <code>S</code> with (6) handled in propagation, but not added to the LP;
<code>S-nolp</code> :	as variant <code>S</code> , but do not add inequalities to LP—instead propagate them;
<code>S-prob</code> :	variant <code>S</code> with probing;
<code>S-orbitmin</code> :	add inequalities (3) for all orbits using the minimal element.

Table 12 reports the results on symmetry breaking inequalities and Table 13 shows a comparison on the instances solved to optimality.

For the Margot instances, only `S-orbitmin` and `S-fund` are able to improve on the default settings. But the running times are still an order of magnitude slower than isomorphism pruning. All other variants are not able to perform better, because no structures like full symmetric subgroups could be found.

For the MIPLIB instances, the results are somewhat inconclusive. Variants `S`, `S-fundnolp`, and `S-orbitmin` seem to perform well, but none dominates the other. Variant `S-orbitmin` increases the number of solved instances by three for `M2010-sym` and for `M2010-bench`. Moreover, with respect to the default, the running time is improved by about 18% for `M2003-sym`, 7% for `M2010-sym`, and 1% for `M2010-bench` with respect to the number of instances solved to optimality by all variants, see Table 13. In fact, all other variants (except `S-fund` on `M2010-sym`) are also able to improve on the default. None of the improvements is statistically significant. Note that three orbitope structures are found for `M2010-sym` and one for `M2010-bench`. Moreover, probing does not improve on variant `S`.

In total, we conclude that adding Inequalities (3) for each orbit (`S-orbitmin`) sometimes performs better than adding Inequalities (4) for each full symmetric group factor (variant `S`) and sometimes worse. Moreover, we checked the behavior of the root-LP bound, which very slightly improves for `S-orbitmin`, but not for variant `S`. We have to leave it to future research to understand these outcomes and possibly improve on these results.

5.7 Experiment 7: projection on the fixed space

The projection on the fixed space is only useful for instances that contain continuous variables. Thus, we do not report results on the Margot testsets. Moreover, we also study the following combination with other variants:

<code>lpfix</code> :	default fixing as described in Sect. 2.7;
<code>lpfix-OP</code> :	<code>lpfix</code> together with orbit probing variant <code>OP</code> ;
<code>lpfix-fundnolp</code> :	<code>lpfix</code> together with variant <code>S-fundnolp</code> ;
<code>lpfix-orbitmin</code> :	<code>lpfix</code> together with variant <code>S-orbitmin</code> .

The results are shown in Tables 14 and 15. Variant `lpfix` is able to fix quite a number of variables. For `M2003-sym` nothing can be gained. For `M2010-sym` and `M2010-bench` all variants slightly improve on the default. In total, the LPs become smaller, but the benefit is obviously not large. Moreover, the performance

Table 12 Comparison of different symmetry breaking inequality variants: given are the shifted geometric means of the number of branch-and-bound nodes (#nodes) and CPU time in seconds (time), the number of instances solved to optimality (#opt), the number of fundamental domains handled (#fund), the number of orbitope structures found (#orb), the number of cyclic groups handled, the total number of inequalities added (#tot), and the shifted geometric mean of the time used for adding inequalities including symmetry computation (sym-time)

Setting	#nodes	Time	#opt	#act	#full	#fund	#orb	#cyc	#tot	#fixed	Sym-time
Margot1 (16)											
Default	253, 743.3	1174.06	5	0	-	-	-	-	-	-	0.00
S	253, 255.2	1173.85	5	0	-	-	-	-	-	-	0.10
S-fund	157, 202.6	940.14	6	16	-	145	-	-	145	-	0.11
S-fundnolp	252, 252.6	1183.89	5	16	-	145	-	-	145	-	0.11
S-nolp	252, 787.9	1174.16	5	0	-	-	-	-	-	-	0.10
S-prob	253, 260.1	1173.11	5	0	-	-	-	-	-	-	0.11
S-orbitmin	103, 895.7	519.73	9	16	-	-	-	16	3794	-	0.13
Margot2 (79)											
Default	969.7	32.64	65	0	-	-	-	-	-	-	0.00
S	970.0	32.71	65	0	-	-	-	-	-	-	0.05
S-fund	809.5	28.89	67	70	-	524	-	-	524	-	0.05
S-fundnolp	987.8	33.40	65	70	-	524	-	-	524	-	0.05
S-nolp	970.0	32.72	65	0	-	-	-	-	-	-	0.05
S-prob	970.0	32.71	65	0	-	-	-	-	-	-	0.04
S-orbitmin	593.2	24.09	68	70	-	-	-	70	10,118	-	0.06

Table 12 continued

Setting	#nodes	Time	#opt	#act	#full	#fund	#orb	#cyc	#tot	#fixed	Sym-time
M2003-sym (18)											
Default	81, 955.4	568.43	13	0	-	-	-	-	-	-	0.00
S	74, 248.9	550.71	12	16	114	-	-	-	168	-	0.12
S-fund	70, 802.7	537.97	12	18	114	91	-	-	259	-	0.12
S-fundnolp	76, 985.9	535.75	13	18	114	91	-	-	259	-	0.12
S-nolp	77, 309.0	535.73	13	16	114	-	-	-	168	-	0.12
S-prob	75, 590.0	554.09	12	16	114	-	-	-	168	65	0.12
S-orbitmin	69, 156.1	537.14	12	18	-	-	-	121	185	-	0.03
M2010-sym (154)											
Default	9061.6	1084.21	60	0	-	-	-	-	-	-	0.00
S	8481.9	1055.29	60	114	9832	-	3	-	19,186	6058	3.19
S-fund	7985.2	1064.99	62	152	9832	2154	3	-	21,340	6058	3.21
S-fundnolp	8874.3	1046.87	61	152	9832	2154	3	-	21,340	6058	3.22
S-nolp	9391.2	1069.33	60	113	9831	-	3	-	19,185	6058	3.21
S-prob	8554.3	1056.96	60	114	9832	-	3	-	19,186	6367	3.31
S-orbitmin	8847.6	1039.67	63	153	-	-	-	9943	23,405	-	1.36
M2010-bench (87)											
Default	16, 375.3	580.00	62	0	-	-	-	-	-	-	0.00
S	15, 016.9	565.00	62	31	1439	-	1	-	5614	3	0.61
S-fund	14, 527.6	556.62	63	39	1439	90	1	-	5704	3	0.61
S-fundnolp	15, 406.9	539.76	63	39	1439	90	1	-	5704	3	0.61
S-nolp	16, 112.8	560.63	63	31	1439	-	1	-	5614	3	0.61
S-prob	15, 142.6	571.88	62	31	1439	-	1	-	5614	33	0.64
S-orbitmin	15, 658.7	556.63	65	39	-	-	-	1457	5723	-	0.55

Table 13 Comparison of different symmetry breaking inequality variants on instances solved to optimality for all selected settings

