

Pruning by Isomorphism in Branch-and-Cut

François Margot

Department of Mathematics, University of Kentucky
Lexington, KY 40506-0027
fmargot@ms.uky.edu

Abstract. The paper presents a Branch-and-Cut for solving $(0, 1)$ integer linear programs having a large symmetry group. The group is used for pruning the enumeration tree and for generating cuts. The cuts are non standard, cutting integer feasible solutions but leaving unchanged the optimal value of the problem. Pruning and cut generation are performed by backtracking procedures using a Schreier-Sims table for representing the group. Applications to the generation of covering designs and error correcting codes are presented.

1 Introduction

Let I^n be the set of all permutations of the ground set $I^n = \{1, \dots, n\}$. A permutation in I^n is represented by an n -vector π , with $\pi[i]$ being the image of i under π . If v is an n -vector and $\pi \in I^n$, let $w = \pi(v)$ denote the vector w obtained by permuting the coordinates of v according to π , i.e.

$$w[\pi[i]] = v[i] \text{ for all } i \in I^n.$$

We consider an ILP problem of the form

$$\begin{aligned} \min \quad & c^T \cdot x \\ \text{s.t.} \quad & Ax \geq b, \\ & x \in \{0, 1\}^n, \end{aligned} \tag{1}$$

where A is an $m \times n$ matrix. For a permutation π of the n variables such that $\pi(c) = c$ and a permutation σ of the m rows of A such that $\sigma(b) = b$, let $A(\pi, \sigma)$ be the matrix obtained from A by permuting its columns according to π and its rows according to σ . Let

$$G = \{ \pi \mid \text{there exists } \sigma \text{ s.t. } A(\pi, \sigma) = A \}.$$

Clearly, G is a permutation group of I^n . Moreover, for $\pi \in G$, a point \bar{x} is feasible (resp. optimal) for the linear relaxation of the ILP (1) if and only if $\pi(\bar{x})$ is feasible (resp. optimal) for that ILP. Hence, G is a symmetry group of the feasible (and of the optimal) set of the ILP.

ILPs with large symmetry groups are natural when formulating classical problems in combinatorics, a.o. problems looking for a family of subsets of a

given set E with specified properties. In most cases, the elements in E are undistinguishable and G is a group with order at least $|E|!$. The problem of scheduling jobs on p parallel identical machines also yields ILPs with a natural symmetry group with at least $p!$ elements. For relatively modest size problems, it turns out that the corresponding ILPs become very difficult (if not impossible) to solve by traditional Branch-and-Cut techniques. The trouble comes from the fact that many subproblems in the enumeration tree will be isomorphic, forcing a wasteful duplication of effort.

In this paper, we assume that an ILP together with its symmetry group G is given. We show how to use G in order to prune efficiently isomorphic subproblems and to help the search by generating isomorphism cuts (cutting integer feasible solutions, but leaving the value of the optimal solution unchanged). This isomorphism pruning is compatible with standard cut generation techniques (Gomory cuts, Lift-and-Project cuts, or specially designed cuts for the problem at hand). The price to pay for the pruning is that the branching variable can no longer be chosen arbitrarily.

While isomorphism rejection in backtracking searches has been used in many applications [1], [3], [4], [7], [8], [10], [11], [12], [13], [15], [17], [18], it is not commonly used in a Branch-and-Cut ($B\mathcal{E}C$) context. In most instances, the symmetry group G is not assumed to be known and the backtracking search has the additional task to produce it. The originality of the proposed approach resides essentially in (i) the possibility of generating isomorphism cuts (that will be shown to be efficient for the covering design problem), and (ii) the development of algorithms for computing orbits and stabilizers of sets under a group, taking advantage of the type of stabilizers and points in the queries needed by the B&C.

Section 2 describes the pruning algorithm and Section 3 presents basic data structures and algorithms for group operations. Section 4 describes the restrictions that can be put on queries for orbits and stabilizers generated during the B&C. Section 5 introduces the isomorphism cuts. Finally, Section 6 presents results on two applications, covering designs and error correcting codes.

We close this section with two basic definitions and a notation:

Let $S \subseteq I^n$. To simplify the notation, we make no difference between a set S and its characteristic vector. The *orbit* of S under G is

$$orb(S, G) = \{S' \subseteq I^n \mid S' = g(S) \text{ for } g \in G\} .$$

The *stabilizer* of S in G is the subgroup of G given by:

$$stab(S, G) = \{g \in G \mid g(S) = S\} .$$

For $1 \leq a \leq b \leq n$, we write $v[a..b]$ the entries $\{v[a], v[a+1], \dots, v[b]\}$ of v as an unordered set.

