

Orbital branching

James Ostrowski · Jeff Linderoth ·
Fabrizio Rossi · Stefano Smriglio

Received: 9 October 2007 / Accepted: 3 February 2009
© Springer and Mathematical Programming Society 2009

Abstract We introduce *orbital branching*, an effective branching method for integer programs containing a great deal of symmetry. The method is based on computing groups of variables that are equivalent with respect to the symmetry remaining in the problem after branching, including symmetry that is not present at the root node. These groups of equivalent variables, called orbits, are used to create a valid partitioning of the feasible region that significantly reduces the effects of symmetry while still allowing a flexible branching rule. We also show how to exploit the symmetries present in the problem to fix variables throughout the branch-and-bound tree. Orbital branching can easily be incorporated into standard integer programming software. Through an empirical study on a test suite of symmetric integer programs, the question as to the most effective orbit on which to base the branching decision is investigated. The resulting method is shown to be quite competitive with a similar method known as *isomorphism pruning* and significantly better than a state-of-the-art commercial solver on symmetric integer programs.

Keywords Integer programming · Symmetry · Branch-and-bound algorithms

J. Ostrowski
Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, USA
e-mail: jao204@lehigh.edu

J. Linderoth (✉)
Department of Industrial and Systems Engineering, University of Wisconsin-Madison, Madison, USA
e-mail: linderoth@wisc.edu

F. Rossi · S. Smriglio
Dipartimento di Informatica, Università di L'Aquila, L'Aquila, Italy
e-mail: fabrizio.rossi@univaq.it

S. Smriglio
e-mail: stefano.smriglio@univaq.it

Mathematics Subject Classification (2000) 90C10

1 Introduction

In this work, we focus on packing and covering integer programs (IP)s of the form

$$\max_{x \in \{0,1\}^n} \left\{ e^T x \mid Ax \leq e \right\} \quad \text{and} \quad (\text{PIP})$$

$$\min_{x \in \{0,1\}^n} \left\{ e^T x \mid Ax \geq e \right\}, \quad (\text{CIP})$$

where $A \in \{0, 1\}^{m \times n}$, and e is a vector of ones of conformal size. Our particular focus is on cases when (CIP) or (PIP) is highly-symmetric, a concept we formalize as follows. Let Π^n be the set of all permutations of $I^n = \{1, \dots, n\}$. Given a permutation $\pi \in \Pi^n$ and a permutation $\sigma \in \Pi^m$, let $A(\sigma, \pi)$ be the matrix obtained by permuting the rows of A by σ and the columns of A by π , i.e., $A(\sigma, \pi) = P_\sigma A P_\pi$, where P_σ and P_π are the permutation matrices corresponding to σ and π , respectively. The *symmetry group* \mathcal{G} of the matrix A is the set of permutations

$$\mathcal{G}(A) \stackrel{\text{def}}{=} \{ \pi \in \Pi^n \mid \exists \sigma \in \Pi^m \text{ such that } A(\sigma, \pi) = A \}.$$

So, for any $\pi \in \mathcal{G}(A)$, if \hat{x} is feasible for (CIP) or (PIP) (or the LP relaxations of (CIP) or (PIP)), then if the permutation π is applied to the coordinates of \hat{x} , the resulting solution, which we denote as $\pi(\hat{x})$, is also feasible. Moreover, the solutions \hat{x} and $\pi(\hat{x})$ have equal objective value.

This equivalence of solutions induced by symmetry is a major factor that might confound the branch-and-bound process. For example, suppose \hat{x} is a (non-integral) solution to an LP relaxation of PIP or CIP, with $0 < \hat{x}_j < 1$, and the decision is made to branch down on variable x_j by fixing $x_j = 0$. If $\exists \pi \in \mathcal{G}(A)$ such that $[\pi(\hat{x})]_j = 0$, then $\pi(\hat{x})$ is a feasible solution for this child node, and $e^T \hat{x} = e^T (\pi(\hat{x}))$, so the relaxation value for the child node will not change. If the cardinality of $\mathcal{G}(A)$ is large, then there are many permutations through which the parent solution of the relaxation can be preserved in this manner, resulting in many branches that do not change the bound on the parent node. Furthermore, symmetric solutions appear again and again all over the tree. Symmetry has long been recognized as a curse for solving integer programs, and auxiliary (often extended) formulations are often sought that reduce the amount of symmetry in an IP formulation [2, 9, 19]. In addition, there is a body of research on valid inequalities that can help exclude symmetric feasible solutions [14, 26]. Kaibel and Pfetsch [11] formalize many of these arguments by defining and studying the properties of a polyhedron known as an orbitope, the convex hull of lexicographically maximal solutions with respect to a symmetry group. Kaibel et al. [10] then use the properties of orbitopes to remove symmetry in partitioning problems.

A different idea, *isomorphism pruning*, introduced by Margot [15, 16] in the context of IP, examines the symmetry group of the problem in order to prune isomorphic

subproblems of the enumeration tree. The branching method introduced in this work, *orbital branching*, also uses the symmetry group of the problem. However, instead of examining this group to ensure that only one node in the equivalence class of the group will be evaluated, the group is used to guide the branching decision. At the cost of potentially evaluating isomorphic subproblems, orbital branching allows for considerably more flexibility in the choice of branching entity than isomorphism pruning. Furthermore, orbital branching can be easily incorporated within a standard integer programming solver and even exploit problem symmetry that may only be locally present at a nodal subproblem.

A preliminary version of this work has been published in the conference proceedings of the the Twelfth Conference on Integer Programming and Combinatorial Optimization (IPCO) [22]. The remainder of the paper is divided into six sections. In Sect. 2 we give some mathematical preliminaries. Orbital branching is introduced and formalized in Sect. 3. Enhancements to orbital branching are discussed in Sect. 4, and a more complete comparison to isomorphism pruning is also presented there. Implementation details are provided in Sect. 5, and computational results are presented in Sect. 6. Conclusions about the impact of orbital branching and future research directions are given in Sect. 7.

2 Preliminaries

Orbital branching is based on elementary concepts from algebra that we recall in this section to make the presentation self-contained. Some definitions are made in terms of an arbitrary permutation group Γ , but for concreteness, the reader may consider the group Γ to be the symmetry group of the matrix $\mathcal{G}(A)$.

For a set $S \subseteq I^n$, the *orbit* of S under the action of Γ is the set of all subsets of I^n to which S can be sent by permutations in Γ , i.e.,

$$\text{orb}(S, \Gamma) \stackrel{\text{def}}{=} \{S' \subseteq I^n \mid \exists \pi \in \Gamma \text{ such that } S' = \pi(S)\}.$$

In orbital branching we are concerned with the orbits of sets of cardinality one, corresponding to decision variables x_j in PIP or CIP. By definition, if $j \in \text{orb}(\{k\}, \Gamma)$, then $k \in \text{orb}(\{j\}, \Gamma)$, i.e., the variable x_j and x_k share the same orbit. Therefore, the union of the orbits

$$\mathcal{O}(\Gamma) \stackrel{\text{def}}{=} \bigcup_{j=1}^n \text{orb}(\{j\}, \Gamma)$$

forms a partition of $I^n = \{1, 2, \dots, n\}$, which we refer to as the orbital partition of Γ , or simply the *orbits* of Γ . The orbits encode which variables are “equivalent” with respect to the symmetry Γ .

The stabilizer of a set $S \subseteq I^n$ in Γ is the set of permutations in Γ that send S to itself.

$$\text{stab}(S, \Gamma) = \{\pi \in \Gamma \mid \pi(S) = S\}.$$

The stabilizer of S is a subgroup of Γ .

Throughout this paper, we display permutations in *cyclic notation*. The expression (p_1, p_2, \dots, p_k) denotes a cycle which sends entry p_i to p_{i+1} for $i = 1, \dots, k - 1$ and sends p_k to p_1 . Permutations may be written as a product of cycles. We will omit all 1-element cycles from our display. For more detailed information on groups, the reader is invited to refer to [3, 7, 24].

We characterize a node $a = (F_1^a, F_0^a)$ of the branch-and-bound enumeration tree by the indices of variables fixed to one F_1^a and fixed to zero F_0^a at node a . The set of free variables at node a is denoted by $N^a = I^n \setminus F_0^a \setminus F_1^a$. At node a , the set of feasible solutions to (CIP) or (PIP) is denoted by $\mathcal{F}(a)$, and the value of an optimal solution for the subtree rooted at node a is denoted as $z^*(a)$.

3 Orbital branching

In this section we introduce orbital branching, an intuitive way to exploit the orbits of the symmetry group $\mathcal{G}(A)$ when making branching decisions. The classical 0–1 branching variable dichotomy does not take advantage of the problem information encoded in the symmetry group. To take advantage of this information in orbital branching, instead of branching on individual variables, orbits of variables are used to create the branching dichotomy. Informally, suppose that at the current subproblem there is an orbit of cardinality k in the orbital partitioning. In orbital branching, the current subproblem is split into $k + 1$ subproblems: the first k subproblems are obtained by fixing to one in turn each variable in the orbit while the $(k + 1)$ st subproblem is obtained by fixing all variables in the orbit to zero. For any pair of variables x_i and x_j in the same orbit, the subproblem created when x_i is fixed to one is essentially equivalent to the subproblem created when x_j is fixed to one. Therefore, we can keep in the subproblem list only *one* representative subproblem, pruning the $(k - 1)$ equivalent subproblems. This is formalized below.

3.1 Orbital branching: description

Let $A(F_1^a, F_0^a)$ be the matrix obtained by removing from the constraint matrix A all columns in $F_0^a \cup F_1^a$ and either all rows intersecting columns in F_1^a (CIP case) or all columns nonorthogonal to columns in F_1^a (PIP case). When we remove columns from the matrix, we do not change the index on any of the remaining columns. Keeping the original indices introduces some notational difficulties that can be overcome by introducing a suitable mapping. Precisely, permutations π acting on the set of free variables N^a can be extended to act on I^n by creating a permutation in I^n which acts identically to π in the set N^a while not moving any fixed elements corresponding to fixed variables. This mapping, $\phi : [0, 1]^n \rightarrow [0, 1]^{|N^a|}$ with $\phi(x)_i = x_i$ for all $i \in N^a$, also maps feasible solutions with respect to A to feasible solutions with respect to $A(F_1^a, F_0^a)$. Therefore, for simplicity of exposition, we will think of all permutations as acting on the set I^n and we will not differentiate solutions which are feasible at node a with solutions feasible with respect to $A(F_1^a, F_0^a)$.