Setting	Margot1 (5)		Margot2 (65)		M2003 -sym (12)		M2010 -sym (57)		M2010 -bench (59)	
	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time
Default	29,761.4	91.8	255.2	6.9	24,462.4	149.7	756.2	138.0	5736.7	246.0
S	29,761.4	91.9	255.2	6.9	21,586.0	128.7	729.7	126.4	5615.5	232.3
S-fund	13,381.5	63.4	225.2	6.6	20,063.4	123.2	792.5	139.9	5861.2	243.5
S-fundhlp	30,062.8	94.7	258.8	7.3	22,433.4	135.0	736.2	129.9	5665.3	234.1
S-nlp	29,761.4	92.0	255.2	6.9	22,410.7	135.1	765.8	133.2	5686.3	234.6
S-prob	29,761.4	91.7	255.2	6.9	22,079.6	130.2	766.2	127.0	5700.9	236.4
S-orbitmin	4001.0	15.1	115.1	3.5	19,737.3	122.9	769.5	128.7	5741.8	242.7

Table 14 Results of LP fixing variants: Given are the shifted geometric means of the number of branch-and-bound nodes (#nodes) and CPU time in seconds (time), the number of instances solved to optimality (#opt), the number of instances in which LP fixing was active (#act), the number of fixed variables, and the shifted geometric mean of the time used for fixing including symmetry computation (lpfix-time)

Setting	#nodes	Time	#opt	#act	#fixed	lpfix-time
M2003-sym (18)						
Default	81,955.4	568.43	13	0	–	0.00
lpfix	85,524.2	583.19	13	4	88	0.04
lpfix-OP	83,132.7	574.66	13	3	12	0.01
lpfix-fundnolp	85,570.8	583.39	13	4	88	0.04
lpfix-orbitmin	85,548.1	583.18	13	4	88	0.04
M2010-sym (154)						
Default	9061.6	1084.21	60	0	–	0.00
lpfix	8890.0	1065.66	61	23	10,676	0.27
lpfix-OP	8871.3	1062.16	61	22	10,619	0.06
lpfix-fundnolp	8921.1	1068.69	61	23	10,676	0.56
lpfix-orbitmin	8916.5	1068.32	61	23	10,676	0.56
M2010-bench (87)						
Default	16,375.3	580.00	62	0	–	0.00
lpfix	15,830.7	565.65	62	8	836	0.11
lpfix-OP	15,782.0	562.53	63	7	779	0.03
lpfix-fundnolp	15,944.2	566.51	63	8	836	0.29
lpfix-orbitmin	15,936.7	567.30	62	8	836	0.29

of the combination `lpfix-orbitmin` is worse than running `S-orbitmin` alone. One reason is that for `lpfix` CVR symmetries are needed. This in general decreases the size of the symmetry group and therefore the number of generated inequalities. However, as a stand-alone this fixing can essentially be performed without causing harm and might be able to improve the performance on instances with very large symmetries on continuous variables.

5.8 Experiment 8: comparison of the different methods

In order to get an overall picture, we compare the best symmetry handling variants in Tables 16 and 17. We obtain the following results:

`Margot1 ISP-NST-first` is the winner, closely followed by `OF-NST-first`. These two are the only variants that solve all instances. They are followed by `OB-NST`, `OB`, `ISP-NST`, `OF-NST`, `ISP`, and `OF`. However, all these variants are extremely fast in comparison to the default settings. In essence, these results prove the theoretical arguments of [38] that orbital branching performs similar to isomorphism pruning. However, the branching rule for isomorphism pruning plays a significant role for these instances and allows to change the ranking.

Note that the good performance of `ISP-NST-first` and `OF-NST-first` is specific to these instances: Randomly permuting the variables renders

Table 15 Comparison of different LP fixing variants on instances solved to optimality for all selected settings

Setting	M2003-sym (13)		M2010-sym (60)		M2010-bench (61)	
	#nodes	Time	#nodes	Time	#nodes	Time
Default	22,322.5	177.5	895.9	158.6	5979.0	263.1
lpfix	23,916.6	185.3	853.1	151.2	5707.3	253.8
lpfix-OP	22,857.8	180.7	856.2	150.8	5726.2	253.3
lpfix-fundnlp	23,916.6	185.4	861.0	152.4	5765.5	254.3
lpfix-orbitmin	23,916.6	185.2	861.0	152.2	5765.5	254.8

ISP-NST-first and OF-NST-first to be slower than ISP-NST. Orbital branching (OB-NST) then slows down as well, but is still somewhat faster than (ISP-NST).

Margot2. Here, the performance of all variants based on isomorphism pruning or orbital branching is quite close together and no variant dominates the others. The isomorphism variants solve the largest number of instances (69).

M2003-sym. Again the performance of the variants based on ISP or OB is quite close together and no variant dominates the others. All these variants (except OB) solve the same number of instances as the default. In terms of time, OB-NST is best with respect to Table 16 and for Table 17, S-orbitmin shows similar performance, but solves one instances less.

M2010-sym. Here, ISP-NST and OF-NST solve 70 instances. All other variants (except the orbital branching variants) still solve more instances than the default. On the instances solved to optimality, OF-NST performs best, followed by ISP-NST, S, OF, S-orbitmin, and ISP. This ranking is supported by the solving time diagram on the left of Fig. 2.

M2010-bench. Variants ISP-NST, and OF-NST solve the largest number of instances (67, i.e., five more than the default settings). In terms of solving time on the instances solved to optimality, S performs best (about 6% faster than the default in Table 17), closely followed by OF-NST, ISP, ISP-NST, and OF. Again, this evaluation is supported by the solving time diagram on the right of Fig. 2.

In total, we conclude that overall ISP-NST and OF-NST are very good variants that improve on the default settings. For the MIPLIB instances also S-orbitmin performs well, although it solves less instances in comparison. Moreover, on these instances, S-orbitmin is faster than S with respect to Table 16 and solves more instances.

5.9 Experiment 9: further instances

We now report on some further conclusions that can be drawn from the instances¹ used by Ghoniem and Sherali [16]. For each instance, a separate file additionally containing

¹ Available at <http://ahmed.ghoniem.info/download/symmetry.zip>.

Table 16 Comparison of different symmetry handling variants

Setting	#nodes	Time	#opt	#act	Method-time	Sym-time
Margot1 (16)						
Default	253, 743.3	1174.06	5	0	0.44	0.00
ISP	1844.1	55.58	15	16	18.47	0.02
ISP-NST	2208.4	48.31	15	16	3.96	0.02
ISP-NST-first	951.3	14.38	16	16	4.74	0.02
OB	2034.1	40.90	15	16	11.16	0.02
OB-NST	2131.0	36.70	15	16	0.92	0.02
OF	2192.5	61.70	15	16	16.26	0.02
OF-NST	2313.5	49.47	15	16	3.49	0.02
OF-NST-first	1089.9	17.78	16	16	6.00	0.02
S	253, 255.2	1173.85	5	16	0.56	0.00
S-orbitmin	103, 895.7	519.73	9	16	0.43	0.00
Margot2 (79)						
Default	969.7	32.64	65	0	0.10	0.00
ISP	517.4	21.75	69	59	5.05	0.00
ISP-NST	521.7	20.95	69	59	2.13	0.00
ISP-NST-first	589.8	21.00	69	59	4.53	0.00
OB	533.4	21.05	68	59	4.31	0.00
OB-NST	553.6	20.91	67	59	0.16	0.00
OF	520.1	21.56	68	59	4.26	0.00
OF-NST	518.0	21.03	68	59	1.89	0.00
OF-NST-first	591.9	21.03	68	59	3.97	0.00
S	970.0	32.71	65	77	0.14	0.00
S-orbitmin	593.2	24.09	68	77	0.13	0.00
M2003-sym (18)						
Default	81, 955.4	568.43	13	0	0.85	0.00
ISP	53, 800.6	526.18	13	16	14.24	0.03
ISP-NST	69, 514.8	524.47	13	16	6.33	0.03
OB	42, 699.9	553.40	12	16	49.40	0.03
OB-NST	50, 106.6	510.20	13	16	4.98	0.03
OF	52, 754.7	524.34	13	16	12.29	0.03
OF-NST	69, 554.7	523.52	13	16	5.59	0.03
S	74, 248.9	550.71	12	18	0.92	0.00
S-orbitmin	69, 156.1	537.14	12	18	0.81	0.00