2 Isomorphism Test and Pruning

The proposed B&C will branch by fixing the value of one variable x_j to 0 or 1. Since the ILP (1) has a large automorphism group G , it is very likely that

several nodes in the enumeration tree will correspond to isomorphic problems. Obviously, solving one of these isomorphic problems and pruning the others would result in huge savings. One important goal is to do so without having to keep in memory the list of all non isomorphic subproblems encountered since the start of the algorithm. One way to achieve this is to define, for each isomorphism class of subproblems, one particular subproblem (called *representative* of the class) that will be solved. Given a subproblem, we then just need to be able to decide if it is a representative or not and, in the latter case, we can prune the corresponding node of the B&C. Some care must be taken to ensure that the representative subproblems form a subtree of the B&C tree including the root. The general approach of isomorphism free generation of combinatorial structures based on representatives was studied by Read [17]. A general theory for isomorphism free generation, developed by McKay, can be found in [15].

Let a be a node of the B&C enumeration tree. Let F_1^a (resp. F_0^a) be the set of indices of variables fixed to 1 (resp. to 0) at a . Let F^a be the set of indices of variables that are not fixed to 0 or 1 at a , variables also called *free* at a . Let b be an other node and let F_1^b, F_0^b and F^b be the corresponding set of indices of fixed and free variables at b . The subproblems associated with nodes a and b of the B&C are isomorphic if and only if there exist a permutation $g \in G$, such that $g(F_i^a) = F_i^b$ for $i = 0, 1$.

Unfortunately, using this isomorphism test to identify subproblems that can be pruned during the B&C would require the storage of a maximal set of non isomorphic subproblems generated so far in the enumeration. Moreover, the computation needed to find if g exists is not trivial and would be required for many pairs of subproblems. Using the definition of a representative, we can use a slight relaxation of the isomorphism test that turns out to be practical. The price to pay for the simplification is that we will no longer be free to branch on any variable of the ILP: At node a , the branching variable will have to be x_f where f is the minimum index in F^a (even if the value of x_f in the current solution of the LP relaxation is 0 or 1). The variable x_f is called the *branching variable* at a . This branching strategy is called *minimum index branching* (MIB). The enumeration tree (containing only nodes that are not pruned) generated by a Branch-and-Bound \mathcal{B} using the LP relaxation of (1) to prune only infeasible subproblems is called the *full enumeration tree* of \mathcal{B} .

A set $S \subseteq I^n$ is a *representative* if S is lexicographically minimum among the sets in its orbit under G . The following property is crucial for the validity of the pruning:

Lemma 1. *Let $S \subseteq I^n$ be a representative under G . Let $S' := S - v$ with $v = \max \{w \in S\}$. Then S' is also a representative.*

Proof. If S' is not a representative, then there exists $g \in G$ such that $g(S')$ is lexicographically smaller than S' . Then $g(S)$ is lexicographically smaller than S , a contradiction. \square

Consider the following *isomorphism pruning* (IP) to be applied on nodes of the enumeration tree of a B&C: If F_1^a is not a representative, then prune node a .

Lemma 2. *Let τ be the full enumeration tree of a B&C \mathcal{B} using MIB. Let S be the nodes in τ that are not pruned by IP. Then*

- (i) S induces a subtree of τ containing the root of τ ;
- (ii) The B&C \mathcal{B}' obtained by adding IP to \mathcal{B} returns the same optimal value as \mathcal{B} .

Proof. (i): Let $a \in S$ and let $b \in \tau$ on the path between the root and a in τ . Then F_1^a is a representative and, by the choice of branching strategy, F_1^b is the set of the $|F_1^b|$ smallest entries in F_1^a . By Lemma 1, F_1^b is a representative, i.e. $b \in S$.

(ii): Let a be a node of τ for which F_1^a is an optimal solution to ILP (1). Then the representative of the orbit of F_1^a under G is a set F^* , and thus there is a node $b \in S$ with $F_1^b = F^*$. By (i), the full enumeration tree of \mathcal{B}' is the subtree induced by S in τ , implying that \mathcal{B}' will process node b at some point, yielding the same optimal value as the one returned by \mathcal{B} . \square

When solving a subproblem a , it is sometimes possible to identify variables that may be set to 0 without affecting the optimal solution returned by a B&C using MIB and IP. Consider the following operations:

- (i) Let b be the father of a in the enumeration tree and let x_f the branching variable at b . If a is the son of b where x_f is set to 0 then set to 0 all free variables in $\text{orb}(f, \text{stab}(F_1^a, G))$.
- (ii) Let $f = \min \{r \in F^a\}$. If $F_1^a \cup f$ is not a representative, then fix to 0 all free variables in $\text{orb}(f, \text{stab}(F_1^a, G))$.

Applying these operations (repeatedly for (ii) if possible, i.e. until no free variable exists or until $F_1^a \cup f$ is a representative) is called a *0-fixing*. The output of the 0-fixing is the value f in (ii) for which $F_1^a \cup f$ is a representative, or $n + 1$ if no such f exists.