Let $O = \{i_1, i_2, \dots, i_{|O|}\} \subseteq N^a$ be an orbit of the symmetry group $\mathcal{G}(A(F_1^a, F_0^a))$. Given a subproblem a , the disjunction

$$x_{i_1} = 1 \vee x_{i_2} = 1 \vee \dots \vee x_{i_{|O|}} = 1 \vee \sum_{i \in O} x_i = 0 \tag{1}$$

splits the search space. In what follows, we show that for any two variables $x_j, x_k \in O$, the two children $a(j)$ and $a(k)$ of a , obtained by fixing, respectively x_j and x_k to 1 have the same optimal solution value. As a consequence, disjunction (1) can be replaced by the binary disjunction

$$x_h = 1 \vee \sum_{i \in O} x_i = 0, \tag{2}$$

where h is a variable in O . Formally, we have Theorem 1.

Theorem 1 *Let O be an orbit in the orbital partitioning $\mathcal{O}(\mathcal{G}(A(F_1^a, F_0^a)))$, and let j, k be two variable indices in O . If $a(j) = (F_1^a \cup \{j\}, F_0^a)$ and $a(k) = (F_1^a \cup \{k\}, F_0^a)$ are the child nodes created when branching on variables x_j and x_k , then $z^*(a(j)) = z^*(a(k))$.*

Proof Let x^* be an optimal solution of $a(j)$ with value $z^*(a(j))$. Obviously x^* is also feasible for a . Since j and k are in the same orbit O , there exists a permutation $\pi \in \mathcal{G}(A(F_1^a, F_0^a))$ such that $\pi(j) = k$. By definition, $\pi(x^*)$ is a feasible solution of a with value $z^*(a(j))$ such that $x_k = 1$. Therefore, $\pi(x^*)$ is feasible for $a(k)$, and $z^*(a(k)) = z^*(a(j))$. \square

In fact, the proof of Theorem 1 establishes a slightly stronger result. Namely, if there exists a solution of value T in $a(j)$, then there exists a (symmetric) solution of value T in $a(k)$. The basic *orbital branching* method is formalized in Algorithm 1.

Algorithm 1 Orbital Branching

Input:	Subproblem $a = (F_1^a, F_0^a)$, non-integral solution \hat{x} .
Output:	Two child subproblems b and c .
<hr/>	
Step 1.	Compute orbital partition $\mathcal{O}(\mathcal{G}(A(F_1^a, F_0^a))) = \{O_1, O_2, \dots, O_p\}$.
Step 2.	Select orbit $O_{j^*}, j^* \in \{1, 2, \dots, p\}$.
Step 3.	Choose arbitrary $k \in O_{j^*}$. Return subproblems $b = (F_1^a \cup \{k\}, F_0^a)$ and $c = (F_1^a, F_0^a \cup O_{j^*})$.

The consequence of Theorem 1 is that the search space is limited, but orbital branching has also the relevant effect of reducing the likelihood of encountering symmetric solutions. Namely, no solutions in the left and right child nodes of the current node will be symmetric with respect to the local symmetry. This is formalized in Theorem 2.

Theorem 2 *Let b and c be any two subproblems in the enumeration tree. Let a be the first common ancestor of b and c . If $a \neq b, a \neq c$, then there $\exists x \in \mathcal{F}(b)$ such that $\exists \pi \in \mathcal{G}(A(F_1^a, F_0^a))$ with $\pi(x) \in \mathcal{F}(c)$.*

Proof Suppose not, i.e., that there $\exists x \in \mathcal{F}(b)$ and a permutation $\pi \in \mathcal{G}(A(F_1^a, F_0^a))$ such that $\pi(x) \in \mathcal{F}(c)$. Let $O_i \in \mathcal{O}(\mathcal{G}(A(F_1^a, F_0^a)))$ be the orbit chosen to branch on at subproblem a . W.l.o.g. we can assume $x_k = 1$ for some $k \in O_i$, that is, b is in the left branch of a . We have that $x_k = [\pi(x)]_{\pi(k)} = 1$, but $\pi(k) \in O_i$. Therefore, by the orbital branching dichotomy, $\pi(k) \in F_0^c$, so $\pi(x) \notin \mathcal{F}(c)$. \square

Note that by using the matrix $A(F_1^a, F_0^a)$, orbital branching attempts to use symmetry found at all nodes in the enumeration tree, not just the symmetry found at the root node. This makes it possible to prune nodes whose corresponding solutions are not symmetric in the original IP.

3.2 Orbital branching: an illustrative example


Example 1 In order to demonstrate the effects of orbital branching, consider the graph $G = (V, E)$ of Fig. 1 and the associated PIP:

$$\begin{aligned} \max \quad & \sum_{i \in V} x_i \\ \text{s.t.} \quad & x_i + x_j \leq 1 \quad \forall \{i, j\} \in E, \\ & x_i \in \{0, 1\} \quad \forall i \in V \end{aligned}$$

which corresponds to computing the stability number of G .

Applying Step 1 of Algorithm 1 at the root subproblem $F_1^a = F_0^a = \emptyset$ results in a group $\mathcal{G}(A)$ containing 4,096 permutations and an orbital partition $\mathcal{O}(\mathcal{G}(A))$ containing two orbits, namely, $O_1 = \{1, \dots, 8\}$ and $O_2 = \{9, \dots, 24\}$. Thanks to the

Fig. 1 Example






Fig. 2 Child subproblems

structure of the matrix A , in which each constraint corresponds to an edge of G , the orbits of $\mathcal{G}(A)$ can be intuitively visualized on the graph.

Step 2 of Algorithm 1 selects an orbit on which to base the branching dichotomy. Suppose the largest orbit O_2 is chosen, and the branching index $k = 9 \in O_2$ is used. Then, two subproblems b and c are generated as follows: $F_1^b = \{9\}$ and $F_0^b = \emptyset$; $F_1^c = \emptyset$ and $F_0^c = \{9, \dots, 24\}$. The structure of subproblems b and c , where fixed variables have been removed, is drawn in Fig. 2.

The advantage of orbital branching over classical branching on a variable is highlighted by completely executing two branch-and-bound algorithms on the PIP of Example 1. We assume that a feasible solution of (optimal) value 8 is found at the root node. In the first algorithm the branching decision is carried out by orbital branching where Step 2 selects the largest orbit. In the second algorithm, ordinary branching is performed on the variable corresponding to the vertex of G with maximum degree in the remaining graph, typically effective for stable set problems [25]. In Figs. 3 and 4, the complete enumeration trees obtained respectively by orbital branching and branching on variable are drawn. At each node a , we report the variables fixed (F_1^a, F_0^a) and the value of the LP relaxation z_{LP} . Orbital branching results in fewer evaluated subproblems: 21 versus 49 for the variable-branching dichotomy.

An insightful explanation of orbital branching’s improved performance is obtained by examining the structure of subproblems. For instance, Fig. 5 shows the graphs remaining at subproblems 9 and 19 of the variable-branching enumeration tree. The graphs are isomorphic, but both subproblems are evaluated when branching on variables. On the contrary, orbital branching breaks such a symmetry at the root subproblem. The complete catalog of graphs and orbital partitions for each subproblem in the orbital branching branch-and-bound tree is reported in the Appendix (Fig. 11). Looking at the catalog of subproblems, one can observe that no isomorphic subproblems are evaluated when orbital branching is used on this example. This is not, however, true in general.

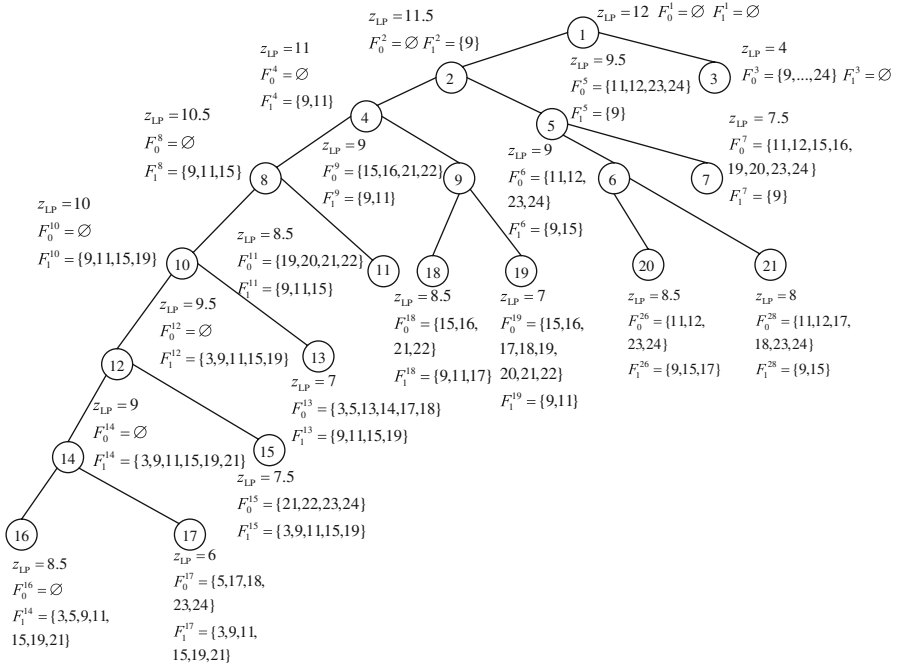


Fig. 3 Enumeration tree with orbital branching

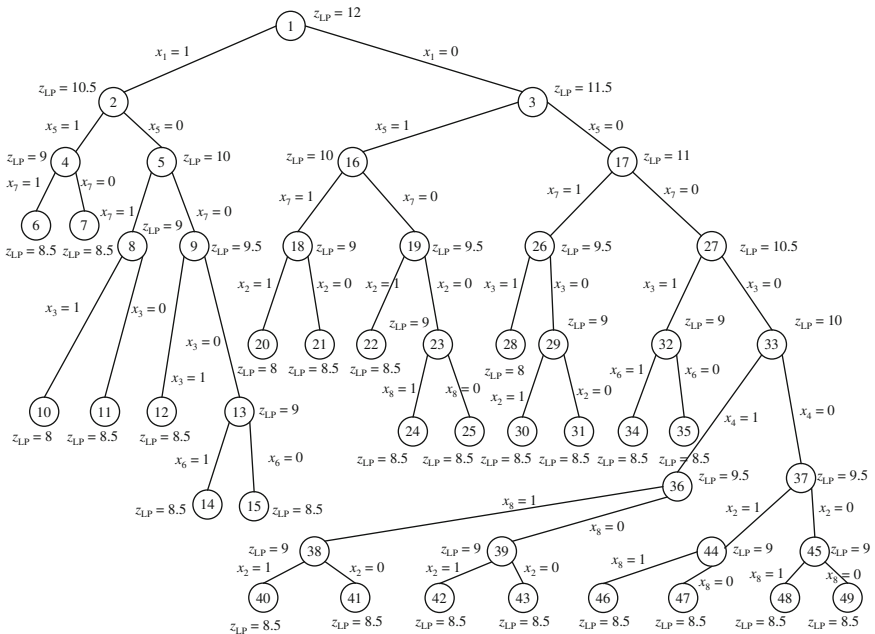


Fig. 4 Enumeration tree with branching on variable




Fig. 5 Isomorphic subproblems from branching on variable

4 Enhancements to orbital branching

In this section, we demonstrate how additional variables may be fixed during branch and bound by considering the implications of symmetry. We also discuss how to perform orbital branching by considering a subgroup of the original symmetry group. We compare orbital branching to a related technique for combating symmetry in integer programs, isomorphism pruning. The section concludes with a brief discussion on how to most effectively employ orbital branching on integer programs whose optimal solution has a large support.