Table 16 continued

Setting	#nodes	Time	#opt	#act	Method-time	Sym-time
M2010-sym (154)						
Default	9061.6	1084.21	60	0	0.39	0.00
ISP	3334.9	1004.48	68	141	86.16	1.22
ISP-NST	5090.2	925.03	70	141	25.98	1.21
OB	3493.9	1276.23	56	140	229.94	1.22
OB-NST	5168.0	1215.22	59	140	50.41	1.22
OF	3560.7	974.22	67	141	77.48	1.22
OF-NST	5301.0	917.86	70	141	22.94	1.22
S	8481.9	1055.29	60	153	3.65	0.00
S-orbitmin	8847.6	1039.67	63	153	1.78	0.00
M2010-bench (87)						
Default	16, 375.3	580.00	62	0	0.41	0.00
ISP	10, 589.4	495.80	66	36	9.39	0.52
ISP-NST	10, 778.4	467.62	67	36	4.43	0.52
OB	12, 040.2	663.83	58	36	25.35	0.52
OB-NST	13, 550.8	629.97	61	36	9.72	0.52
OF	10, 216.5	489.68	66	36	8.69	0.52
OF-NST	10, 978.2	469.00	67	36	4.15	0.52
S	15, 016.9	565.00	62	40	1.03	0.00
S-orbitmin	15, 658.7	556.63	65	40	0.96	0.00

the cutting planes discussed by Ghoniem and Sherali is available. Thus, a comparison to specific symmetry handling inequalities is possible. This testset contains three types of instances, which we discuss in turn.

The testset contains 10 instances from the “doubles tennis problem”, for which an unclassified group operates on a matrix. For these instances, adding the inequalities by Ghoniem and Sherali yields a speedup of about a factor of 30 compared to the default settings. These results can even be improved by using isomorphism pruning, orbital branching, and orbital fixing, yielding a speedup of about a factor of 5 compared to using these inequalities, see Table 18.

The 40 instances from the “noise dosage problem” have symmetric groups operating on a matrix, but do not contain binary variables. Thus, using our implementations of isomorphism pruning, orbital branching, and orbital fixing do not apply any reductions. Adding symmetry breaking inequalities slightly improves on the default settings, but is clearly dominated by using the inequalities in [16].

Finally, the 120 instances for “wagon load-balancing problem” have unclassified symmetry groups and contain mainly binary variables. The particular structure of these instances make all our symmetry handling methods extremely slow. Both computing stabilizers and lexicographic representatives requires a very large computation time and memory. Thus, the inequalities in [16] are clearly superior. Moreover, even if

Table 17 Comparison of different symmetry handling variants on instances solved to optimality for all selected settings

Setting	MargotL1 (14)		MargotL2 (65)		M2003-sym (12)		M2010-sym (49)		M2010-bench (53)	
	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time	#nodes	Time
Default	–	–	255.2	6.9	24,462.4	149.7	712.2	138.8	7184.2	288.2
ISP	838.4	30.8	113.8	4.0	17,540.4	129.9	523.9	138.8	6148.0	270.3
ISP-NST	961.4	25.0	113.2	3.7	17,540.4	129.1	535.0	122.5	6189.6	270.4
ISP-NST-first	492.9	6.9	120.1	3.6	–	–	–	–	–	–
OB	842.5	18.6	113.6	3.5	14,408.0	129.9	766.1	174.6	6782.4	284.4
OB-NST	881.9	16.2	116.8	3.5	14,439.5	122.6	843.7	161.7	6901.9	283.0
OF	953.9	32.6	111.4	3.9	17,579.0	129.0	526.3	127.5	6169.0	270.5
OF-NST	1026.8	26.1	111.6	3.7	17,579.0	128.6	528.9	120.3	6192.8	270.1
OF-NST-first	499.4	7.6	122.2	3.6	–	–	–	–	–	–
S	–	–	255.2	6.9	21,586.0	128.7	686.8	127.2	7029.2	269.9
S-orbitmin	–	–	115.1	3.5	19,737.3	122.9	710.0	129.3	6963.9	277.5




Fig. 2 Solving time diagram for selected variants on M2010-sym and M2010-bench: The x-axis depicts solving time, the y-axis shows the number of instances solved within the given time

Table 18 Comparison of different symmetry handling variants on the doubles tennis problem instances by Ghoniem and Sherahi (10 instances)

Setting	#nodes	Time	#opt
Cuts from [16]	689.4	12.0	10
Default	194, 205.5	357.5	7
ISP-NST	48.0	2.4	10
OF-NST	54.0	2.4	10
OB-NST	150.8	2.1	10

the local symmetry group is computed (which can be done efficiently), our methods fail to improve on the case where the inequalities of [16] are used. Obviously, these inequalities help to improve the lower bounds using further cutting planes, while the other techniques do not allow to do this.

We have to leave the development of general symmetry handling methods for such instances to future research.

6 Conclusion

The nine experiments above allow to reach the following conclusions:

- Isomorphism pruning and orbital fixing allow for a speed-up of two orders of magnitude for solving special instances like the instances of Margot, but also allow a decent speed-up on instances like the ones in the MIPLIB.
- Orbital branching also performs very well on the particular instances of Margot, but does not perform too well on the MIPLIB instances, mainly because the branching rule needs to be coordinated with the other rules.
- The instances of Margot, used both to test isomorphism pruning and orbital fixing in the literature, have a quite special structure. In particular, first index branching significantly helps isomorphism pruning on these instances, while it does not help the default.
- Adding Inequalities (3) for each orbit seems to be a quite good idea for general MIPLIB instances, but also adding (4) performs well. However, these two methods are usually slower than isomorphism pruning and orbital fixing. Moreover, these

methods are relatively sensitive to changes in the model and their behavior is not completely understood yet.

- The fixing of continuous variables slightly speeds up LP-solving, but does not result in a significant overall speed-up. It can nevertheless be applied, since it also does not create overhead and might help for very symmetric instances involving continuous variables.
- The performance of orbit probing is somewhat disappointing.
- Symmetry detection is not a bottleneck, while the computation of stabilizers and lexicographical tests can be.