Remark 3. Trivially, the variables set to 0 during a 0-fixing at node a have all an index larger than the maximum index M in F_1^a , since all variables in F^a have index larger than M . \square

Lemma 4. *Consider a B&C \mathcal{B} using MIB and IP and let \mathcal{B}' be the B&C obtained by adding 0-fixing in \mathcal{B} . Then the optimal values returned by \mathcal{B} and \mathcal{B}' are equal.*

Proof. Let a be a node of the full enumeration tree τ of \mathcal{B} for which F_1^a is an optimal solution to ILP (1). Then F_1^a is a representative. Assume that no node b in the full enumeration tree τ' of \mathcal{B}' has $F_1^b = F_1^a$. Hence there exists a node $c \in \tau'$ such that F_1^c contains the $|F_1^c|$ smallest indices in F_1^a and, during the

0-fixing at c , one of the variables in $F_1^a - F_1^c$ is fixed to 0. Assuming that c is chosen as close as possible to the root, we then have $j \in \text{orb}(f, \text{stab}(F_1^c, G))$ for some $j \in F_1^a - F_1^c$ and $f \in F_0^c$ with

$$\max\{r \in F_1^c\} < f < m := \min\{r \in (F_1^a - F_1^c)\} \leq j .$$

The first inequality comes from Remark 3 and the second one from the fact that $F_1^c \cup m$ is a representative: If m is fixed to 0 during the 0-fixing at c , then it is from a $f < m$ and if m is not fixed to 0, then all f considered during the 0-fixing are smaller than m .

Thus there exists $g \in \text{stab}(F_1^c, G)$ such that $g[j] = f$. Then $g(F_1^c \cup j) = F_1^c \cup f$ which is lexicographically smaller than $F_1^c \cup m$, proving that F_1^a is not a representative as $F_1^c \cup j \subseteq F_1^a$, a contradiction. \square

It remains to show how to compute $\text{orb}(f, \text{stab}(F_1^a, G))$ and how to test if a set is a representative or not. This will be covered in Section 4. In the remainder of the paper, the B&C is assumed to use MIB, IP and 0-fixing. The operations performed at node a in the enumeration tree are thus:

```

r := 0-fixing(a);
Repeat until a criterion is met
    solve the LP relaxation;
    generate cuts;
If r < n + 1 then create two sons of a by fixing x_r to 0 or 1;

```

3 Group Representation and Basic Algorithms

Essentially two options are available to represent a permutation group G : The explicit representation or a representation by generators. The explicit representation simply store in a list each permutation in G . A representation by generators store only a subset $\{g_1, \dots, g_k\}$ of the permutations in G , with the property that any permutation in G can be written as a product of permutations in the subset. If $|G|$ is small, the explicit representation might work well, but in most cases of interest a representation by generators is required. The operations of interest listed above are also, usually, faster with the representation by generators.

We use the *Schreier-Sims* representation of G (also called *strong generators*) [1], [2], [3], [4], [10], [11], [12]. Let

$$\begin{aligned} G_0 &= G \\ G_1 &= \{g \in G_0 \mid g[1] = 1\} \\ G_2 &= \{g \in G_1 \mid g[2] = 2\} \\ &\dots \\ G_n &= \{g \in G_{n-1} \mid g[n] = n\} . \end{aligned}$$

G_1 is simply the stabilizer of 1 in G and G_i is the stabilizer of i in G_{i-1} . It follows that G_0, G_1, \dots, G_n are nested subgroups of G .

For $k = 1, \dots, n$, let $\text{orb}(k, G_{k-1}) = \{j_1, \dots, j_p\}$ be the orbit of k under G_{k-1} . Then for each $1 \leq i \leq p$, let h_{k, j_i} be a permutation in G_{k-1} sending k on j_i , i.e. $h_{k, j_i}[k] = j_i$. Let $U_k = \{h_{k, j_1}, \dots, h_{k, j_p}\}$. Note that U_k is never empty as $\text{orb}(k, G_{k-1})$ always contains k .

Arrange the permutations in the sets U_k , $k = 1, \dots, n$ in an $n \times n$ table T , with

$$T_{k,j} = \begin{cases} h_{k,j} & \text{if } j \in \text{orb}(k, G_{k-1}), \\ \emptyset & \text{otherwise.} \end{cases}$$

The table T is called the Schreier-Sims representation of G . This table is not uniquely defined, as there is usually a choice for the permutations included in the sets U_k . However, the general shape of the table (i.e. which entries are empty or not) is fixed.

Remark 5. It is more efficient to implement the table as a vector of ordered lists instead of a 2-dimensional table, as most entries in the table are usually empty. However, algorithms are simpler to describe and understand for the 2-dimensional table. The actual implementation uses a vector of ordered lists. \square

Remark 6. The most interesting property of this representation of G is that each $g \in G$ can be uniquely written as

$$g = g_1 \cdot g_2 \cdot \dots \cdot g_n \tag{2}$$

with $g_i \in U_i$ for $i = 1, \dots, n$. Hence the permutations in the table form a set of generators of G . It is called a strong set of generators, since the equation (2) shows that $g \in G$ can be expressed as a product of at most n permutations in the sets.