4.1 Orbital fixing

In orbital branching, all variables fixed to zero and one are removed from the constraint matrix at every node in the enumeration tree. As Theorem 2 demonstrates, using orbital branching in this way ensures that any two nodes are not equivalent with respect to the symmetry found at their first common ancestor. It is possible however, for two child subproblems to be equivalent with respect to a symmetry group found elsewhere in the tree. In order to combat this type of symmetry we perform *orbital fixing*, which works as follows.

Consider the symmetry group $\mathcal{G}(A(F_1^a, \emptyset))$ at node a . If there exists an orbit O in the orbital partition $\mathcal{O}(\mathcal{G}(A(F_1^a, \emptyset)))$ that contains variables such that $O \cap F_0^a \neq \emptyset$ and $O \cap N^a \neq \emptyset$, then all variables in O can be fixed to zero. In the following theorem, we show that such variable setting (orbital fixing) excludes feasible solutions only if there exists a feasible solution of the same objective value to the left of the current node in the branch and bound tree. (We assume that the enumeration tree is oriented so that the branch with an additional variable fixed at one is the left branch).

To aid in our development, we introduce the concept of a *focus node*. For $x \in \mathcal{F}(a)$, we call node $b(a, x)$ a focus node of a with respect to x if $\exists y \in \mathcal{F}(b)$ such that $e^T x = e^T y$ and b is found to the left of a in the tree.

Theorem 3 Let $\{O_1, O_2, \dots, O_q\}$ be an orbital partitioning of $\mathcal{G}(A(F_1^a, \emptyset))$ at node a , and let the set

$$S \stackrel{\text{def}}{=} \{j \in N^a \mid \exists k \in F_0^a \text{ and } j, k \in O_\ell \text{ for some } \ell \in \{1, 2, \dots, q\}\}$$

be the set of free variables that share an orbit with a variable fixed to zero at a . If $x \in \mathcal{F}(a)$ with $x_i = 1$ for some $i \in S$, then either there exists a focus node for a with respect to x or x is not an optimal solution.

Proof Let $S \neq \emptyset$. Then, there exist $j \in F_0^a$ and $i \in S$ such that $i \in \text{orb}(j, \mathcal{G}(A(F_1^a, \emptyset)))$, i.e., there exists a $\pi \in \mathcal{G}(A(F_1^a, \emptyset))$ with $\pi(i) = j$. W.l.o.g., suppose that j is any of the first such variables fixed to zero on the path from the root node to a and let c be the first subproblem in which j is fixed. Let $\rho(c)$ be the parent node of c . By our choice of j as the first fixed variable, $\{\pi(i) \mid x_i = 1 \text{ and } \pi \in \mathcal{G}(A(F_1^a, \emptyset))\} \cap F_0^{\rho(c)} = \emptyset$. Therefore, $\pi(x)$ is not feasible in a since it does not satisfy the bounds, but is feasible in $\rho(c)$ and has the same objective value of x .

The variable x_j could have been fixed either (i) as a result of a branching decision, or (ii) it was deduced that no optimal solution exists with $x_j = 1$ at node $\rho(c)$ (and the fixing applied to the child nodes), or (iii) by orbital fixing (at ρ).

- (i) If j was fixed by orbital branching then the left child of $\rho(c)$ has $x_h = 1$ for some $h \in \text{orb}(j, \mathcal{G}(A(F_1^{\rho(c)}, F_0^{\rho(c)})))$. Let $\pi' \in \mathcal{G}(A(F_1^{\rho(c)}, F_0^{\rho(c)}))$ have $\pi'(j) = h$. Then $\pi'(\pi(x))$ is feasible in the left node with the same objective value of x . The left child node of $\rho(c)$ is then the focus node of a with respect to x .
- (ii) If it was deduced that no optimal solution feasible at $\rho(c)$ exists with $x_j = 1$, then, since $\pi(x)$ is feasible in $\rho(c)$ with $x_j = 1$, and π preserves objective value, x cannot be an optimal solution.
- (iii) Lastly, j could have been fixed by orbital fixing. This implies that the set S is nonempty in $\rho(c)$ and the argument can be repeated until the first ancestor d of a is reached such that F_0^d does not contain variables fixed by orbital fixing. Therefore, a sequence of permutations π^1, \dots, π^r has been found such that $\pi^r \pi^{r-1} \dots \pi^1 \pi(x)$ is feasible in d and has the same value of x .

Then, either argument (i) or (ii) can be applied, that is, either there is a focus node f of d with respect to $\pi^r \pi^{r-1} \dots \pi^1 \pi(x)$ (which would also be a focus node for a with respect to x), or j was fixed by an optimality condition (which implies $\pi^r \pi^{r-1} \dots \pi^1 \pi(x)$ and thus x are not optimal). □

There may be elements in S which do not share an orbit with j . One can show that these elements can also be fixed by adding the fixed variables to F_0 , updating S , and repeating the argument. As long as S is non-empty, each iteration will fix at least one variable.

An immediate consequence of Theorem 3 is that for all $i \in F_0^a$ and for all $j \in \text{orb}(i, \mathcal{G}(A(F_1^a, \emptyset)))$ one can set $x_j = 0$. We update orbital branching to include orbital fixing in Algorithm 2.

In orbital fixing, the set S of additional variables set to zero depends on F_0^a . Variables may appear in F_0^a due to a branching decision or due to traditional methods for

Algorithm 2 Orbital Branching with Orbital Fixing

Input:	Subproblem $a = (F_1^a, F_0^a)$ (with free variables $N^a = I^n \setminus F_1^a \setminus F_0^a$), fractional solution \hat{x} .
Output:	Two child nodes b and c .
Step 1.	Compute orbital partition $\mathcal{O}(\mathcal{G}(A(F_1^a, \emptyset))) = \{\hat{O}_1, \hat{O}_2, \dots, \hat{O}_q\}$. Let $S \stackrel{\text{def}}{=} \{j \in N^a \mid \exists k \in F_0^a \text{ and } (j \cap k) \in \hat{O}_\ell \text{ for some } \ell \in \{1, 2, \dots, q\}\}$.
Step 2.	Compute orbital partition $\mathcal{O}(\mathcal{G}(A(F_1^a, F_0^a))) = \{O_1, O_2, \dots, O_p\}$.
Step 3.	Select orbit $O_{j^*}, j^* \in \{1, 2, \dots, p\}$.
Step 4.	Choose arbitrary $k \in O_{j^*}$. Return child subproblems $b = (F_1^a \cup \{k\}, F_0^a \cup S)$ and $c = (F_1^a, F_0^a \cup O_{j^*} \cup S)$.

variable fixing in integer programming, e.g., reduced cost fixing or implication-based fixing. Orbital fixing, then, gives a way to *enhance* traditional variable-fixing methods by including the symmetry present at a node of the branch and bound tree.

Example (continued) When orbital branching with orbital fixing is applied to the PIP of Example 1, it generates the enumeration tree drawn in Fig. 6.

Orbital fixing is performed at subproblem 6, a node that has $F_0^6 = \{11, 12, 23, 24\}$ and $F_1^6 = \{9, 15\}$. The group $\mathcal{G}(A(F_1^6, \emptyset))$ yields the orbits: $\{2, 3\}\{5, 8\}\{6, 7\}\{11, 12, 13, 14\}\{17, 18, 23, 24\}\{19, 20, 21, 22\}$. The orbit $\{11, 12, 13, 14\}$ contains variables that have already been set to zero: $\{11, 12, 13, 14\} \cap F_0^a = \{13, 14\}$. Therefore, the variables x_{13} and x_{14} are fixed to 0 by orbital fixing. In the same way, looking at the orbit $\{17, 18, 23, 24\}$, orbital fixing sets variables x_{17} and x_{18} to 0. All the variables fixed to 0 by orbital fixing are underlined in Fig. 6.

The effect of orbital fixing is clear at subproblem 6, where the optimal value of the LP relaxation reduces from 9 to 7, as compared to the algorithm without orbital fixing, avoiding further branching (see the tree of Fig. 3).

The example also helps illustrate the existence of a focus node if orbital fixing is performed (Theorem 3). Define a as the subproblem found at node 6. The set of variables fixed by orbital fixing is $S = \{13, 14, 17, 18\}$. Consider the solution $x \in \mathcal{F}(a)$: $x_2 = x_5 = x_8 = x_9 = x_{13} = x_{15} = x_{19} = x_{21} = 1$, and all other variables set to 0. Following the proof of Theorem 3, we have $i = 13$ and $j \in \text{orb}(\{i\}, \mathcal{G}(A(F_1^a, \emptyset)))$, i.e., $j = 12$. A permutation $\pi \in \mathcal{G}(A(F_1^a, \emptyset))$ such that $\pi(i) = j$ is: $[(2, 3), (12, 13), (11, 14)]$. We have $\bar{x} = \pi(x)$, that is, $\bar{x}_2 = \bar{x}_5 = \bar{x}_8 = \bar{x}_{12} = \bar{x}_{19} = \bar{x}_{21} = 1$, and all other variables set to 0. Notice that $\bar{x} \notin \mathcal{F}(a)$, since $x_{12} = 1$. By definition, subproblem 5 is the subproblem c in the proof of Theorem 3, and subproblem 2 is the subproblem $\rho(c)$. Then, we have $h = 11$ and π' can be defined as: $(11, 12)$. Finally, $\tilde{x} = \pi'(\pi(x))$ is: $\tilde{x}_3 = \tilde{x}_7 = \tilde{x}_9 = \tilde{x}_{11} = \tilde{x}_{15} = \tilde{x}_{17} = \tilde{x}_{19} = \tilde{x}_{23} = 1$, and all other variables set to 0. This is feasible for subproblem 4. Thus, 4 is a focus node for a .