This suggests the following future research:

- The performance of the symmetry handling methods should be improved by tuning parameters that automatically avoid symmetry handling if this is likely to not help or would take a long time. Moreover, methods to avoid very large computation times like on the instances of Ghoniem and Sherali should be developed if possible.
- The special structure of matrix actions of symmetric groups should be exploited to speed-up both the detection of symmetries and the computation of stabilizers and the like.
- Possibly orbit probing can be turned into an interesting method.
- It seems that particular inequalities can help to improve the dual bound – see the instances of Ghoniem and Sherali. More research is needed to find general purpose inequalities with the same behavior.

In conclusion, it seems that there are several ways in which the handling of symmetries might improve in the future, even though many techniques are already available and used in practice.

Acknowledgements We thank Tobias Achterberg for interesting discussions on the topic and Christopher Hojny for helpful comments. We also thank the editor and referees for their helpful comments that helped to improve this paper. Furthermore, the first author acknowledges support of the German Research Foundation (DFG) within the Collaborative Research Center 666.

Appendix: List of MIPLIB 2010 symmetries

The following Table 19 lists details about the symmetry groups of MIPLIB 2010 instances. The second column shows the logarithm to the base 10 of the order of the symmetry groups. The third column presents the percentage of variables that lie in an orbit of at least size two, i.e., variables on which the symmetry group acts non-trivially. The fourth column shows the groups which the symmetry group is a direct product of. The following notation is used:

- S_k denotes a symmetric group of degree k in coordinate action;
- $M(G, \ell)$ represents the matrix action of group G on ℓ points.
- “unknown” denotes groups whose type could not be determined by `PermLib`.

The computations for this table were performed on an Intel i7 CPU with 3.40 GHz and 32 GB of memory and a time limit of 10 h. For some instances, the computation or analysis ran into the memory or time limit. These instances are marked with “–”. Instances without symmetry are not shown.

Table 19 List of (possibly) symmetric MIPLIB 2010 instances before presolving. Additionally the \log_{10} of the order of their symmetry groups and the number of variables involved in a nontrivial symmetry is shown

Name	$\log_{10} G $	#vars	Factors
30_70_45_095_100	0.4	0.0	S_2
acc-tight4	4.7	97.8	1 Unknown
ash608gpia-3col	0.8	100.0	$M(S_3, 3651)$
atlanta-ip	4450.0	40.5	$(S_2)^{5898}, (S_3)^{1441}, (S_4)^{284}, (S_5)^{175}, (S_6)^{165}, (S_7)^{88}$
bab3	70.1	9.5	$(S_2)^{10}, (S_5)^{16}, (S_6)^4, (S_{10})^2, (M(S_2, 12328))^3,$ 4 Unknown
bab5	1.3	17.9	$(M(S_2, 968))^4$
beasleyC3	0.7	0.6	$(M(S_2, 8))^2$
biella1	217.3	11.7	$(S_2)^{261}, (S_3)^{30}, (S_4)^{20}, (S_5)^{13}, (S_6)^2, (S_7)^5, S_8,$ $(S_9)^2, (S_{10})^2, S_{11}$
blp-ar98	0.7	0.0	$(S_2)^2$
blp-ic97	0.7	0.0	$(S_2)^2$
bnatt350	503.7	43.3	$(S_2)^{42}, (S_3)^{81}, (S_4)^{88}, (S_5)^{76}, (S_6)^{39}, (S_7)^{10}$
bnatt400	588.0	43.6	$(S_2)^{50}, (S_3)^{83}, (S_4)^{98}, (S_5)^{91}, (S_6)^{42}, (S_7)^{16}, S_8$
circ10-3	1.4	100.0	1 Unknown
co-100	736.0	6.3	$(S_2)^{837}, (S_3)^{399}, (S_4)^3, (S_5)^3, S_6, (S_7)^2, S_{12}, S_{17},$ $S_{19}, S_{21}, S_{28}, S_{29}, S_{30}$
core2536-691	3.7	0.1	$(S_2)^5, S_5$
core4872-1529	68.9	2.0	$(S_2)^{211}, (S_3)^2, S_6, M(S_2, 8), M(S_2, 14), M(S_2, 28)$
cov1075	6.6	100.0	1 Unknown
datt256	3.7	0.0	$(S_2)^{12}$
dc1c	356.6	19.8	$(S_2)^{819}, (S_3)^{64}, (S_4)^{24}, (S_5)^2, (S_6)^8$
dc11	794.9	11.4	$(S_2)^{1640}, (S_3)^{174}, (S_4)^{69}, (S_5)^{18}, (S_6)^9, (S_7)^2$
dg012142	89.2	3.1	S_{64}
dolom1	468.2	19.4	$(S_2)^{907}, (S_3)^{71}, (S_4)^{25}, (S_5)^5, (S_7)^2, S_{15}, S_{33}, S_{34}$
ds-big	468.1	1.8	$(S_2)^{1516}, (S_3)^{15}$
eilB101	0.8	0.1	S_3
enlight13	0.4	92.3	$M(S_2, 312)$
enlight14	0.4	92.9	$M(S_2, 364)$
enlight15	0.4	93.3	$M(S_2, 420)$
enlight16	0.4	93.8	$M(S_2, 480)$
enlight9	0.4	88.9	$M(S_2, 144)$
ex1010-pi	1482.1	21.7	$(S_2)^{1324}, (S_3)^{366}, (S_4)^{148}, (S_5)^{70}, (S_6)^{32}, (S_7)^{27},$ $(S_8)^{20}, (S_9)^{13}, (S_{10})^3, (S_{11})^5, S_{12}, (S_{16})^2$
ex10	1.0	100.0	1 Unknown
ex9	30.2	100.0	1 Unknown
glass4	0.4	0.6	S_2

