

New refiners for permutation group search

Christopher Jefferson, Markus Pfeiffer

*University of St Andrews
School of Computer Science
North Haugh
St Andrews
KY16 9SX
Scotland*

Rebecca Waldecker

*Martin-Luther-Universität Halle-Wittenberg
Institut für Mathematik
06099 Halle
Germany*

Abstract

We describe how orbital graphs can be used to improve the practical performance of many algorithms for permutation groups, including intersection and stabilizer problems. First we explain how orbital graphs can be integrated in partition backtrack, the current state of the art algorithm for many permutation group problems. We then show how our algorithms perform in practice, demonstrating improvements of several orders of magnitude for some problems.

Keywords: Backtrack search, refiners, permutation groups, algorithmic group theory, computational algebra, partition backtrack.

1. Introduction

Permutation groups are one of the most natural and convenient representations of finite groups. They have proved particularly useful as a basis for practical computations, and systems such as GAP (2016) and Magma (Bosma et al. (1997)) provide efficient implementations of many algorithms to solve a range of different problems.

Given a permutation group acting on a finite set Ω , some problems – for example checking if a group contains a particular permutation or computing the size of the group – can be done in time polynomial in the size of Ω . There are other problems on permutation groups where no theoretically efficient algorithm is known: examples of such problems are permutation group

Email addresses: caj21@st-andrews.ac.uk (Christopher Jefferson), markus.pfeiffer@st-andrews.ac.uk (Markus Pfeiffer), rebecca.waldecker@mathematik.uni-halle.de (Rebecca Waldecker)

URL: <http://caj.host.cs.st-andrews.ac.uk/> (Christopher Jefferson),

<https://www.morphism.de/~markusp/> (Markus Pfeiffer),

<http://conway1.mathematik.uni-halle.de/~waldecker/index-english.html> (Rebecca Waldecker)

Preprint submitted to Elsevier

May 26, 2017

intersection, computation of centralizers and normalizers, set stabilizers, and partition stabilizers. It is known (see for example Chapter 3 in Seress (2003)) that these problems are at least as difficult as graph isomorphism, so it is unlikely that a theoretically efficient algorithm will be found.

However, there is a history of algorithms which are of practical usefulness. The current state of the art algorithm for the problems listed above is partition backtrack (see Leon (1991)), which extends the partition refinement ideas for graph isomorphism introduced by McKay (1980). For most problems partition backtrack provides much better performance than previous techniques (see for example Butler (1983)).

However, this does not mean that all permutation group problems can be solved easily or quickly using partition backtrack – and this is where our orbital graph methods come into play. The ideas presented in this paper improve performance by several orders of magnitude for many problems.

The concept of a **refiner** is at the heart of partition backtrack, which should not be viewed as a single monolithic algorithm, but as a combination of a number of smaller algorithms. Refiners are used to detect and skip parts of the search which are superfluous: the better the refiner, the more search can be skipped.

We explain briefly how refiners work in practice and show how orbital graphs lead to very useful refiners.

These new refiners can then improve all partition backtrack implementations.

This article is organized as follows: Section 2 includes all the necessary notation and gives examples, Section 3 introduces orbital graphs and collects some of their properties. In Section 4 we start with a brief description of backtrack search and the roles of refiners before we look at specific refiners. We illustrate the method of refinements via the orbital graph in an example by hand before we go deeper into the theory and prove results about the usefulness of orbital graphs for refiners.

Finally we explain and present our experiments and show their outcomes in Section 5. We briefly discuss the limits of this method – for highly transitive groups, the orbital graph does not give useful information, and there are also cases where the costs of calculating the orbital graphs outweigh the advantages in the search. Regardless, for difficult search problems in large groups our methods prove to be very effective. At the end of the paper we will comment on some related questions and further problems.

Acknowledgements.

All authors thank the DFG (**Wa 3089/6-1**) and the EPSRC CCP CoDiMa (**EP/M022641/1**) for supporting this work. The first author would like to thank the Royal Society, and the EPSRC (**EP/M003728/1**). The second author would like to acknowledge support from the OpenDreamKit Horizon 2020 European Research Infrastructures Project (#676541). The third author wishes to thank the Computer Science Department of the University of St Andrews for its hospitality during numerous visits, and Karin Helbich for the pictures in this article.

2. Notation, basic results and examples

We mostly use standard notation for permutation groups and related objects and refer the reader to references such as Dixon and Mortimer (1996).

Throughout we let Ω be a finite set and $n \in \mathbb{N}$. We use $\text{Sym}(\Omega)$ as notation for the **symmetric group on Ω** and \mathcal{S}_n for the symmetric group on $\{1, \dots, n\}$. If Δ is a subset of $\{1, \dots, n\}$, then we denote by G_Δ the **setwise stabilizer in G of the set Δ** .