4.2 Using a subgroup of the original symmetry group

We delay discussion of the computation of the symmetry groups $\mathcal{G}(A(F_1^a, F_0^a))$ until Sect. 5.1, but we simply note at this point that all known algorithms which compute the

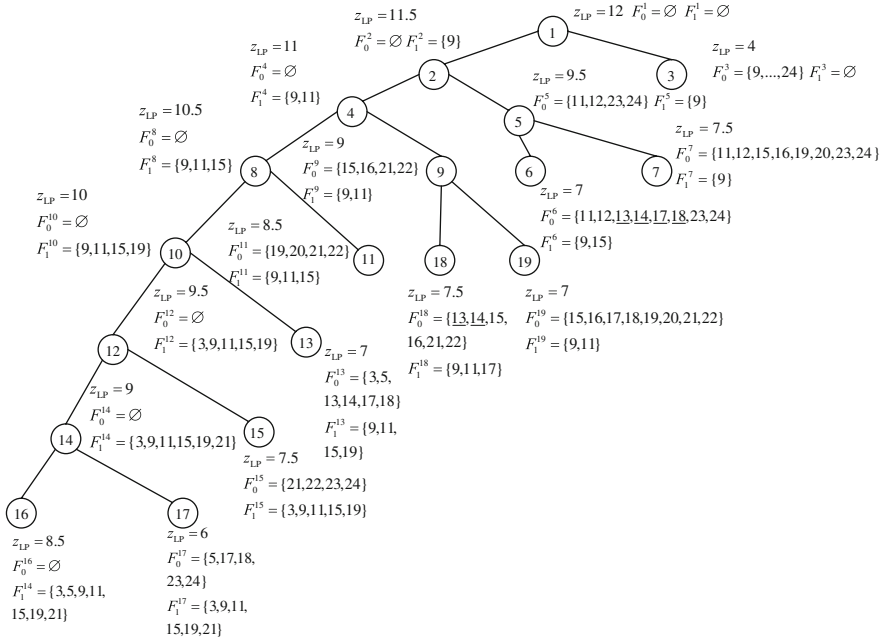


Fig. 6 Enumeration tree with orbital branching and orbital fixing

symmetry group of a given graph have exponential running time. Thus, computing the symmetry group $\mathcal{G}(A(F_1^a, F_0^a))$ at each node a may be computationally prohibitive. We will show via computational results in Sect. 6 that this is often not the case. In the case that recomputing the full symmetry group $\mathcal{G}(A(F_1^a, F_0^a))$ is too costly, there is an alternative. Instead, orbital branching can use the symmetry group $\text{stab}(F_1^a, \mathcal{G}(A))$ to create orbits at every node in the tree. In this method, the original, global, symmetry group $\mathcal{G}(A)$ is only computed once, at the root node, and the stabilizers are computed given the original symmetry group. This is typically more computationally efficient than re-computing the symmetry groups from scratch. In the following, in order to distinguish between the two symmetry groups that could be used in orbital branching at node a , we will refer to $\text{stab}(F_1^a, \mathcal{G}(A))$ as the *global symmetry group* (because we use only the symmetry group found at the root node), and $\mathcal{G}(A(F_1^a, F_0^a))$ as the *local symmetry group*.

The decreased computational overhead in orbit calculations when using the global symmetry group comes at a price. As Theorems 4 and 5 demonstrate, the orbits from the global group $\text{stab}(F_1^a, \mathcal{G}(A))$ are a subdivision of the orbits used for orbital fixing and orbital branching, so the branching dichotomy and fixing mechanisms are weaker.

Theorem 4 *If $O \in \mathcal{O}(\text{stab}(F_1^a, \mathcal{G}(A)))$ and $O \cap F_1^a = \emptyset$, then $\exists O' \in \mathcal{O}(\mathcal{G}(A(F_1^a, \emptyset)))$ with $O \subseteq O'$.*

Proof We prove Theorem 4 by proving the following equivalent statement: if $\exists \pi \in \text{stab}(F_1^a, \mathcal{G}(A))$ with $\pi(i) = j$ and $i, j \notin F_1^a$, then $\exists \pi' \in \mathcal{G}(A(F_1^a, \emptyset))$ with $\pi(i) = j$.

Let $\pi \in \text{stab}(F_1^a, \mathcal{G}(A))$ be such that $\pi(i) = j$ and $i, j \notin F_1^a$. Since $\pi(F_1^a) = F_1^a$ we can restrict π by ignoring its action on F_1^a . Let C^a be the collection of inequalities which have been removed (become redundant) either at a or any parent of a . Since $\pi \in \text{stab}(F_1^a, \mathcal{G}(A))$, there exists a $\sigma \in \Pi^m$ such that $A(\sigma, \pi) = A$. Each constraint, $c^T x \geq (\leq) 1$, in C^a contains at least one variable in F_1^a . This constraint gets mapped to $\sigma(c)^T \pi(x) \geq (\leq) 1$, a constraint still containing at least one variable in F_1^a , hence a constraint in C^a . We can then restrict σ by ignoring its action on the constraint set C^a . Call the pair of restricted permutations π' and σ' . These permutations act on the same set of variables and constraints as $\mathcal{G}(A(F_1^a, \emptyset))$. We also have that $A(F_1^a, \emptyset)(\pi', \sigma') = A(F_1^a, \emptyset)$, so $\pi' \in \mathcal{G}(A(F_1^a, \emptyset))$ with $\pi'(i) = j$. \square

Orbital fixing does not change the result of Theorem 5. Specifically, if S^a is the set of indices of variables fixed to zero by orbital fixing at node a , then the orbits from the group $\mathcal{G}(A(F_1^a, \emptyset))$ are a subdivision of orbits from the group $\mathcal{G}(A(F_1^a, F_0^a \cup S^a))$.

Theorem 5 *Let S^a be the set of variables fixed to zero by orbital fixing at node a . If $O \in \mathcal{O}(\mathcal{G}(A(F_1^a, \emptyset)))$, $\exists O' \in \mathcal{O}(\mathcal{G}(A(F_1^a, F_0^a \cup S^a)))$ with $O \subseteq O'$.*

Proof We prove Theorem 5 by proving the equivalent statement that if $\exists \pi \in \mathcal{G}(A(F_1^a, \emptyset))$ with $\pi(i) = j$ and $i, j \notin S^a$, then $\exists \pi' \in \mathcal{G}(A(F_1^a, F_0^a \cup S^a))$ with $\pi'(i) = j$. Let $\pi \in \mathcal{G}(A(F_1^a, \emptyset))$ with $\pi(i) = j$. We can restrict π by ignoring its actions on the set $F_0^a \cup S^a$. Call the restricted permutation π' . Let C^a be the collection of inequalities which have been removed (become redundant) either at a or any parent of a . We know that there exists a $\sigma \in \Pi^{m-|C^a|}$ such that $A(F_1^a, \emptyset)(\pi, \sigma) = A(F_1^a, \emptyset)$. Since $A(F_1^a, F_0^a)$ contains the same rows as $A(F_1^a, \emptyset)$, we have that $A(F_1^a, F_0^a)(\pi', \sigma) = A(F_1^a, F_0^a)$. \square

4.3 Comparison to isomorphism pruning

The fundamental idea behind isomorphism pruning is that for each node $a = (F_1^a, F_0^a)$, the orbits $\text{orb}(F_1^a, \mathcal{G}(A))$ of the “equivalent” sets of variables to F_1^a are computed. If there is a node $b = (F_1^b, F_0^b)$ elsewhere in the enumeration tree such that $F_1^b \in \text{orb}(F_1^a, \mathcal{G}(A))$, then the node a need not be evaluated—the node a is pruned by isomorphism. A very distinct and powerful advantage of this method is that *no* nodes whose sets of variables fixed to 1 are isomorphic will be evaluated. One disadvantage of this method is that computing $\text{orb}(F_1^a, \mathcal{G}(A))$ can require significant computational effort. Further the set $\text{orb}(F_1^a, \mathcal{G}(A))$ may contain many equivalent subsets to F_1^a , and the entire enumeration tree must be compared against this list to ensure that a is not isomorphic to any other node b . In a series of papers, Margot offers a way around this second disadvantage [15, 16]. The key idea introduced is to declare one *unique representative* among the members of $\text{orb}(F_1^a, \mathcal{G}(A))$, and if F_1^a is not the unique representative, then the node a may safely be pruned. The oracle that checks if F_1^a a unique representative among $\text{orb}(F_1^a, \mathcal{G}(A))$ runs in polynomial time. The disadvantage of the method is ensuring that the unique representative occurs *somewhere* in the branch and bound tree requires a relatively inflexible branching rule. Namely, *all* child nodes at a fixed depth must be created by branching on the *same* variable.

Orbital branching does not suffer from this inflexibility. By not focusing on pruning *all* isomorphic nodes, but rather eliminating the symmetry through branching, orbital branching offers a great deal more flexibility in the choice of branching entity. Another advantage of orbital branching is that by using the symmetry group $\mathcal{G}(A(F_1^a, F_0^a))$, symmetry *introduced* as a result of the branching process is also exploited.

Both methods allow for the use of traditional integer programming methodologies such as cutting planes and fixing variables based on considerations such as reduced costs and implications derived from preprocessing. In isomorphism pruning, for a variable fixing to be valid, it must be that *all* non-isomorphic optimal solutions are in agreement with the fixing. Orbital branching does not suffer from this limitation. A powerful idea in both methods is to combine the variable fixing with symmetry considerations in order to fix many additional variables. This idea is called *orbit setting* in [16] and *orbital fixing* in this work (see Sect. 4.1).

4.4 Reversing orbital branching

One of the advantages of orbital branching is that the “right” branch, in which all variables in the branching orbit O are fixed to zero, typically changes the optimal value of the LP relaxation significantly, and the left branch, in which one variable in O is fixed to one also has a significant impact on the problem. In some classes of PIP or CIP, fixing a variable to zero can have more impact than fixing a variable to one. This is typically true in instances in which the number of ones in an optimal solution is larger than $1/2$ the number of variables. In such cases, orbital branching would be much more efficient if all variables were complemented, or equivalently if the orbital branching dichotomy (2) was replaced by its complement. Margot [16] also makes a similar observation for the isomorphism pruning algorithm and solves the complemented versions of such instances. In orbital branching, we opt for the former way of exploiting this fact, and the “left” branch fixes one variable to zero, and orbital fixing fixes variables to one instead of zero.