Table 19 continued

Name	$\log_{10} G $	#vars	Factors
gmu-35-40	801.6	39.3	$(S_2)^{18}, (S_3)^{16}, (S_4)^3, (S_5)^3, S_{363}$
gmu-35-50	1837.7	44.5	$(S_2)^{18}, (S_3)^{16}, (S_4)^3, (S_5)^3, S_{742}$
gmut-75-50	–	–	–
gmut-77-40	40, 582.7	47.2	$(S_2)^{12}, (S_3)^7, (S_4)^{20}, (S_5)^{31}, S_{11198}$
go19	3.5	99.8	$S_4, 1$ Unknown
iis-bupa-cov	2.2	2.0	S_3, S_4
iis-pima-cov	35.5	4.2	S_{32}
lectsched-1	624.4	12.4	$(S_2)^3, (S_3)^3, S_{193}, (M(S_2, 4))^{797}, (M(S_3, 6))^{18}, (M(S_4, 8))^6$
lectsched-1-obj	624.4	12.4	$(S_2)^3, (S_3)^3, S_{193}, (M(S_2, 4))^{797}, (M(S_3, 6))^{18}, (M(S_4, 8))^6$
lectsched-2	408.6	11.7	$(S_2)^3, S_{148}, (M(S_2, 4))^{462}, (M(S_3, 6))^6, (M(S_4, 8))^4$
lectsched-3	543.9	10.9	$(S_2)^3, (S_3)^3, S_{184}, (M(S_2, 4))^{633}, (M(S_3, 6))^{15}$
lectsched-4-obj	217.1	13.2	$S_{93}, (M(S_2, 4))^{230}, (M(S_3, 6))^3, M(S_4, 8)$
liu	0.4	0.2	S_2
macrophage	61.2	25.0	$(S_2)^{146}, (M(S_2, 4))^{15}, (M(S_2, 6))^9, (M(S_2, 12))^2, M(S_2, 14), M(S_4, 24), M(S_6, 48), M(S_2 \times S_2, 14), 2$ Unknown
map06	–	–	–
map10	–	–	–
map14	–	–	–
map18	–	–	–
map20	–	–	–
maxgasflow	5.5	1.8	$(M(S_2, 6))^{11}, M(S_2, 12), M(S_2, 24), M(S_3, 12), M(S_3, 18)$
mcsched	4.6	5.2	$(M(S_2, 6))^{15}$
methanosarcina	762.7	64.7	$(S_2)^{2505}, M(S_2 \wr S_2, 28), 4$ Unknown
mkc	77.2	61.3	$(S_2)^9, (S_3)^4, (M(S_2, 4))^{15}, (M(S_2, 6))^3, (M(S_2, 8))^2, M(S_2, 10), (M(S_2, 12))^2, M(S_2, 24), M(S_2, 28), (M(S_3, 6))^6, M(S_3, 9), (M(S_3, 12))^2, M(S_3, 18), (M(S_3, 24))^2, M(S_3, 42), (M(S_4, 8))^4, M(S_4, 48), (M(S_5, 10))^2, M(S_6, 60), M(S_7, 14), M(S_8, 16), 5$ Unknown
momentum3	36.7	0.9	$(M(S_3, 6))^2, M(S_5, 10), M(S_6, 12), M(S_7, 14), (M(S_9, 18))^3, M(S_{13}, 26)$
msc98-ip	2673.2	33.2	$(S_2)^{609}, (S_3)^{266}, (S_4)^{219}, (S_5)^{198}, (S_6)^{210}, (S_7)^{108}, (S_8)^{112}, (S_9)^2, (S_{10})^6, (M(S_2, 14))^2, M(S_2, 24), M(S_2, 26), M(S_2, 34), M(S_2, 38), S_2 \wr S_2$
mzzv11	46.7	3.0	$(S_2)^{155}$

Table 19 continued

Name	$\log_{10} G $	#vars	Factors
n3seq24	9.9	11.4	$M(S_2, 48), (M(S_2, 50))^4, (M(S_2, 54))^2, (M(S_2, 62))^4, M(S_2, 72), (M(S_2, 90))^4, (M(S_2, 110))^2, M(S_2, 142), (M(S_2, 180))^2, M(S_2, 296), (M(S_2, 490))^2, M(S_2, 798), M(S_2, 1018), M(S_2, 3878), M(S_2, 4772), M(S_4, 180)$
nag	0.4	0.1	S_2
neos-1109824	1.7	100.0	1 Unknown
neos-1171692	19.8	100.0	$M(S_{21}, 1638)$
neos-1171737	32.5	100.0	$M(S_{30}, 2340)$
neos-1224597	504.2	100.0	$(S_5)^{17}, (S_{10})^3, S_{35}, 2 \text{ Unknown}$
neos-1311124	78.9	100.0	$M(S_{21}, 84), (M(S_{21}, 168))^2, M(S_{21}, 672)$
neos-1337307	3.8	87.5	$M(S_7, 2485)$
neos-1396125	0.8	78.3	$M(S_3, 909)$
neos13	0.4	0.1	S_2
neos-1426635	26.3	100.0	$M(S_{10}, 40), (M(S_{10}, 80))^2, M(S_{10}, 320)$
neos-1426662	53.4	100.0	$M(S_{16}, 64), (M(S_{16}, 128))^2, M(S_{16}, 512)$
neos-1429212	-	-	-
neos-1436709	39.2	100.0	$M(S_{13}, 52), (M(S_{13}, 104))^2, M(S_{13}, 416)$
neos-1440460	22.3	100.0	$M(S_9, 36), (M(S_9, 72))^2, M(S_9, 288)$
neos-1442119	43.8	100.0	$M(S_{14}, 56), (M(S_{14}, 112))^2, M(S_{14}, 448)$
neos-1442657	34.8	100.0	$M(S_{12}, 48), (M(S_{12}, 96))^2, M(S_{12}, 384)$
neos-1601936	432.5	9.2	$S_{72}, 1 \text{ Unknown}$
neos-1605061	103.8	1.8	S_{72}
neos-1605075	103.8	1.7	S_{72}
neos-1620770	3.8	97.2	1 Unknown
neos18	247.9	36.1	$(S_2)^6, (S_3)^2, (S_4)^2, (M(S_2, 4))^{28}, (M(S_2, 8))^3, (M(S_3, 6))^2, (M(S_3, 12))^3, (M(S_4, 8))^{19}, (M(S_4, 16))^3, M(S_6, 12), M(S_8, 16), (S_2 \wr S_2)^2, (S_2 \wr S_9)^4, (S_2 \wr S_{11})^2, M(S_2 \wr S_{10}, 40), M(S_2 \wr S_6, 24), M(S_2 \wr S_7, 28), M(S_2 \wr S_6, 24), M(S_2 \wr S_9, 36), M(S_2 \wr S_9, 36), M(S_2 \wr S_{11}, 44), M(S_2 \wr S_8, 32), M(S_2 \wr S_7, 28), M(S_2 \wr S_{10}, 40), M(S_2 \wr S_3, 12), M(S_2 \wr S_{10}, 40), M(S_2 \wr S_{13}, 52), M(S_2 \wr S_2, 8), M(S_2 \wr S_2, 8), M(S_2 \wr S_5, 20), M(S_2 \wr S_3, 12), M(S_2 \wr S_3, 12), M(S_2 \wr S_{10}, 40), M(S_2 \wr S_4, 16), M(S_2 \wr S_3, 12), M(S_2 \wr S_9, 36), M(S_2 \wr S_2, 8), M(S_2 \wr S_2, 8), M(S_2 \wr S_2, 8), M(S_2 \wr S_2, 8)$
neos-476283	39.9	1.1	$S_7, S_{27}, (M(S_2, 4))^{10}, (M(S_2, 24))^2, S_5 \wr S_2$
neos-555424	146.4	99.9	$M(S_{10}, 20), (M(S_{10}, 300))^2, M(S_{20}, 40), M(S_{30}, 60), M(S_{10} \wr S_3, 90), M(S_{10} \wr S_3, 90), 1 \text{ Unknown}$
neos-631710	313.7	99.7	$M(S_{11}, 6094), M(S_{15}, 8325), M(S_{15}, 8340), M(S_{16}, 8864), M(S_{18}, 9990), M(S_{19}, 10564), M(S_{24}, 13344), M(S_{28}, 15540), M(S_{33}, 18282), M(S_{39}, 21645), M(S_{40}, 22160), M(S_{42}, 23352)$