Given a permutation $g \in G$, it is easy to find the n permutations g_1, \dots, g_n of equation (2): the permutations g_2, \dots, g_n all stabilize point 1, forcing g_1 to be $T[1, g[1]]$. Then, as g_3, \dots, g_n all stabilize point 2, we must have $(g_1 \cdot g_2)[2] = g[2]$, i.e. $g_2[2] = (g_1^{-1} \cdot g)[2]$ and thus $g_2 = T[2, (g_1^{-1} \cdot g)[2]]$. A similar reasoning yields g_3, \dots, g_n . \square

It is possible to make a small generalization of the presentation by ordering the points of the ground set in an arbitrary order β , called the *base* of the table. In that case, the subgroups $G(\beta)_k$ for $k = 1, \dots, n$ are defined as the stabilizer of $\beta[k]$ in $G(\beta)_{k-1}$, with $G(\beta)_0 = G$. The corresponding table is denoted by $T(\beta)$. Row k of $T(\beta)$ corresponds to the element k , $U(\beta)_k$ is the set of non empty entries in row k of $T(\beta)$ and $J(\beta)_k$ denotes the set of indices $\{j \in I^n \mid T(\beta)[k, j] \neq \emptyset\}$, also called the *basic orbit* of k in T , following the terminology of [12]. When the base β is fixed, we sometimes drop the qualifier (β) in these symbols, but from now on each table T is defined with respect to a base.

Remark 7. For any $k \in \{1, \dots, n\}$, replacing rows $\beta[1], \dots, \beta[k-1]$ in $T(\beta)$ by identity rows yields a Schreier-Sims representation of $G(\beta)_{k-1}$. Hence the permutations on rows $\beta[k], \dots, \beta[n]$ of $T(\beta)$ form a set of generators of $G(\beta)_{k-1}$. \square

Two natural questions arise: How can we create the table $T(\beta)$, knowing the group G either explicitly or by a family of generators, and how can we change the base β of the representation. Algorithms for performing these operations can be found in [1], [3], [4], [10], [11], [12]. The implemented algorithm to create the table is closest to [10], with worst case complexity in $O(n^6)$. The algorithm for changing the base is essentially the algorithm in [4], blended with basic algorithms of [10], taking advantage of the fact that the base changes that arise during the course of the B&C are of a particular kind: Suppose that $T(\beta)$ is the table at a node of the enumeration tree. As we will see in Section 4, the base β' for any of its sons can be obtained through a few applications of the following operation (called *downing of a point v*): Assume that $v = \beta[k]$ and let $k \leq r \leq n$. Let β' be the permutation obtained from β by moving the entry v to position r of β' , keeping the other entries in the same order as in β . The worst case complexity of this modified base change algorithm is in $O(n^4)$ for downing a single point.

An algorithm with worst case complexity of $O(n^6)$ or even $O(n^4)$ might seem impractical for values of $n \geq 100$. It turns out that these bounds are very pessimistic and that the amount of time spent in performing these group operations during the B&C stays well below 5% of the total cpu time in typical applications.

4 Orbits, Stabilizers and Representatives

We are interested in performing the following operations that were mentioned in Section 2: Computing the orbit of a point in the stabilizer of a set and deciding if a set is lexicographically minimum in its orbit under G .

Standard algorithms for computing orbits under a group G' [2], [10] usually return a partition of the ground set into orbits. Here, we only need the orbit of a particular point v and it seems a waste of effort to compute all the disjoint orbits and keep only the one we are interested in. If G' was given by a Schreier-Sims table, it would be of course possible to make a change of basis so that v becomes the first entry of the basis, since, by definition, non empty entries in row v of the table will be the orbit of v under $G'_0 = G'$. In our particular case, however, G' is the stabilizer of a set in G and building the table for G' would be quite expensive.

We thus devised a backtracking algorithm for computing the orbit of a single point in the stabilizer of a set in G . It takes advantage of the fact that we might assume that the basis β of the group at node a of the enumeration tree has the following structure: Variables set to 1 at a (i.e. F_1^a) come first in β , then the free variables (F^a), and then the variables set to 0 at a (F_0^a).

The data structure associated with group G at node a of the B&C is the following:

<code>table:</code> T	<code>base:</code> β
<code>integer:</code> $fixed_one$	<code>vector:</code> $part_zero$.

The table T is just a Schreier-Sims representation of the group with base β . The variable $fixed_one$ gives the number of variables in F_1^a and

$$F_1^a = \beta[1..fixed_one] \quad \text{with} \quad \beta[1] < \cdots < \beta[fixed_one] .$$

The vector $part_zero$ is used to store information about variables fixed to 0. For $i = 1, \dots, fixed_one$, $\beta[part_zero[i]..n]$ are the variables that have been fixed to 0 before $\beta[i]$ was set to 1. For $i = fixed_one + 1$, $\beta[part_zero[i]..n] = F_0^a$, i.e. all the variables currently fixed to 0 at a . The remaining variables (the free ones) appear in β in increasing order of their index, after variables in F_1^a and before variables in F_0^a . Note that this structure of β is easy to maintain throughout the B&C: When the 0-fixing is performed (or a variable is set to 0 by branching), free variables in a set U become fixed to 0. To update the table, simply use the procedure for changing the base by downing a point, moving one by one the variables in U . When a variable is fixed to 1 by branching, it is always the free variable with smallest index, and the basis (and thus the table) remains the same.