5 Implementation

The orbital branching method has been implemented using the user application functions of MINTO v3.1 [21]. The branching dichotomy of Algorithm 1 or 2 is implemented in the `appl_divide()` method, and reduced cost fixing is implemented in `appl_bounds()`. The entire implementation, including code for all the branching rules subsequently introduced in Sect. 5.2 consists of slightly over 1,000 lines of code. All advanced IP features of MINTO were used, including *clique inequalities*, which can be useful for instances of (PIP). In this section, we discuss the features of the implementation that are specific to orbital branching—the computation of the symmetry groups and orbital branching rules.

5.1 Computing $\mathcal{G}(\cdot)$

Computation of the symmetry groups required for orbital branching and orbital fixing is done by computing the automorphism group of a related graph. Recall that the automorphism group $\text{Aut}(G(V, E))$ of a graph $G = (V, E)$, is the set of permutations of V that leave the incidence matrix of G unchanged, i.e.,

$$\text{Aut}(G(V, E)) = \left\{ \pi \in \Pi^{|V|} \mid \{i, j\} \in E \Leftrightarrow \{\pi(i), \pi(j)\} \in E \right\}.$$

The matrix A whose symmetry group is to be computed is transformed into a bipartite graph $G(A) = (N, M, E)$ where vertex set $N = \{1, 2, \dots, n\}$ represents the variables, vertex set $M = \{n+1, n+2, \dots, n+m\}$ represents the constraints, and edge $(i, j) \in E$ if and only if $a_{ij} = 1$. Under this construction, feasible solutions to (PIP) are subsets of the vertices $S \subseteq N$ such that each vertex $i \in M$ is adjacent to *at most* one vertex $j \in S$. In this case, we say that S *packs* M . Feasible solutions to (CIP) correspond to subsets of vertices $S \subseteq N$ such that each vertex $i \in M$ is adjacent to *at least* one vertex $j \in S$, or S *covers* M . Since applying members of the automorphism group preserves the incidence structure of a graph, if S packs (covers) M , and $\pi \in \text{stab}(M, \text{Aut}(G(A)))$, then there exists a $\sigma \in \Pi^M$ such that $\sigma(M) = M$ and $\pi(S)$ packs (covers) $\sigma(M)$. This implies that if $\pi \in \text{stab}(M, \text{Aut}(G(A)))$, then the restriction of π to N must be an element of $\mathcal{G}(A)$, i.e., using the graph $G(A)$, one can find elements of symmetry group $\mathcal{G}(A)$. In particular, we compute the orbital partition of the stabilizer of the constraint vertices M in the automorphism group of $G(A)$, i.e.,

$$\mathcal{O}(\text{stab}(M, \text{Aut}(G(A)))) = \{O_1, O_2, \dots, O_p\}.$$

The orbits O_1, O_2, \dots, O_p in the orbital partition are such that if $i \in M$ and $j \in N$, then i and j are not in the same orbit. We can then refer to these orbits as *variable* orbits and *constraint* orbits. In orbital branching, we are concerned only with the variable orbits.

There are several software packages that can compute the automorphism groups required to perform orbital branching. The program *nauty* [18], by McKay, has been shown to be quite effective [5], and we use *nauty* in our orbital branching implementation.

The complexity of computing the automorphism group of a graph is not known to be polynomial time. However, *nauty* was able to compute the symmetry groups of our problems very quickly, generally faster than solving an LP at a given node. One explanation for this phenomenon is that the running time of *nauty*'s backtracking algorithm is correlated to the size of the symmetry group being computed. For example, computing the automorphism group of the clique on 2,000 nodes takes 85 s, while graphs of comparable size with little or no symmetry require fractions of a second. The orbital branching procedure quickly reduces the symmetry group of the child subproblems, so explicitly recomputing the group by calling *nauty* is computationally very feasible. In the table of results presented in the Appendix, we state explicitly the time required in computing automorphism groups by *nauty*.

5.2 Branching rules

The orbital branching rule introduced in Sect. 3 leaves significant freedom in choosing the orbit on which to base the branching (Step 2 of Algorithm 1). In this section, we discuss mechanisms for deciding on which orbit to branch. As input to the branching decision, we are given a fractional solution \hat{x} and orbits O_1, O_2, \dots, O_p (consisting of all currently free variables) of the orbital partition $\mathcal{O}(\mathcal{G}(A(F_1^a, F_0^a)))$ for the subproblem at node a . Output of the branching decision is an index j^* of an orbit on which to base the orbital branching. We tested six different branching rules.

Rule 1: Branch largest: The first rule chooses to branch on the largest orbit O_{j^*} :

$$j^* \in \arg \max_{j \in \{1, \dots, p\}} |O_j|.$$

Rule 2: Branch largest LP solution: The second rule branches on the orbit O_{j^*} whose variables have the largest total solution value in the fractional solution \hat{x} :

$$j^* \in \arg \max_{j \in \{1, \dots, p\}} \hat{x}(O_j).$$

Rule 3: Strong branching: The third rule is a strong branching rule. For each orbit j , two tentative child nodes are created and their bounds z_j^+ and z_j^- are computed by solving the resulting linear programs. The orbit j^* for which the product of the change in linear program bounds is largest is used for branching:

$$j^* \in \arg \max_{j \in \{1, \dots, p\}} \left(|e^T \hat{x} - z_j^+| \right) \left(|e^T \hat{x} - z_j^-| \right). \tag{3}$$

We also tested using a combination of the bound changes

$$j^* \in \arg \max_{j \in \{1, \dots, p\}} \left(3 \min \left(|e^T \hat{x} - z_j^+|, |e^T \hat{x} - z_j^-| \right) + \max \left(|e^T \hat{x} - z_j^+|, |e^T \hat{x} - z_j^-| \right) \right),$$

as suggested by Linderoth and Savelsbergh [12], but the computational results using Eq. 3 were slightly stronger.

Note that if one of the potential child nodes in the strong branching procedure would be pruned, either by bound or by infeasibility, then the bounds on the variables may be fixed to their values on the alternate child node. We refer to this as *strong branching fixing*, and in the computational results in the Appendix, we report the number of variables fixed in this manner. As discussed at the end of Sect. 4.1, variables fixed by strong branching fixing may result in additional variables being fixed by orbital fixing.

Rule 4: Break symmetry left: This rule is similar to *strong branching*, but instead of fixing a variable and computing the change in objective value bounds, we fix a variable and compute the change in the size of the symmetry group. Specifically, for each orbit j , we compute the size of the symmetry group in the resulting left branch

if orbit j (including variable index i_j) was chosen for branching, and we branch on the orbit that reduces the symmetry by as much as possible:

$$j^* \in \arg \min_{j \in \{1, \dots, p\}} (|\mathcal{G}(A(F_1^a \cup \{i_j\}, F_0^a))|).$$

Rule 5: Keep symmetry left: This branching rule is the same as **Rule 4**, except that we branch on the orbit for which the size of the child's symmetry group would remain the largest:

$$j^* \in \arg \max_{j \in \{1, \dots, p\}} (|\mathcal{G}(A(F_1^a \cup \{i_j\}, F_0^a))|).$$

Rule 6: Branch max product left: This rule attempts to combine the fact that we would like to branch on a large orbit at the current level and also keep a large orbit at the second level on which to base the branching dichotomy. For each orbit O_1, O_2, \dots, O_p , the orbits $P_1^j, P_2^j, \dots, P_q^j$ of the symmetry group $\mathcal{G}(A(F_1^a \cup \{i_j\}, F_0^a))$ of the left child node are computed for some variable index $i_j \in O_j$. We then choose to branch on the orbit j^* for which the product of the orbit size and the largest orbit of the child subproblem is largest:

$$j^* \in \arg \max_{j \in \{1, \dots, p\}} \left(|O_j| \left(\max_{k \in \{1, \dots, q\}} |P_k^j| \right) \right).$$

6 Computational experiments

In this section, we give empirical evidence of the effectiveness of orbital branching, we investigate the impact of choosing the orbit on which branching is based, and we demonstrate the positive effect of orbital fixing. The computations are based on the instances whose characteristics are given in Table 1. The instances beginning with `cod` are used to compute maximum cardinality binary error correcting codes [13], the instances whose names begin with `cov` are covering designs [20], the instance `f5` is the “football pool problem” on five matches [8], and the instances `sts` are used to compute the incidence width of the well-known Steiner-triple systems [6]. The `cov` formulations have been strengthened with a number of Schönheim inequalities, as derived by Margot [17]. The `sts` instances typically have roughly two-third of the variables equal to one in an optimal solution, so for these instances, we reverse the orbital branching dichotomy, as explained in Sect. 4.4. All instances, save for `f5`, are available from Margot's web site: <http://wpweb2.tepper.cmu.edu/fmargot/lpsym.html>.

The computations comparing different orbital branching rules were run on machines with AMD Opteron processors clocked at 1.8 GHz and having 2 GB of RAM. The COIN-OR software `CLP` was used to solve the linear programs at nodes of the branch and bound tree. For each instance, the (known) optimal solution value was set a priori to aid pruning and reduce the random impact of finding a feasible solution in the search. Nodes were searched in a depth-first fashion. When the size of the maximum

Table 1 Symmetric integer programs

Name	Variables	Group size
cod83	256	10,321,920
cod93	512	185,794,560
cod105	1,024	3,715,891,200
cov1053	252	3,628,800
cov1054	2,252	3,628,800
cov1075	120	3,628,800
cov1076	120	3,628,800
cov954	126	362,880
f5	243	933,120
sts45	45	360
sts63	63	72,576
sts81	81	1,965,150,720

orbit in the orbital partitioning is less than or equal to two, nearly all of the symmetry in the problem has been eliminated by the branching procedure, and there is little use in performing orbital branching. In this case, we use MINTO's default branching strategy [12]. If orbital branching is not performed at a node, then there is little likelihood that it will be effective at the node's children. In this case, we save the computational overhead of re-computing the symmetry group, and simply allow MINTO to choose a branching variable. The CPU time was limited in all cases to four hours, and a limit of one million nodes evaluated was imposed.

Table 2 shows the results of an experiment designed to compare the performance of the six different orbital branching rules introduced in Sect. 5.2. In this experiment, reduced cost fixing, orbital fixing, and the local symmetry group $\mathcal{G}(A(F_1^a, F_0^a))$ were used, and the CPU time required (in s) and number of nodes required for orbital branching to solve each instance in the test suite for the six different is reported. A “–” in the table indicates that the instance was not solved within the four hour or one million node limit. A complete table showing the number of nodes, CPU time, CPU time computing automorphism groups, the number of variables fixed by reduced cost fixing, orbital fixing, and strong branching fixing, and the deepest tree level at which orbital branching was performed for a variety of parameter settings is shown in Table 5 in the Appendix.