Table 19 continued

Name	$\log_{10} G $	#vars	Factors
neos6	6.6	94.9	$M(\mathcal{S}_{10}, 8340)$
neos-738098	21.0	93.5	1 Unknown
neos-777800	1.0	100.0	1 Unknown
neos-785912	9.8	91.3	1 Unknown
neos788725	0.7	100.0	1 Unknown
neos808444	0.4	11.5	$M(\mathcal{S}_2, 2288)$
neos-820146	106.2	100.0	1 Unknown
neos-820157	557.8	100.0	1 Unknown
neos-824661	936.3	100.0	1 Unknown
neos-824695	936.3	100.0	1 Unknown
neos-826650	475.2	98.6	2 Unknown
neos-826694	973.0	99.3	2 Unknown
neos-826812	141.1	99.3	1 Unknown
neos-826841	41.5	98.5	1 Unknown
neos-849702	3.7	100.0	1 Unknown
neos-859770	769.1	98.9	$(\mathcal{S}_2)^{11}, \mathcal{S}_6, \mathcal{S}_{12}, \mathcal{S}_{20}, \mathcal{S}_{25}, \mathcal{S}_{32}, 1$ Unknown
neos-885086	56.1	100.0	$M(\mathcal{S}_{45}, 4860)$
neos-885524	–	–	–
neos-911880	7.6	100.0	$(M(\mathcal{S}_3, 111))^6, M(\mathcal{S}_6, 222)$
neos-932816	598.6	85.3	$\mathcal{S}_{293}, M(\mathcal{S}_4, 17628)$
neos-933638	1184.6	93.9	$\mathcal{S}_{42}, \mathcal{S}_{495}, 3$ Unknown
neos-933966	1183.6	95.7	$\mathcal{S}_{42}, \mathcal{S}_{495}, M(\mathcal{S}_2, 1218), M(\mathcal{S}_3, 1827), 3$ Unknown
neos-934278	1709.5	81.4	$\mathcal{S}_{389}, \mathcal{S}_{396}, M(\mathcal{S}_2, 1056), 5$ Unknown
neos-935627	5817.1	75.9	$\mathcal{S}_{2022}, M(\mathcal{S}_2, 236), 5$ Unknown
neos-935769	5817.5	80.9	$\mathcal{S}_{2022}, M(\mathcal{S}_3, 354), M(\mathcal{S}_3, 360), M(\mathcal{S}_3, 387),$ $M(\mathcal{S}_3, 405), 5$ Unknown
neos-937511	5261.9	79.7	$\mathcal{S}_{1853}, M(\mathcal{S}_3, 408), M(\mathcal{S}_3, 423), M(\mathcal{S}_3, 465),$ $M(\mathcal{S}_3, 480), 5$ Unknown
neos-937815	5335.7	72.6	$\mathcal{S}_{1876}, M(\mathcal{S}_2, 264), M(\mathcal{S}_2, 266), M(\mathcal{S}_2, 276),$ 5 Unknown
neos-941262	5570.5	76.3	$\mathcal{S}_{1948}, M(\mathcal{S}_2, 170), M(\mathcal{S}_2, 204), M(\mathcal{S}_2, 220),$ $M(\mathcal{S}_2, 248), (M(\mathcal{S}_2, 282))^2, M(\mathcal{S}_2, 292), M(\mathcal{S}_2, 298),$ $M(\mathcal{S}_2, 310), 4$ Unknown
neos-941313	–	–	–
neos-948126	4969.9	78.9	$\mathcal{S}_{1764}, M(\mathcal{S}_2, 204), M(\mathcal{S}_2, 222), M(\mathcal{S}_2, 234),$ $M(\mathcal{S}_2, 238), M(\mathcal{S}_2, 282), M(\mathcal{S}_2, 298), 5$ Unknown
neos-952987	16.5	0.3	$(\mathcal{S}_2)^{36}, \mathcal{S}_9$
neos-957389	33.8	75.1	$M(\mathcal{S}_2, 126), (M(\mathcal{S}_4, 280))^2, M(\mathcal{S}_{10}, 690),$ $M(\mathcal{S}_{10}, 700), 1$ Unknown

Table 19 continued

Name	$\log_{10} G $	#vars	Factors
neos-984165	4280.8	75.8	$(S_2)^{11}$, S_{1549} , $M(S_2, 152)$, $M(S_2, 206)$, $M(S_2, 218)$, $M(S_2, 222)$, $M(S_2, 240)$, $M(S_2, 260)$, $M(S_2, 266)$, $M(S_2, 286)$, 4 Unknown
noswot	0.4	40.6	$M(S_2, 52)$
ns1111636	–	–	–
ns1116954	553.4	99.9	2 Unknown
ns1158817	–	–	–
ns1208400	2.8	97.8	1 Unknown
ns1456591	36.8	99.8	S_{20} , $M(S_{20}, 8360)$
ns1606230	103.8	1.7	S_{72}
ns1631475	31.7	0.9	$(S_2)^{105}$
ns1702808	2.9	100.0	$M(S_6, 804)$
ns1758913	0.4	3.0	$M(S_2, 540)$
ns1830653	15.9	1.1	S_{18}
ns1853823	0.7	56.1	1 Unknown
ns1854840	2.0	88.0	1 Unknown
ns1856153	0.4	92.5	$M(S_2, 11102)$
ns1905797	1.4	100.0	$M(S_4, 18192)$
ns1905800	1.1	100.0	1 Unknown
ns1952667	69.7	3.4	$(S_2)^{222}$, S_3 , S_5
ns2081729	1.0	90.8	$(M(S_2, 200))^3$
ns2118727	1161.2	0.5	$(S_2)^{20}$, $(S_3)^{45}$, $(S_4)^2$, $(S_5)^{10}$, S_{235} , S_{308} , $(M(S_2, 4))^{18}$
ns2124243	–	–	–
ns2137859	755.8	96.0	S_{321} , 1 Unknown
ns894236	8.7	2.4	$M(S_{12}, 228)$
ns903616	17.2	2.3	$M(S_4, 92)$, $M(S_{18}, 414)$
nsr8k	100.5	1.2	$(S_2)^{198}$, $(S_3)^7$, S_{32}
p2m2p1m1p0n100	32.1	92.0	$(S_2)^6$, $(S_3)^8$, $(S_4)^4$, $(S_5)^3$, $(S_6)^3$, S_7
p6b	1.4	100.0	1 Unknown
pigeon-10	119.0	93.5	S_{30} , S_{60} , 1 Unknown
pigeon-11	135.3	94.4	S_{33} , S_{66} , 1 Unknown
pigeon-12	152.0	95.2	S_{36} , S_{72} , 1 Unknown
pigeon-13	169.0	95.8	S_{39} , S_{78} , 1 Unknown
pigeon-19	277.6	97.8	S_{57} , S_{114} , 1 Unknown
protfold	0.7	98.1	1 Unknown
pw-myciel4	1.1	86.9	1 Unknown
qju	1.7	100.0	1 Unknown
queens-30	1.0	100.0	1 Unknown