Definition 1.

- We introduce the set $O\text{Part}(\Omega)$ of **all ordered partitions of Ω** . By an **ordered partition** we mean an ordered list of disjoint subsets (called **cells**) of Ω whose union is all of Ω .

The notation will be explained in Example 2.

- We say that Q is **finer than** P if for all i and j in Ω , if i and j are in a common cell of Q then they are also in a common cell of P . Conversely, we say that P is **coarser** than Q in this situation.

The above relation defines a partial order on $O\text{Part}(\Omega)$, which we denote by $Q \preceq P$. Note that, for all $P \in O\text{Part}(\Omega)$, the definition implies that $P \preceq P$.

We call an ordered partition **discrete** if and only if every element of Ω is in a cell by itself, and we call an ordered partition **trivial** if it only has one cell.

- We denote the **meet of two ordered partitions P and Q** by $P \wedge Q$. The meet is defined as follows: Two elements are in a cell of $P \wedge Q$ if and only if they are in the same cell of P and in the same cell of Q . For every cell of $P \wedge Q$ there is a unique pair of cells of P and Q that its elements are from, and the cells of $P \wedge Q$ are ordered lexicographically with respect to these pairs.

- If P is an ordered partition of Ω and $g \in G$, then we write P^g for the ordered partition that we obtain by applying g to the elements in the cells of P . In this sense, we denote by G_P the **stabilizer of P in G** , i.e. the set of permutations in G that map P to itself.

- Important ordered partitions come from the action of subgroups of $\text{Sym}(\Omega)$ on Ω . We call them **ordered orbit partitions**. If $H \leq \text{Sym}(\Omega)$ and $P \in O\text{Part}(\Omega)$, then we say that P is an ordered orbit partition for H if and only if the cells of P are exactly the orbits of H on Ω . Note that, if H has more than one orbit, then there are several distinct ordered orbit partitions, differing only by the ordering of the cells.

- Suppose that P is an ordered orbit partition of $\{1, \dots, n\}$ into $k \in \mathbb{N}$ cells $\Delta_1, \dots, \Delta_k$.

Then we write $\text{Sym}(P)$ for the subgroup $\text{Sym}(\Delta_1) \times \dots \times \text{Sym}(\Delta_k)$ of $\text{Sym}(\Omega)$, i.e. the stabilizer of the ordered partition P in $\text{Sym}(\Omega)$. We use $\text{Co}(P, Q)$ for the set of permutations in $\text{Sym}(\Omega)$ that map P to Q . Note that $\text{Co}(P, Q)$ will either be empty, or a coset of $\text{Sym}(P)$. Moreover $\text{Co}(P, P) = \text{Sym}(P)$.

This is why our partitions must be ordered – ordered partitions allow us to represent any coset of $\text{Sym}(P)$, for any ordered partition P .

We remark that $\text{Sym}(P)$ stabilizes each cell of the ordered partition, so elements of $\text{Sym}(P)$.

- We use $\text{Sym}(P)$ if we are interested in the elements in the cells of the ordered partition. If we are not, for example because we are only interested in the structure or size of the group in question, then we take the sizes l_1, \dots, l_k of the cells, in order, and we write $\text{Sym}(l_1, \dots, l_m)$ for the direct product $\mathcal{S}_{l_1} \times \dots \times \mathcal{S}_{l_k}$. It will be a convention that we do not write down factors of size 1 in this direct product, so we omit cells of size 1. This is because $\text{Sym}(P)$ fixes all cells of size one, hence the permutations of these cells are trivial and do not contribute to our group $\text{Sym}(P)$ or to the (isomorphic) group $\text{Sym}(l_1, \dots, l_m)$.

Example 2. We look at the ordered partition $P := [1, 2, 3, 4|5, 6, 7]$ of $\Omega := \{1, 2, 3, 4, 5, 6, 7\}$.

Then $P = [2, 1, 3, 4|5, 6, 7]$, because the ordering of elements within a cell is irrelevant, but $P \neq [5, 6, 7|1, 2, 3, 4]$, as the ordering of cells is relevant. The ordering of cells is not relevant for \preceq , therefore $P \preceq [5, 6, 7|1, 2, 3, 4]$.

Let $Q := [1, 2, 3|4|5, 6, 7]$. Then $Q \preceq P$ and $Q \preceq [4, 5, 6, 7|3, 2, 1]$.