In order to succinctly present many of our computational results, we use performance profiles of Dolan and Moré [4]. A performance profile is a relative measure of the effectiveness of one solution method in relation to a group of solution methods on a fixed set of problem instances. A performance profile for a solution method m is essentially a plot of the probability that the performance of m (measured in this case with CPU time or number of nodes evaluated) on a given instance in the test suite is within a factor of β of the *best* method for that instance. Methods whose corresponding profile lines are the highest are the most effective. Figure 7 shows a performance profile of the results of the first experiment, using CPU times from Table 2 as a measure of performance, and Fig. 8 shows the performance profile using the number of nodes as a measure of performance.

Table 2 CPU time and nodes for orbital branching using local symmetry group

Instance	Rule 1		Rule 2		Rule 3		Rule 4		Rule 5		Rule 6	
	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes
cod83	11	193	4	57	5	21	6	143	8	195	5	105
cod93	1,677	16,439	1,557	14,461	2,368	161	3,269	37,297	242	1,577	399	3,503
cod105	239	9	238	7	345	7	255	17	424	23	229	9
cov954	5	249	4	153	24	63	8	237	17	449	5	217
cov1053	103	3,437	617	20,725	768	777	346	15,321	105	3,139	90	2,859
cov1054	–	–	–	–	–	–	–	–	181	1,249	–	–
cov1075	69	461	50	495	216	71	–	–	210	381	128	543
cov1076	–	–	–	–	–	–	–	–	1,560	31,943	–	–
f5	64	1,829	80	2,573	668	123	42	995	34	717	64	1,835
sts45	8	4,917	8	4,683	95	1,417	8	4,571	8	4,507	8	4,917
sts63	93	33,785	91	32,627	1,132	3,157	1,630	666,623	161	9,993	137	31,261
sts81	127	11,323	164	25,739	13,465	11,291	–	–	434	83,961	–	–
Times best	2	0	5	1	0	10	1	0	5	2	3	0




Fig. 7 CPU time performance profile of branching rules

In terms of CPU time, the most effective branching method is **Rule 5**—the method that keeps the size of the symmetry group large on the left branch. (This method gives the “highest” line in Fig. 7). In fact, this branching method is the only one that is able to solve all of the instances in the test suite within the four hour time limit. This result is somewhat surprising. Anecdotally, symmetry has long been thought to be a significant hurdle for solving integer programs. One might expect that methods in which symmetry was *removed* as quickly as possible would have been the most effective.




Fig. 8 Nodes evaluated performance profile of branching rules

Our results go counter to this intuition. Instead, if effective methods for *exploiting* problem symmetry (like those in orbital branching) are present, the results indicate that one should attempt to keep a large amount of symmetry in the subproblems.


Another interesting result in this computational experiment is that the number of nodes required to solve the instance is nearly always the smallest using **Rule 3**—strong branching. (Note the height of the `strong-branch` line in Fig. 8). In fact, if strong branching can solve the problem instance within the time limit, then it requires the fewest number of nodes. However, since many linear programs are solved to determine the best branching orbit, strong branching quite often is the slowest of the branching methods in terms of CPU time. We view the performance of strong branching as a positive demonstration of the fact that exploiting the flexibility offered by orbital branching (as opposed to the more static branching rules offered by isomorphism pruning) may for some instances have a positive impact on running time. More sophisticated orbital branching implementations may choose to implement advanced “hybrid” branching rules that perform limited strong branching, like the reliability branching of Achterberg et al. [1].

A second experiment was aimed at measuring the impact of using the the global symmetry group $\text{stab}(F_1^a, \mathcal{G}(A))$ (discussed in Sect. 4.2) instead of the local symmetry group $\mathcal{G}(A(F_1^a, F_0^a))$ when making a branching decision. Table 3 shows the CPU time (in s) orbital branching, equipped with reduced cost fixing and orbital fixing, required on the instances in the test suite, for the different branching rules employing the global symmetry group.

Again, branching **Rule 5** that keeps symmetry on the left child node, was the most effective. A side-by-side comparison of Tables 2 and 3 indicates that in general using the global symmetry group is more effective than attempting to exploit symmetry that may only be locally present at a node. Figure 9 shows a performance profile comparing the CPU time required to solve the instances using branching **Rule 5** with both

Table 3 CPU time and nodes for orbital branching using global symmetry group

Instance	Rule 1		Rule 2		Rule 3		Rule 4		Rule 5		Rule 6	
	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes
cod83	10	193	3	57	5	21	1	25	1	25	5	105
cod93	1,677	16,439	1,556	14,461	2,361	161	166	1,361	167	1,361	396	3,503
cod105	237	9	237	7	359	7	234	11	242	11	237	9
cov954	5	249	4	153	23	63	13	373	6	249	5	217
cov1053	103	3,437	619	20,903	761	777	280	11,271	240	9,775	89	2,859
cov1054	–	–	–	–	–	–	–	–	179	1,249	–	–
cov1075	55	461	42	495	202	71	–	–	152	381	95	543
cov1076	–	–	–	–	–	–	–	–	1,415	31,943	–	–
f5	64	1,829	79	2,573	664	123	44	1,125	45	1,125	64	1,835
sts45	8	4,917	8	4,683	50	1,287	8	4,709	8	4,709	8	4,917
sts63	104	43,349	90	36,579	101	1,377	20	5,533	20	5,533	81	30,133
sts81	29	5,823	28	5,649	73	573	39	6,293	39	6,293	–	–
Times best	1	0	4	1	0	10	6	0	5	2	2	0

**Fig. 9** Performance profile of local versus global symmetry groups

the local and global symmetry groups. Surprisingly, the improved performance of the global symmetry group comes not only from improved efficiency of the branching calculations, but in many cases the number of *nodes* is reduced, which can be seen by comparing the relevant columns in Tables 2 and 3. These computational results run counter to Theorem 4, which states that orbits from the global symmetry group are a subdivision of orbits from the local group. Since the orbits of the local group are no smaller, one would expect that orbital branching's enumeration tree would also be smaller in this case.

A third comparison worthy of note is the impact of performing orbital fixing, as introduced in Sect. 4.1. Using branching **Rule 5**, each instance in Table 1 was run both with and without orbital fixing. Figure 10 shows a performance profile comparing the results in the two cases. The results shows that orbital fixing has a *significant* positive impact.

The final comparison we make in the paper is between different branch and bound techniques for solving the symmetric test instances. Five different algorithms were compared: the isomorphism pruning algorithm of Margot, orbital branching (using branching **Rule 5** and the global symmetry group), MINTO's default algorithm, CPLEX v11.2 without symmetry handling, and CPLEX v11.2 with symmetry handling. The symmetry handling was changed in CPLEX by setting the option `symmetry` to 0 or 5, respectively. As orbital branching is implemented in the MINTO framework, the MINTO default results demonstrate the direct impact of orbital branching on the symmetric test instances. In this experiment, we dropped the one million node limit, and imposed only a CPU time limit of 4 h. Table 4 summarizes the results of the comparison. The results for isomorphism pruning are taken directly from the paper of Margot using the most sophisticated of his branching rules "BC4" [16] running on a HP-B2000 workstation with a 500 MHz CPU. The paper [16] does not report results on $\epsilon 5$. The CPLEX results were obtained on an Intel Core 2 Duo P8600 clocked at 2.40 GHz, as this is the machine on which a CPLEX license was available. Since the results were obtained on three different computer architectures and each used a different LP solver for the child subproblems, the CPU times should be interpreted appropriately. Regardless of the impact of the CPU on performance, the results clearly show that isomorphism pruning and orbital branching are the most effective methods for these symmetric instances.




Fig. 10 Performance profile of impact of orbital fixing

Table 4 Comparison of different solvers on test instances

Instance	Isomorphism pruning		Orbital branching		MINTO default		CPLEX v11.2 w/o Sym		CPLEX v11.2 w/Sym	
	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes
cod83	19	33	1	25	14,400	346,963	14,400	1,496,664	205	54,518
cod93	651	103	167	1,361	14,400	43,305	14,400	749,100	14,400	983,205
cod105	2,000	15	242	11	14,400	4	14,400	28,313	223	215
cov954	24	126	6	249	1,180	39,213	64	33,908	5	1,476
cov1053	35	111	240	9,775	14,400	207,430	14,400	2,293,458	858	1,31,000
cov1054	130	108	179	1,249	14,400	32,679	14,400	180,512	14,400	1,75,311
cov1075	118	169	152	381	14,400	144,064	8,120	1,608,160	20	3,002
cov1076	3,634	5,121	1,415	31,943	14,400	180,147	14,400	3,04,498	14,400	3,25,598
f5	N/A	N/A	45	1,125	14,400	1,76,611	14,390	2,428,603	111	15,891
sts45	31	513	8	4,709	223	58,683	11	58,368	8	44,024
sts63	120	1,247	20	5,533	14,400	1,475,000	2,827	9,506,961	989	3,207,124
sts81	68	199	39	6,293	14,400	15,57,850	14,400	28,314,500	14,400	26,575,877

7 Conclusions

In this work, we presented a simple way to capture and exploit the symmetry of an integer program when branching. We showed through a set of experiments that the new method, orbital branching, outperforms CPLEX, a state-of-the-art solver, when a high degree of symmetry is present. Orbital branching also seems to be of comparable quality to the isomorphism pruning method of Margot [16]. Further, we feel that the simplicity and flexibility of orbital branching make it an attractive candidate for further study. Continuing research includes techniques for further reducing the number of isomorphic nodes that are evaluated, for implementing the natural extension of the orbital branching method on more general classes of integer programs, on developing branching mechanisms that combine the child bound improvement and change in symmetry in a meaningful way, and on exploiting symmetry when branching on constraints instead of variables [23].

Acknowledgments The authors would like to thank Kurt Anstreicher and François Margot for inspiring and insightful comments on this work. In particular, the name *orbital branching* was suggested by Kurt. The comments of two anonymous referees and the associate editor helped improve the paper considerably. Author Linderoth would like to acknowledge support from the US National Science Foundation (NSF) under grant DMI-0522796, by the US Department of Energy under grant DE-FG02-05ER25694, and by IBM, through the faculty partnership program. Author Ostrowski is supported by the NSF through the IGERT Grant DGE-9972780.

Appendix

See Tables 5 and 6.