Table 19 continued

Name	$\log_{10} G $	#vars	Factors
rail01	2170.5	5.6	$(S_2)^{931}, (S_3)^{579}, (S_4)^{204}, (S_5)^{170}, (S_6)^{28}, (S_7)^{77}, (S_8)^4, (S_{11})^3, (S_{12})^6, (S_{13})^6, (S_{14})^9, (S_{15})^9, (S_{16})^6$
rail02	–	–	–
rail03	–	–	–
rail507	256.5	2.6	$(S_2)^{788}, (S_3)^{21}, S_6$
ramos3	9.2	100.0	1 Unknown
rococoB10-011000	56.1	1.0	S_{45}
rococoC10-001000	62.8	4.1	$(S_2)^{44}, S_{41}$
rococoC11-011100	73.2	0.8	S_{55}
rococoC12-111000	411.3	8.5	$(S_5)^{34}, (S_6)^{62}, (S_9)^{14}, S_{62}$
rvb-sub	31.1	0.6	$(S_2)^{103}$
satellites1-25	60.3	4.4	$(S_2)^{200}$
satellites2-60-fs	607.0	5.1	$(S_2)^6, 125$ Unknown
sct1	3336.1	64.1	$(S_2)^4, (S_3)^{211}, (S_4)^{594}, (S_5)^{17}, (S_6)^{121}, (S_7)^5, S_{471}, (M(S_2, 4))^{327}, (M(S_4, 8))^{23}, 9$ Unknown
sct32	397.0	49.1	$(S_2)^{150}, (S_3)^{15}, S_4, (M(S_2, 34))^{11}, (M(S_3, 6))^{106}, (M(S_3, 51))^4, (M(S_4, 8))^8, (M(S_4, 68))^4, M(S_5, 10), (M(S_5, 85))^3, (M(S_6, 12))^{13}, (M(S_6, 102))^2, (M(S_7, 14))^3, M(S_7, 119), M(S_{24}, 408), 27$ Unknown
sct5	3063.1	76.5	$(S_2)^{314}, S_4, (S_5)^{65}, (S_6)^{64}, (S_7)^9, (S_8)^6, (S_9)^{10}, (S_{10})^{13}, (M(S_2, 4))^4, (M(S_3, 6))^3, (M(S_4, 8))^{161}, (M(S_5, 10))^7, (M(S_6, 12))^3, (M(S_7, 14))^3, (M(S_8, 16))^5, (M(S_9, 18))^3, 10$ Unknown
seymour-disj-10	18.6	8.8	$(S_2)^{41}, (S_3)^4, (S_4)^2, M(S_2, 4)$
seymour	234.5	19.9	$(S_2)^{43}, (S_3)^4, (S_4)^3, S_6, S_{117}, M(S_2, 4), S_2 \wr S_{15}, S_3 \wr S_2$
shipsched	0.4	0.5	$M(S_2, 70)$
shs1023	0.4	13.5	$M(S_2, 60198)$
siena1	146.5	4.7	$(S_2)^{271}, (S_3)^{14}, (S_5)^2, S_7, S_{11}, S_{34}$
sing161	0.7	10.8	$M(S_2, 27388), M(S_2, 56024)$
sing245	0.4	5.3	$M(S_2, 12554)$
sing2	0.4	13.7	$M(S_2, 4332)$
stp3d	178.9	1.6	$(M(S_2, 4))^{382}, (M(S_2, 6))^{72}, (M(S_2, 8))^{116}, (M(S_2, 12))^{24}$
sts405	7.3	100.0	1 Unknown
sts729	19.9	100.0	1 Unknown
swath	816.6	7.1	$(S_4)^{20}, S_{83}, S_{320}$
t1717	–	–	–
t1722	–	–	–
tanglegram1	–	–	–

Table 19 continued

Name	$\log_{10} G $	#vars	Factors
tanglegram2	6073.1	95.5	$(S_2)^{45}, (S_3)^{26}, (S_4)^{25}, (S_5)^{11}, (S_6)^{10}, (S_7)^3, (S_8)^9, (S_9)^9, (S_{10})^8, (S_{11})^6, (S_{12})^7, (S_{14})^4, (S_{15})^{10}, (S_{16})^3, S_{19}, (S_{20})^6, (S_{22})^7, S_{25}, (S_{27})^2, S_{28}, S_{30}, (S_{31})^2, (S_{36})^2, S_{40}, S_{41}, S_{44}, (S_{48})^3, S_{49}, S_{50}, S_{55}, S_{59}, (S_{60})^2, S_{89}, S_{98}, S_{108}, S_{112}, S_{130}, S_{154}, S_{193}, S_{236}, S_{287}, S_{297}, S_{425}, (M(S_2, 4))^3, (M(S_3, 6))^2, M(S_5, 10), M(S_2 \wr S_6, 24), 4 \text{ Unknown}$
timtab1	0.4	0.5	S_2
toll-like	140.2	37.8	$(S_2)^{230}, (M(S_2, 4))^{10}, (M(S_2, 6))^2, (M(S_2, 8))^2, (M(S_2, 10))^2, M(S_2, 12), M(S_2, 30), M(S_5, 30), 9 \text{ Unknown}$
uc-case11	0.8	5.9	$M(S_3, 2016)$
uc-case3	0.7	14.2	$M(S_2, 1008), M(S_2, 4366)$
uct-subprob	29.1	14.0	$(M(S_2, 4))^{58}, M(S_2, 14), (M(S_3, 6))^8, M(S_4, 8), M(S_7, 14)$
unitcal_7	304.7	52.3	$(S_2)^{672}, (M(S_2, 3030))^3, 1 \text{ Unknown}$
van	16, 059.6	39.5	S_{4928}
vpphard2	–	–	–
vpphard	–	–	–
wachplan	1.9	96.6	1 Unknown
zib54-UUE	0.4	2.4	$M(S_2, 126)$

References

1. Achterberg, T.: Constraint Integer Programming. Ph.D. Thesis, TU Berlin (2007)
2. Achterberg, T.: SCIP: solving constraint integer programs. *Math. Program. Comput.* **1**(1), 1–41 (2009)
3. Achterberg, T., Koch, T., Martin, A.: MIPLIB 2003. *Oper. Res. Lett.* **34**(4), 361–372 (2006)
4. Atamtürk, A., Nemhauser, G.L., Savelsbergh, M.W.P.: Conflict graphs in integer programming. *Eur. J. Oper. Res.* **121**, 40–55 (2000)
5. Berthold, T.: Heuristic Algorithms in Global MINLP Solvers. Ph.D. Thesis, TU Berlin (2014)
6. Berthold, T., Pfetsch, M.E.: Detecting orbitopal symmetries. In: Fleischmann, B., Borgwardt, K.H., Klein, R., Tuma, A. (eds.) *Operations Research Proceedings 2008*, pp. 433–438. Springer, Berlin (2009)
7. Bödi, R., Herr, K., Joswig, M.: Algorithms for highly symmetric linear and integer programs. *Math. Program.* **137**(1–2), 65–90 (2013)
8. Bremner, D., Pasechnik, D.V., Rehn, T., Schürmann, A., Dutour Sikirić, D.: Computing symmetry groups of polyhedra. *LMS J. Comput. Math.* **17**(1), 565–581 (2014)
9. Christophel, P.M., Güzelsoy, M., Pólik, I.: New symmetries in mixed-integer linear optimization symmetry heuristics and complement-based symmetries. In: Technical Report, Optimization Online. http://www.optimization-online.org/DB_HTML/2014/07/4466.html (2014). Accessed June 2018
10. Darga, P.T., Katebi, H., Liffiton, M., Markov, I., Sakallah, K.: Saucy. <http://vlsicad.eecs.umich.edu/BK/SAUCY/> (2012). Accessed June 2018
11. Faenza, Y., Kaibel, V.: Extended formulations for packing and partitioning orbitopes. *Math. Oper. Res.* **34**(3), 686–697 (2009)
12. Fischetti, M., Libertini, L.: Orbital shrinking. In: Mahjoub, A., Markakis, V., Milis, I., Paschos, V.T. (eds.) *Combinatorial Optimization, Lecture Notes in Computer Science*, vol. 7422, pp. 48–58. Springer, Berlin (2012)