The backtracking procedure computes the orbit of $\beta[k]$ in the stabilizer of the points in $\beta[1..k-1]$. Due to the particular structure of the base β , this is exactly the operation of computing $orb(f, stab(F_1^a, G))$ with $f = \min \{r \in F^a\}$ needed in Section 2 if we use $k = |F_1^a| + 1$.

The procedure is a recursive procedure `orb_in_stab()` (described below) called from the following initializing procedure:

```

orbit_in_stabilizer( $n, T, \beta, k$ )
/* Returns the orbit of  $\beta[k]$  in  $stab(\beta[1..(k-1)], G)$  where  $G$ 
is the group represented by  $T$  with base  $\beta$  */

 $J_k$  = basic orbit of  $\beta[k]$  in  $T$ ;
 $ident$  = identity permutation in  $\Pi^n$ ;
 $remain := \beta[1..k-1]$ ;
 $orbit := J_k$ ;
orb_in_stab( $n, T, \beta, k, J_k, ident, remain, orbit, 1$ );
return( $orbit$ );

```

The parameters of the call to `orb_in_stab()` have the following interpretation: $perm$ is a permutation in G sending $\beta[1..ind-1]$ on a subset $B \subseteq \beta[1..k-1]$; $remain$ is the set $perm^{-1}(\beta[1..k-1] - B)$; J_k is the basic orbit of $\beta[k]$ in T ; $orbit$ is the set of points currently known in the orbit of $\beta[k]$ in $stab(\beta[1..(k-1)], G)$; ($orbit$ is passed by reference during the recursive calls;) ind refers to the point $\beta[ind]$ being treated during the current call.

```

orb_in_stab( $n, T, \beta, k, J_k, perm, remain, orbit, ind$ )

  For each  $i \in remain$  do
     $h := T[\beta[ind], i]$ ;
    If  $h \neq \emptyset$  then
       $loc\_remain := remain - i$ ;
       $loc\_remain := h^{-1}(loc\_remain)$ ;
       $loc\_perm := perm \cdot h$ ;
      If  $ind < k - 1$  then
        orb_in_stab( $n, T, \beta, k, J_k, loc\_perm, loc\_remain, orbit, ind + 1$ );
      else
        For each  $j \in J_k$  do  $orbit := orbit \cup perm[j]$ ;

```

Proposition 8. *The algorithm `orb_in_stabilizer()` is correct.*

Proof. Let $S = \beta[1..(k-1)]$. If $k = 1$, $stab(\emptyset, G) = G$ and the orbit of $\beta[1]$ in G is J_1 , as returned by the algorithm. Otherwise, we have $k \geq 2$. By Remark 2, $stab(S, G)$ is generated by all permutations g such that $g(S) = S$ with

$$g = g_1 \cdots g_{k-1} \cdot g_k \cdot h$$

and $g_i \in U_{\beta[i]}$ for $i = 1, \dots, k$, $h \in G_k$. Since $h[\beta[k]] = \beta[k]$,

$$orb(\beta[k], stab(S, G)) = \{v \in I^n \mid v = (g_1 \cdots g_k)[\beta[k]], g_i \in U_{\beta[i]} \text{ for } i = 1, \dots, k, g(S) = S\}.$$

Assume that $g_i = T[\beta[i], j_i]$ for $i = 1, \dots, k-1$. The condition $g(S) = S$ implies $j_1 \in S$. Moreover, if $k \geq 3$ then $(g_1 \cdot g_2)[2] \in S - j_1$ and thus $g_2[2] \in g_1^{-1}(S - j_1)$. In general, for index $2 \leq ind \leq k-1$, we have

$$g_{ind}[ind] \in (g_{ind-1}^{-1}(\dots(g_2^{-1}((g_1^{-1}(S - j_1)) - j_2)) - \dots - j_{ind-1})) . \quad (3)$$

Note that the set in (3) is exactly the parameter *remain* of the call to `orb_in_stab()` with value *ind* as last parameter. That procedure simply selects an index in this set, update *perm* and *remain* and calls itself recursively with *ind + 1* until *ind = k - 1* or no permutation *h* is found. In the former case, $g_1 \cdots g_{k-1}(S) = perm(S) = S$ and it adds $perm[j]$ for all $j \in J_k$. This amounts to compute $(perm \cdot g_k)[\beta[k]]$ for all $g_k \in U_{\beta[k]}$. In the later case, the algorithm backtracks to *ind - 1*, since no permutation in G stabilizes S with the current choice of permutations g_1, \dots, g_{ind-1} . Since at each level in the recursion, all possible choices for g_{ind} are explored, the algorithm indeed returns the desired orbit. \square