Table 5 Performance of orbital branching rules (Local Symmetry) on symmetric IPs

Instance	Branching rule	Time	Nodes	Nauty calls	Nauty time	#Fixed by RCF	#Fixed by OF	#Fixed by SBF	Deepest orbital level
cod83	Branch largest	10.6	193	14	0.4	233	588	0	7
cod83	Branch largest LP	3.6	57	18	0.5	328	864	0	7
cod83	Strong branch	5.3	21	9	0.4	16	762	412	6
cod83	Break symmetry	6.2	143	41	1.3	325	548	0	15
cod83	Keep symmetry	8.2	195	73	2.3	251	942	0	18
cod83	Max prod. orbit size	4.8	105	19	0.6	69	642	0	11
cod93	Branch largest	1677.3	16,439	15	2.2	205,636	1,060	0	7
cod93	Branch largest LP	1557.1	14,461	13	1.9	201,292	348	0	7
cod93	Strong branch	2367.9	161	79	8.2	437	2,400	13,478	15
cod93	Break symmetry	3268.8	37,297	557	58.4	106,725	6,202	0	26
cod93	Keep symmetry	242.5	1,577	303	32.9	11,473	2,422	0	44
cod93	Max prod. orbit size	398.9	3,503	59	6.8	41,907	704	0	25
cod105	Branch largest	239.2	9	3	6.4	0	0	0	3
cod105	Branch largest LP	237.9	7	2	4.2	0	0	0	2
cod105	Strong branch	344.7	7	2	4.2	0	1,024	1,532	2
cod105	Break symmetry	254.5	17	7	15.0	0	1,020	0	7
cod105	Keep symmetry	423.9	23	10	21.6	216	1,228	0	8
cod105	Max prod. orbit size	229.5	9	3	6.1	1	960	0	3
cov954	Branch largest	5.3	249	20	1.2	818	304	0	12
cov954	Branch largest LP	3.8	153	11	0.7	638	0	0	6
cov954	Break symmetry	8.4	237	69	4.4	423	272	0	11
cov954	Strong branch	24.0	63	30	1.9	65	160	1,724	11
cov954	Keep symmetry	17.3	449	170	11.0	677	948	0	15
cov954	Max prod. orbit size	4.8	217	18	1.1	699	132	0	11
cov1053	Branch largest	103.5	3,437	55	1.9	0	1,094	0	17
cov1053	Branch largest LP	616.7	20,725	61	2.1	0	988	0	19
cov1053	Strong branch	768.4	777	387	14.1	0	2,834	16,462	43
cov1053	Break symmetry	345.8	15,321	800	28.7	0	2,418	0	35
cov1053	Keep symmetry	105.4	3,139	520	18.7	0	1,696	0	31
cov1053	Max prod. orbit size	90.2	2,859	71	2.5	0	1,466	0	20
cov1054	Branch largest	14400.0	105,500	10	1.7	0	0	0	7
cov1054	Branch largest LP	14400.0	104,126	6	1.1	56	88	0	5
cov1054	Strong branch	14400.0	846	430	79.3	0	220	12,846	57
cov1054	Break symmetry	14400.0	110,116	1	0.2	0	0	0	0
cov1054	Keep symmetry	181.3	1,249	103	18.5	0	454	0	15

Table 5 continued

Instance	Branching rule	Time	Nodes	Nauty calls	Nauty time	#Fixed by RCF	#Fixed by OF	#Fixed by SBF	Deepest orbital level
cov1054	Max prod. orbit size	14400.0	104,172	12	2.0	0	176	0	8
cov1075	Branch largest	68.6	461	43	44.3	1,333	900	0	13
cov1075	Branch largest LP	49.8	495	22	23.3	1,400	520	0	9
cov1075	Strong branch	215.5	71	34	37.4	126	92	1,858	10
cov1075	Break symmetry	14400.0	408,822	1	0.8	862,268	0	0	0
cov1075	Keep symmetry	209.7	381	181	189.8	413	962	0	15
cov1075	Max prod. orbit size	128.4	543	98	102.0	1,028	1,090	0	21
cov1076	Branch largest	14400.0	504,396	40	34.0	495,631	388	0	9
cov1076	Branch largest LP	14400.0	498,573	20	15.8	631,691	222	0	7
cov1076	Strong branch	14400.0	4,989	2,498	2327.4	2,798	1,256	71,682	27
cov1076	Break symmetry	14400.0	496,533	1	0.7	720,913	0	0	0
cov1076	Keep symmetry	1559.9	31,943	786	657.0	21,902	960	0	20
cov1076	Max prod. orbit size	14400.0	498,258	133	110.2	638,795	532	0	18
f5	Branch largest	64.1	1,829	35	0.6	9,710	430	0	11
f5	Branch largest LP	79.8	2,573	31	0.6	7,660	252	0	8
f5	Strong branch	668.2	123	60	1.1	169	736	8,610	15
f5	Break symmetry	42.5	995	117	2.5	3,515	1,356	0	14
f5	Keep symmetry	34.5	717	78	1.5	2,102	598	0	14
f5	Max prod. orbit size	64.3	1,835	38	0.7	9,678	418	0	13
sts45	Branch largest	8.1	4,917	4	0.4	1	0	0	2
sts45	Branch largest LP	7.8	4,683	6	0.6	3	0	0	3
sts45	Strong branch	94.5	1,417	707	43.0	0	0	7,984	16
sts45	Break symmetry	7.6	4,571	11	0.7	1	0	0	4
sts45	Keep symmetry	8.1	4,507	16	1.3	2	0	0	6
sts45	Max prod. orbit size	8.1	4,917	4	0.4	1	0	0	2
sts63	Branch largest	92.7	33,785	15	9.1	19	0	0	7
sts63	Branch largest LP	91.4	32,627	17	12.6	7	0	0	9
sts63	Strong branch	1132.1	3,157	1,577	913.6	0	0	16,858	24
sts63	Break symmetry	1630.3	666,623	10	6.9	720	126	0	43
sts63	Keep symmetry	160.8	9,993	155	135.7	12	0	0	11
sts63	Max prod. orbit size	136.8	31,261	73	57.3	48	0	0	10
sts81	Branch largest	127.0	11,323	28	84.6	0	0	0	13
sts81	Branch largest LP	164.0	25,739	20	68.7	5	0	0	13
sts81	Strong branch	13465.4	11,291	5,644	12074.9	1	0	62,098	30
sts81	Break symmetry	3422.7	1,000,000	5	2.4	235	0	0	4
sts81	Keep symmetry	434.1	83,961	38	128.0	8	0	0	15
sts81	Max prod. orbit size	3370.8	1,000,000	1	0.1	200	0	0	0

Table 6 Performance of orbital branching rules (Global symmetry) on symmetric IPs

Instance	Branching rule	Time	Nodes	Nauty calls	Nauty time	#Fixed by RCF	#Fixed by OF	#Fixed by SBF	Deepest orbital level
cod83	Branch largest	10.4	193	14	0.3	233	588	0	7
cod83	Branch largest LP	3.4	57	18	0.4	328	864	0	7
cod83	Strong branch	5.2	21	9	0.3	16	762	412	6
cod83	Break symmetry	1.3	25	11	0.3	37	906	0	7
cod83	Keep symmetry	1.3	25	11	0.3	37	906	0	7
cod83	Max prod. orbit size	4.6	105	19	0.4	69	642	0	11
cod93	Branch largest	1677.2	16,439	15	1.7	205,636	1,060	0	7
cod93	Branch largest LP	1555.7	14,461	13	1.5	201,292	348	0	7
cod93	Strong branch	2361.0	161	79	3.8	437	2,400	13,478	15
cod93	Break symmetry	165.7	1,361	80	8.4	7,397	3,378	0	14
cod93	Keep symmetry	167.2	1,361	80	8.4	7,397	3,378	0	14
cod93	Max prod. orbit size	395.7	3,503	59	4.1	41,907	704	0	25
cod105	Branch largest	237.3	9	3	5.5	0	0	0	3
cod105	Branch largest LP	237.2	7	2	3.6	0	0	0	2
cod105	Strong branch	359.1	7	2	3.6	0	1,024	1,532	2
cod105	Break symmetry	234.1	11	4	7.1	0	1,020	0	4
cod105	Keep symmetry	242.5	11	4	7.1	0	1,020	0	4
cod105	Max prod. orbit size	237.5	9	3	5.3	1	960	0	3
cov954	Branch largest	5.0	249	20	0.9	818	304	0	12
cov954	Branch largest LP	3.6	153	11	0.5	638	0	0	6
cov954	Strong branch	23.5	63	30	1.3	65	160	1,724	11
cov954	Break symmetry	12.7	373	150	7.1	632	524	0	13
cov954	Keep symmetry	6.2	249	42	1.9	748	48	0	11
cov954	Max prod. orbit size	4.6	217	18	0.8	699	132	0	11
cov1053	Branch largest	102.6	3,437	55	1.2	0	1,094	0	17
cov1053	Branch largest LP	619.3	20,903	56	1.1	0	988	0	19
cov1053	Strong branch	761.0	777	387	7.6	0	2,830	16,464	43
cov1053	Break symmetry	280.5	11,271	1,276	23.7	0	3,454	0	33
cov1053	Keep symmetry	240.2	9,775	248	4.7	0	724	0	25
cov1053	Max prod. orbit size	89.2	2,859	71	1.6	0	1,466	0	20
cov1054	Branch largest	14400.0	105,846	10	1.4	0	0	0	7
cov1054	Branch largest LP	14400.0	104,161	6	0.9	56	88	0	5
cov1054	Strong branch	14400.0	846	430	52.7	0	220	12,846	57
cov1054	Break symmetry	14400.0	110,307	1	0.2	0	0	0	0
cov1054	Keep symmetry	178.8	1,249	103	15.2	0	454	0	15
cov1054	Max prod. orbit size	14400.0	104,184	12	1.7	0	176	0	8
cov1075	Branch largest	54.6	461	43	30.6	1,333	900	0	13
cov1075	Branch largest LP	41.9	495	22	15.7	1,400	520	0	9
cov1075	Strong branch	201.6	71	34	23.9	126	92	1,858	10
cov1075	Break symmetry	14400.0	410,572	1	0.8	865,517	0	0	0