13. Friedman, E.J.: Fundamental domains for integer programs with symmetries. In: *Combinatorial Optimization and Applications*, volume 4616 of *Lecture Notes in Computer Science*, pp. 146–153. Springer, Berlin (2007)
14. GAP—Groups, Algorithms, Programming: A system for computational discrete algebra. Version 4.5. <http://www.gap-system.org/> (2013). Accessed June 2018
15. Gatermann, K., Parrilo, P.: Symmetry groups, semidefinite programs, and sums of squares. *J. Pure Appl. Algebra* **192**(1–3), 95–128 (2004)
16. Ghoniem, A., Sherali, H.D.: Defeating symmetry in combinatorial optimization via objective perturbations and hierarchical constraints. *IIE Trans.* **43**, 575–588 (2011)
17. Herr, K.: *Core Sets and Symmetric Convex Optimization*. Ph.D. Thesis, TU Darmstadt (2013)
18. Herr, K., Rehn, T., Schürmann, A.: Exploiting symmetry in integer convex optimization using core points. *Oper. Res. Lett.* **41**(3), 298–304 (2013)
19. Herr, K., Rehn, T., Schürmann, A.: On lattice-free orbit polytopes. *Discrete Comput. Geom.* **53**(1), 144–172 (2015)
20. Johnson, D.S.: The NP-completeness column. *ACM Trans. Algorithms* **1**(1), 160–176 (2005)
21. Junttila, T., Kaski, P.: bliss: A tool for computing automorphism groups and canonical labelings of graphs. <http://www.tcs.hut.fi/Software/bliss/> (2012). Accessed June 2018
22. Kaibel, V., Pfetsch, M.E.: Packing and partitioning orbitopes. *Math. Program.* **114**, 1–36 (2008)
23. Kaibel, V., Peinhardt, M., Pfetsch, M.E.: Orbitopal fixing. *Discrete Optim.* **8**(4), 595–610 (2011)
24. Koch, T., Achterberg, T., Andersen, E., Bastert, O., Berthold, T., Bixby, R.E., Danna, E., Gamrath, G., Gleixner, A.M., Heinz, S., Lodi, A., Mittelmann, H., Ralphs, T., Salvagnin, D., Steffy, D.E., Wolter, K.: *MIPLIB 2010: mixed integer programming library version 5*. *Math. Program. Comput.* **3**(2), 103–163 (2011)
25. Lang, S.: *Algebra*, 3rd edn. Springer, New York (2005)
26. Liberti, L.: Automatic generation of symmetry-breaking constraints. In: *Combinatorial Optimization and Applications*, volume 5165 of *Lecture Notes in Computer Science*, pp. 328–338. Springer, Berlin (2008)
27. Liberti, L.: Reformulations in mathematical programming: automatic symmetry detection and exploitation. *Math. Program.* **131**(1–2), 273–304 (2012)
28. Liberti, L., Ostrowski, J.: Stabilizer-based symmetry breaking constraints for mathematical programs. *J. Glob. Optim.* **60**, 183–194 (2014)
29. Margot, F.: Pruning by isomorphism in branch-and-cut. *Math. Program.* **94**, 71–90 (2002)
30. Margot, F.: Exploiting orbits in symmetric ILP. *Math. Program.* **98**(1–3), 3–21 (2003)
31. Margot, F.: Small covering designs by branch-and-cut. *Math. Program.* **94**(2–3), 207–220 (2003)
32. Margot, F.: Symmetric ILP: Coloring and small integers. *Discrete Optim.* **4**(1), 40–62 (2007)
33. Margot, F.: Symmetry in integer linear programming. In: Jünger, M., Liebling, T., Naddef, D., Nemhauser, G.L., Pulleyblank, W., Reinelt, G., Rinaldi, G., Wolsey, L. (eds.) *50 Years of Integer Programming 1958–2008*, chapter 17, pp. 647–681. Springer, Berlin (2010)
34. McKay, B.D.: The nauty program. <http://cs.anu.edu.au/people/bdm/nauty/> (2012). Accessed June 2018
35. Mittelmann, H.D., Salvagnin, D.: On solving a hard quadratic 3-dimensional assignment problem. *Math. Program. Comput.* **7**, 219–234 (2015)
36. Ostrowski, J., Linderoth, J., Rossi, F., Smriglio, S.: Orbital branching. In: *Integer Programming and Combinatorial Optimization*, volume 4513 of *Lecture Notes in Computer Science*, pp. 104–118. Springer, Berlin (2007)
37. Ostrowski, J.: *Symmetry in Integer Programming*. Ph.D. Thesis, Lehigh University (2009)
38. Ostrowski, J., Linderoth, J., Rossi, F., Smriglio, S.: Orbital branching. *Math. Program.* **126**(1), 147–178 (2011)
39. Padberg, M.W.: Facets and rank of integer polyhedra. In: Jünger, M., Reinelt, G. (eds.) *Facets of Combinatorial Optimization*, pp. 23–58. Springer, Berlin (2013)
40. Puget, J.-F.: Automatic detection of variable and value symmetries. In: Beek, P. (ed.) *Principles and Practice of Constraint Programming—CP 2005*. *Lecture Notes in Computer Science*, vol. 3709, pp. 475–489. Springer, Berlin (2005)
41. Read, R.C., Corneli, D.G.: The graph isomorphism disease. *J. Graph Theory* **1**(4), 339–363 (1977)
42. Rehn, T.: *Exploring Core Points for Fun and Profit: A Study Of Lattice-free Orbit Polytopes*. Ph.D. Thesis, University of Rostock (2014)
43. Rehn, T.: *PerMLib: Permutation Computation Library*. <http://www.geometrie.uni-rostock.de/software/> (2013). Accessed June 2018

44. Salvagnin, D.: A Dominance Procedure For Integer Programming. Master's Thesis, University of Padova (2005)
45. Salvagnin, D.: Orbital shrinking: a new tool for hybrid MIP/CP methods. In: Gomes, C., Sellmann, M. (eds.) Proceedings of Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CPAIOR), pp. 204–215. Springer, Berlin, Heidelberg (2013)
46. Savelsbergh, M.W.P.: Preprocessing and probing techniques for mixed integer programming problems. *ORSA J. Comput.* **6**(4), 445–454 (1994)
47. SCIP–Solving Constraint Integer Programs: <http://scip.zib.de> (2015). Accessed June 2018
48. Seress, Á.: Permutation Group Algorithms. Cambridge University Press, Cambridge (2003)
49. Sherali, H., Smith, J.C.: Improving discrete model representations via symmetry considerations. *Manag. Sci.* **47**(10), 1396–1407 (2001)