Remark 9. As observed in the justification above, the set in (3) is the current set *remain*. A weaker statement about this set is that $remain \subseteq perm^{-1}(S)$, as

$$perm^{-1} = g_{ind-1}^{-1} \cdots g_2^{-1} \cdot g_1^{-1} .$$

\square

Let us now turn to the question of deciding if set $S = \beta[1..k]$ is the lexicographically minimum set in $orb(S, G)$. Note that for $k = |F_1^a| + 1$, this is exactly the same question as deciding if $F_1^a \cup f$ is a representative, with $f = \min \{r \in F^a\}$ mentioned in Section 2. We assume that β has the structure stated at the beginning of this section.

```

first_in_orbit( $n, T, k$ )
/* Returns ‘‘true’’ if and only if  $\beta[1..k]$  is
lexicographically minimum in  $orb(\beta[1..k], G)$  */

  ident := identity permutation in  $\Pi^n$ ;
  remain :=  $\beta[1..k]$ ;
  is_first := true;
  f_in_orb( $n, T, k, ident, remain, 1, is\_first$ );
  return( $is\_first$ );

```

The parameters $perm$, $remain$ and ind in the call to `f_in_orb()` are similar to the same parameters in the call of `orb_in_stab()`. The parameter is_first is passed by reference and is used to stop the procedure as soon as it is known that $\beta[1..k]$ is not lexicographically minimum in $orb(\beta[1..k], G)$.

```

f_in_orb( $n, T, k, perm, remain, ind, is\_first$ )

  If  $is\_first = false$  then return;
  For each  $i \in remain$  do
     $h := T[\beta[ind], i]$ ;
    If  $h \neq \emptyset$  then
       $loc\_remain := remain - i$ ;
       $loc\_remain := h^{-1}(loc\_remain)$ ;
       $loc\_perm := perm \cdot h$ ;
      For each  $j \in loc\_remain$  do
        If  $\beta^{-1}[j] \geq part\_zero[ind + 1]$  then
           $is\_first := false$ ;
          return;
      If  $ind < k$  then
        f_in_orb( $n, T, k, loc\_perm, loc\_remain, ind + 1, is\_first$ );

```

Proposition 10. *The algorithm `first_in_orbit()` is correct.*

Proof. Suppose that the condition

$$\beta^{-1}[j] \geq part_zero[ind + 1]$$

in procedure `f_in_orb()` is satisfied. This condition means that there exists a point j in loc_remain that has been fixed to 0 before fixing $\beta[t]$ to 1 and (if $t \geq 2$)

after fixing $\beta[t-1]$ to 1, for some $t \leq \text{ind} + 1$. Let

$$S := \text{perm}(\beta[1..\text{ind}]) \subseteq \beta[1..k] \quad \text{i.e.} \quad \text{perm}^{-1}(S) = \beta[1..\text{ind}] .$$

Moreover, as pointed out in Remark 9 (the algorithms are similar, so this remark holds here too), $\text{loc_remain} \subseteq \text{perm}^{-1}(\beta[1..k])$ and since it is disjoint from S , we have

$$j = \text{perm}^{-1}[\beta[s]] \quad \text{for some } s \in \{\text{ind} + 1, \dots, k\} .$$

Since j was fixed to 0 before setting $\beta[t]$ to 1, we have, for some $w < \beta[t]$,

$$j \in \text{orb}(w, \text{stab}(\beta[1..t-1], G)) .$$

Hence there exists a permutation

$$p \in \text{stab}(\beta[1..t-1], G) \quad \text{with} \quad p(j) = w .$$

Let $T := \text{perm}(\beta[1..t-1]) \subseteq S$. As $p(\beta[1..t-1]) = \beta[1..t-1]$, we have

$$(p \cdot \text{perm}^{-1})(T) = \beta[1..t-1] \quad \text{and} \quad (p \cdot \text{perm}^{-1})[\beta[s]] = w < \beta[t] .$$

Thus $(p \cdot \text{perm}^{-1})(\beta[1..t-1] \cup \beta[s]) = \beta[1..t-1] \cup w$ is lexicographically smaller than $\beta[1..t]$. It follows that when the algorithm returns “false”, the set $\beta[1..k]$ is indeed not lexicographically minimal in its orbit under G .