Table 6 continued

Instance	Branching rule	Time	Nodes	Nauty calls	Nauty time	#Fixed by RCF	#Fixed by OF	#Fixed by SBF	Deepest orbital level
cov1075	Keep symmetry	152.2	381	181	133.0	413	962	0	15
cov1075	Max prod. orbit size	95.2	543	98	69.2	1,028	1,090	0	21
cov1076	Branch largest	14400.0	504,849	40	26.2	496,164	388	0	9
cov1076	Branch largest LP	14400.0	496,393	20	13.1	628,579	222	0	7
cov1076	Strong branch	14400.0	5,280	2,642	1692.9	2,971	1,288	76,298	27
cov1076	Break symmetry	14400.0	495,919	1	0.7	719,961	0	0	0
cov1076	Keep symmetry	1415.0	31,943	786	516.2	21,902	960	0	20
cov1076	Max prod. orbit size	14400.0	497,593	133	86.5	637,905	532	0	18
f5	Branch largest	63.8	1,829	35	0.4	9,710	430	0	11
f5	Branch largest LP	79.5	2,573	31	0.4	7,660	252	0	8
f5	Strong branch	664.4	123	60	0.4	169	736	8,610	15
f5	Break symmetry	44.5	1,125	292	4.6	2,983	2,994	0	17
f5	Keep symmetry	44.6	1,125	292	4.6	2,983	2,994	0	17
f5	Max prod. orbit size	63.9	1,835	38	0.5	9,678	418	0	13
sts45	Branch largest	8.1	4,917	4	0.4	1	0	0	2
sts45	Branch largest LP	7.8	4,683	6	0.5	3	0	0	3
sts45	Strong branch	49.9	1,287	642	3.2	0	148	7,150	16
sts45	Break symmetry	7.9	4,709	16	0.8	0	0	0	6
sts45	Keep symmetry	7.8	4,709	16	0.7	0	0	0	6
sts45	Max prod. orbit size	8.1	4,917	4	0.4	1	0	0	2
sts63	Branch largest	103.8	43,349	14	1.7	17	32	0	7
sts63	Branch largest LP	90.1	36,579	16	1.7	19	32	0	9
sts63	Strong branch	101.4	1,377	687	10.5	0	676	6,710	24
sts63	Break symmetry	20.1	5,533	93	5.5	1	308	0	11
sts63	Keep symmetry	20.1	5,533	93	5.6	1	308	0	11
sts63	Max prod. orbit size	81.1	30,133	53	4.6	47	176	0	8
sts81	Branch largest	28.9	5,823	46	5.7	0	410	0	14
sts81	Branch largest LP	27.9	5,649	41	6.0	0	562	0	14
sts81	Strong branch	73.5	573	285	19.8	0	1,112	2,514	22
sts81	Break symmetry	39.1	6,293	112	14.3	0	670	0	17
sts81	Keep symmetry	38.9	6,293	112	14.3	0	670	0	17
sts81	Max prod. orbit size	3382.5	1,000,000	1	0.1	200	0	0	0






Fig. 11 Example 1: Structure of subproblems and orbits in orbital branching

6




$\{2\}\{3\}\{5\}\{8\}\{6, 7\}$
 $\{13, 14\}\{17, 18\}\{19, 20, 21, 22\}$

7




$\{2, 8\}\{3, 7\}\{4, 6\}\{5\}$
 $\{13, 14, 21, 22\}\{17, 18\}$

8




$\{3, 13, 14\}\{5, 8\}\{6, 7\}$
 $\{17, 18, 23, 24\}\{19, 20, 21, 22\}$

9




$\{3, 8\}\{4, 7\}\{5, 6\}$
 $\{13, 14, 23, 24\}\{17, 18, 19, 20\}$

10



$\{3, 13, 14, 5, 17, 18\}\{7, 8\}$
 $\{21, 22, 23, 24\}$

11



$\{3, 13, 14\}\{5, 8\}\{6, 7\}$
 $\{17, 18, 23, 24\}$

Fig. 11 continued

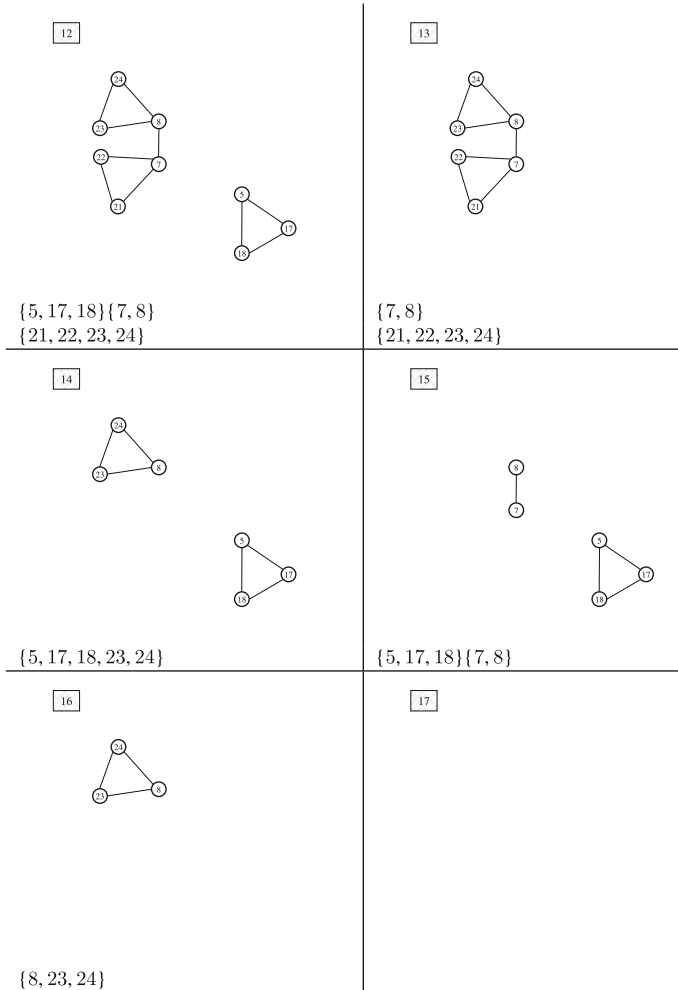


Fig. 11 continued




Fig. 11 continued

References

1. Achterberg, T., Koch, T., Martin, A.: Branching rules revisited. *Oper. Res. Lett.* **33**, 42–54 (2004)
2. Barnhart, C., Johnson, E.L., Nemhauser, G.L., Savelsbergh, M.W.P., Vance, P.H.: Branch and price: column generation for solving huge integer programs. *Oper. Res.* **46**, 316–329 (1998)
3. Cameron, P.J.: *Permutation Groups*. London Mathematical Society, London (1999)
4. Dolan, E., Moré, J.: Benchmarking optimization software with performance profiles. *Math. Program.* **91**, 201–213 (2002)
5. Foggia, P., Sansone, C., Vento, M.: A performance comparison of five algorithms for graph isomorphism. In: *Proc. 3rd IAPR-TC15 Workshop Graph-Based Representations in Pattern Recognition*, pp. 188–199 (2001)
6. Fulkerson, D.R., Nemhauser, G.L., Trotter, L.E.: Two computationally difficult set covering problems that arise in computing the 1-width of incidence matrices of Steiner triples. *Math. Program. Study* **2**, 72–81 (1973)
7. Grove, L.C., Benson, C.T.: *Finite Reflection Groups*. Springer, Heidelberg (1985)
8. Hamalainen, H., Honkala, I., Litsyn, S., Östergård, P.: Football pools—a game for mathematicians. *Am. Math. Monthly* **102**, 579–588 (1995)
9. Holm, S., Sørensen, M.: The optimal graph partitioning problem: solution method based on reducing symmetric nature and combinatorial cuts. *OR Spectrum* **15**, 1–8 (1993)
10. Kaibel, V., Peinhardt, M., Pfetsch, M.E.: Orbital fixing. In: *IPCO 2007: The Twelfth Conference on Integer Programming and Combinatorial Optimization*, pp. 74–88. Springer, Heidelberg (2007)
11. Kaibel, V., Pfetsch, M.E.: Packing and partitioning orbitopes. *Math. Program.* **114**, 1–36 (2008)
12. Linderroth, J.T., Savelsbergh, M.W.P.: A computational study of search strategies in mixed integer programming. *INFORMS J. Comput.* **11**, 173–187 (1999)

13. Litsyn, S.: An updated table of the best binary codes known. In: Pless, V.S., Huffman, W.C. (eds.) *Handbook of Coding Theory* volume 1, pp. 463–498. Elsevier, Amsterdam (1998)
14. Macambira, E.M., Maculan, N., de Souza, C.C.: Reducing symmetry of the SONET ring assignment problem using hierarchical inequalities. Technical Report ES-636/04, Programa de Engenharia de Sistemas e Computação, Universidade Federal do Rio de Janeiro (2004)
15. Margot, F.: Pruning by isomorphism in branch-and-cut. *Math. Program.* **94**, 71–90 (2002)
16. Margot, F.: Exploiting orbits in symmetric ILP. *Math. Program. Ser. B* **98**, 3–21 (2003)
17. Margot, F.: Small covering designs by branch-and-cut. *Math. Program.* **94**, 207–220 (2003)
18. McKay, B.D.: *Nauty User's Guide (Version 1.5)*. Australian National University, Canberra (2002)
19. Méndez-Díaz, I., Zabala, P.: A branch-and-cut algorithm for graph coloring. *Discrete Appl. Math.* **154**(5), 826–847 (2006)
20. Mills, W.H., Mullin, R.C.: Coverings and packings. In: *Contemporary Design Theory: A Collection of Surveys*, pp. 371–399. Wiley, New York (1992)
21. Nemhauser, G.L., Savelsbergh, M.W.P., Sigismondi, G.C.: MINTO, a Mixed INTEger Optimizer. *Oper. Res. Lett.* **15**, 47–58 (1994)
22. Ostrowski, J., Linderoth, J., Rossi, F., Smriglio, S.: Orbital branching. In: *IPCO 2007: The Twelfth Conference on Integer Programming and Combinatorial Optimization. Lecture Notes in Computer Science*, vol. 4517, pp. 104–118. Springer, Heidelberg (2007)
23. Ostrowski, J., Linderoth, J., Rossi, F., Smriglio, S.: Constraint orbital branching. In: Lodi, A., Panconesi, A., Rinaldi, G. (eds.) *IPCO 2008: The Thirteenth Conference on Integer Programming and Combinatorial Optimization. Lecture Notes in Computer Science*, vol. 5035, pp. 225–239 (2008)
24. Rotman, J.J.: *An Introduction to the Theory of Groups*, 4th edn. Springer, Heidelberg (1994)
25. Sewell, E.C.: A branch-and-bound algorithm for the stability number of a sparse graph. *INFORMS J. Comput.* **10**, 438–447 (1998)
26. Sherali, H.D., Smith, J.C.: Improving zero-one model representations via symmetry considerations. *Manage. Sci.* **47**(10), 1396–1407 (2001)