Suppose now that the set $\beta[1..k]$ is not lexicographically minimal in its orbit under G . Hence, there is a smallest index $1 \leq t \leq k$ such that $\beta[1..t-1]$ is lexicographically minimal in its orbit under G , but $\beta[1..t]$ is not. Let $p \in G$ such that $p(\beta[1..t])$ is lexicographically smaller than $\beta[1..t]$. By Remark 6, we can write

$$p = h_1 \cdots h_n$$

with $h_i \in U_{\beta[i]}$ for $i = 1, \dots, n$. Observe that the choice of t implies that

$$p(\beta[1..t]) = \beta[1..t-1] \cup w \quad \text{for some } w < \beta[t]$$

and w fixed to 0 before setting $\beta[t]$ to 1. We have

$$p^{-1}(\beta[1..t-1] \cup w) = \beta[1..t] \quad \text{and thus} \quad p^{-1}[w] = \beta[s] \text{ for some } s \in \{1, \dots, t\} .$$

During the recursive calls to $\text{f_in_orb}()$, a permutation perm will occur with $\text{perm}[\beta[i]] = p^{-1}[\beta[i]]$ for $i = 1, \dots, t-1$, namely

$$\text{perm} = h_1 \cdots h_{t-1}$$

with $h_i = T[\beta[i], p^{-1}[\beta[i]]]$ for $i = 1, \dots, t-1$. Let $z := \text{perm}^{-1}[\beta[s]]$. Observe that

$$(\text{perm}^{-1} \cdot p^{-1})[w] = z \quad \text{and} \quad \text{perm}^{-1} \cdot p^{-1} \in \text{stab}(\beta[1..t-1], G) .$$

Hence $z \in \text{orb}(w, \text{stab}(\beta[1..t-1], G))$ and z was fixed to 0 with w (or earlier). It follows that loc_remain contains z and that $\beta^{-1}[z] \geq \text{part_zero}[t]$, implying that the algorithm will return “false”. \square

Crude bounds on the worst case complexities for these two backtracking procedures are $O(n \cdot k!)$ and $O(n \cdot (k + 1)!)$, respectively, but they turn out to be orders of magnitude faster on average, making them practical. (Values of k in the range of 20 to 40 with $n \geq 200$ appear routinely in applications and are handled efficiently).

Remark 11. For clarity, the algorithms `orb_in_stab()` and `first_in_orbit()` were presented separately, but it is possible to take advantage of their similarities to merge them into one single recursive procedure. \square

5 Isomorphism Inequalities

Let a be a node of the enumeration tree and H^a be the set of variables that are not fixed to 0 at node a . Suppose that there exists $J \subseteq H^a$ such that the representative J^* of the orbit of J under $\text{stab}(F_1^a, G)$ is lexicographically smaller than F_1^a . Then if a node b in the descendants of a with $J \subseteq F_1^b$ exists, this node will be pruned by IP. Hence, the *isomorphism inequality*

$$\sum_{j \in J} x_j \leq |J| - 1 \tag{4}$$

is valid in the subtree rooted at a . Moreover, if the whole restricted enumeration tree is explored by a depth-first search, always selecting first the son d where the branching variable is set to 1, then the sets F_1^d are enumerated in lexicographic order, starting with the smallest one. It follows that if an inequality (4) is generated at a , it is valid for the rest of the enumeration, i.e. it can be considered global.

The (exact) separation algorithm for the isomorphism inequalities is similar to the backtracking procedure for testing if a set is lexicographically minimal in its orbit under a stabilizer. Crude estimates for its worst case complexity is in $O(n \cdot |H^a|!)$ but, in practice, it is able to handle efficiently instances with $|H^a| \geq 100$ and $n \geq 200$. The algorithm can also be turned to an heuristic separation algorithm by working with a subset $H \subseteq H^a$ instead of H^a , if needed.

6 Applications

We use the software ABACUS (version 2.2) developed by Thienel [9] as generic implementation of all B&C steps (isomorphism pruning excepted) and the LP solver is CPLEX6.6. We briefly describe preliminary results obtained on two applications: covering designs and error correcting codes.

Let V be a set of elements of cardinality v and let k and t be integers such that $v \geq k \geq t \geq 0$. Let \mathcal{K} be the set of all k -subsets of V and \mathcal{T} be the set of all t -subsets of V . A (v, k, t) -covering design is a collection \mathcal{C} of sets in \mathcal{K} such that each $t \in \mathcal{T}$ is contained in at least one set of \mathcal{C} . A (v, k, t) -covering design \mathcal{C} is *minimum* if the cardinality of \mathcal{C} is as small as possible.

Covering designs have a long history and have applications in statistics, coding theory and combinatorics, among others. Numerous theorems give the value of a minimum covering design under certain assumptions on the parameters (see the survey [16]). Yet, for particular values of the parameters, only lower and upper bounds are available. A case point is the $(10, 5, 4)$ -covering design, for which a lower bound of 50 and an upper bound of 51 are known [6].

Running the described B&C algorithm for the $(10, 5, 4)$ -covering design problem, pruning nodes as soon as their associated LP relaxation has value strictly larger than 50, we obtain a proof that no solution better than the best known solution of 51 exists (see [14] for the ILP formulation and details). The ILP has 252 variables, 384 inequalities and the symmetry group G has order $10! = 3'628'800$. The average number of non empty entries in the Schreier-Sims table over all nodes of B&C is about 550. There are only 313 nodes in the enumeration tree and the cpu time (in seconds) is distributed as follows (the machine used is an HP-J5000 running HP-UX10.20 with two 440MHz PA-8500 RISC CPUs): Total cpu time: 75.99, LP cpu time: 66.13, Pool separation for inactive inequalities: 0.97, Separation for isomorphism inequalities: 4.72, Representative test algorithm: 2.34.

Although the separation for isomorphism inequalities might seem time consuming, this should be balanced with the fact that not using these inequalities makes the B&C enumeration tree grow from 313 nodes to well above 400. (These numbers and running times are better than those in [14] where no 0-fixing and a less general isomorphism pruning were used). It is worth noting that proving that this ILP has no solution with value 50 is not doable by the B&C of CPLEX6.6.

An error correcting binary code with distance d and word length w is a collection \mathcal{C} of binary w -vectors such that the Hamming distance between any pair of vectors in \mathcal{C} is at least d [5]. The maximum number of vectors in \mathcal{C} is denoted by $A(w, d)$. Here also, for small values of w and d , only bounds on $A(w, d)$ are known. For example, $72 \leq A(10, 3) \leq 79$. A simple set partitioning problem with one variable per binary w -vector with at least three 1's yields an ILP with a group of order $w!$. This ILP for finding $A(8, 3)$ is difficult for the B&C of CPLEX6.6 as it needs about half a million nodes and 4 hours CPU to prove optimality of a given solution. The isomorphism pruning algorithm described in this paper, however, does it in 95 nodes and 13 seconds CPU. The ILP has 219 variables, 347 inequalities and the symmetry group G has order $8! = 40'320$. The average number of non empty entries in the Schreier-Sims table over all nodes of the B&C is about 315. The cpu time (in seconds) is distributed as follows: Total cpu time: 13.46, LP cpu time: 11.80, Pool separation for inactive inequalities: 0.04, Separation for isomorphism inequalities: 0.04, Representative test algorithm: 0.79.

References

1. Butler G. "Computing in Permutation and Matrix Groups II: Backtrack Algorithm", *Mathematics of Computation* 39 (1982), 671-680.

2. Butler G., *Fundamental Algorithms for Permutation Groups, Lecture Notes in Computer Science* 559, Springer (1991).
3. Butler G., Cannon J.J., “Computing in Permutation and Matrix Groups I: Normal Closure, Commutator Subgroups, Series”, *Mathematics of Computation* 39 (1982), 663-670.
4. Butler G., Lam W.H., “A General Backtrack Algorithm for the Isomorphism Problem of Combinatorial Objects”, *J. Symbolic Computation* 1 (1985), 363-381.
5. Conway J.H., Sloane N.J.A., *Sphere Packings, Lattices and Groups*, Springer (1993).
6. Etzion T., Wei V., Zhang Z., “Bounds on the Sizes of Constant Weight Covering Codes”, *Designs, Codes and Cryptography* 5 (1995), 217-239.
7. Gibbons P.B., “Computational Methods in Design Theory”, in: *The CRC Handbook of Combinatorial Designs*, Colbourn C.J., Dinitz J.H. (eds.), CRC Press (1996), 718-740.
8. Ivanov A.V., “Constructive Enumeration of Incidence Systems”, *Annals of Discrete Mathematics* 26 (1985), 227-246.
9. Jünger M., Thienel S., “Introduction to ABACUS – A Branch-And-CUt System”, *Operations Research Letters* 22 (1998), 83-95.
10. Kreher D.L., Stinson D.R., *Combinatorial Algorithms, Generation, Enumeration, and Search*, CRC Press (1999).
11. Leon J.S., “On an Algorithm for Finding a Base and a Strong Generating Set for a Group Given by Generating Permutations”, *Mathematics of Computation* 35 (1980), 941-974.
12. Leon J.S., “Computing Automorphism Groups of Combinatorial Objects”, in *Computational Group Theory*, Atkinson M.D. (ed.), Academic Press (1984), 321-335.
13. Luetolf C., Margot F., “A Catalog of Minimally Nonideal Matrices”, *Mathematical Methods of Operations Research* 47 (1998), 221-241.
14. Margot F., “Small Covering Designs by Branch-and-Cut”, Research report 2000-27, Department of Mathematics, University of Kentucky.
15. McKay D., “Isomorph-free Exhaustive Generation”, *Journal of Algorithms* 26 (1998), 306-324.
16. Mills W.H., Mullin R.C., “Coverings and Packings”, in: *Contemporary Design Theory: A collection of Surveys*, Dinitz H., Stinson D.R. (eds.), Wiley (1992), 371-399.
17. Read R.C., “Every One a Winner or How to Avoid Isomorphism Search When Cataloguing Combinatorial Configurations”, *Annals of Discrete Mathematics* 2 (1978), 107-120.
18. Seah E., Stinson D.R., “An Enumeration of Non-isomorphic One-factorizations and Howell Designs for the Graph K_{10} minus a One-factor”, *Ars Combinatorica* 21 (1986), 145-161.