Now we calculate the meet of P and $[1, 2|5, 3|7, 4, 6]$. For each element $i \in \Omega$, we find the index of the cells which contain i . This generates the pairs:

$1 - (1, 1), 2 - (1, 1), 3 - (1, 2), 4 - (1, 3), 5 - (2, 2), 6 - (2, 3), 7 - (2, 3)$.

Therefore $P \wedge [1, 2|5, 3|7, 4, 6] = [1, 2|3|4|5|6, 7]$.

We look at its stabilizer in \mathcal{S}_7 of the ordered partition P . If $H_1 \leq \mathcal{S}_7$ is the subgroup stabilizing the set $\{1, 2, 3, 4\}$ and $H_2 \leq \mathcal{S}_7$ is the subgroup stabilizing the set $\{5, 6, 7\}$, then $\text{Sym}(P)$ is exactly $H_1 \times H_2$.

Finally, we look at the subgroup $H := \langle (123), (57) \rangle$ of \mathcal{S}_7 and we write down two of its ordered orbit partitions: $[123|4|57|6]$ and $[4|6|57|123]$.

It will be important later that taking meets of ordered partitions of Ω is compatible with the action of elements of $\text{Sym}(\Omega)$ on Ω .

Lemma 3. Let $G \leq \text{Sym}(\Omega)$, let P and Q be ordered partitions of Ω , and let $g \in G$.

Then $(P \wedge Q)^g = P^g \wedge Q^g$.

Proof. For any ordered partition P and elements i and j of Ω we write $P(i, j)$ to denote that i and j are in the same cell of P . First we note the following: If $g \in \text{Sym}(\Omega)$, then $P^g(i, j)$ if and only if $P(i^{g^{-1}}, j^{g^{-1}})$, and, by definition, $(P \wedge Q)(i, j)$ if and only if $P(i, j)$ and $Q(i, j)$.

Now we see:

$$\begin{aligned} (P \wedge Q)^g(i, j) &\Leftrightarrow (P \wedge Q)(i^{g^{-1}}, j^{g^{-1}}) \\ &\Leftrightarrow P(i^{g^{-1}}, j^{g^{-1}}) \text{ and } Q(i^{g^{-1}}, j^{g^{-1}}), \\ &\Leftrightarrow P^g(i, j) \text{ and } Q^g(i, j) \\ &\Leftrightarrow (P^g \wedge Q^g)(i, j). \end{aligned}$$

For an ordered orbit partition P denote by $\Delta_P(i)$ the cell that the point i is in. We note that it holds that $\Delta_{P^g}(i) = \Delta_P(i^{g^{-1}})$.

For a meet $P \wedge Q$ of two ordered partitions we can describe the cell $\Delta_{P \wedge Q}(i)$ by the pair $[\Delta_P(i), \Delta_Q(i)]$. So we argue as follows:

$$\begin{aligned} \Delta_{(P \wedge Q)^g}(i) &= \Delta_{(P \wedge Q)}(i^{g^{-1}}) \\ &= [\Delta_P(i^{g^{-1}}), \Delta_Q(i^{g^{-1}})] \\ &= [\Delta_{P^g}(i), \Delta_{Q^g}(i)] \\ &= \Delta_{P^g \wedge Q^g}(i). \end{aligned}$$

This implies that $(P \wedge Q)^g = P^g \wedge Q^g$. □

2.1. Orbital graphs

Definition 4. For the purposes of this paper, a **digraph** Γ is a pair $\Gamma = (V, A)$ where V denotes the set of **vertices (or points)** and A denotes the set of **arcs**, i.e. directed edges. If $x, y \in V$, then an arc from x to y in Γ will be denoted by (x, y) .


We refer the reader to Bang-Jensen and Gutin (2008) for standard notation and for the definitions of connected components, graph isomorphisms etc.

Let G be a group of permutations on a set Ω and let $x, y \in \Omega$ be distinct elements, chosen in this order. We look at the digraph $\Gamma = (\Omega, A)$ where the set of arcs A is the set of all pairs (x^g, y^g) where $g \in G$. This digraph is called the **orbital graph of G with base-pair** (x, y) .


Example 5. If we build an orbital graph for S_3 , then for each base-pair we obtain the complete digraph on $\{1, 2, 3\}$. The reason is that this group is 2-transitive: Given $x, y \in \{1, 2, 3\}$ such that (x, y) is a base-pair, and given any distinct $a, b \in \{1, 2, 3\}$, there exists some $g \in S_3$ such that $i^g = a$ and $j^g = b$. Therefore (a, b) is also an arc, and this means that all possible arcs exist.

For groups that are not 2-transitive, the choice of the base-pair becomes much more important. Let $G := \langle (123), (45), (46) \rangle \leq S_6$.


Then, starting with the base-pair $(1, 2)$, we obtain the following digraph:



But, starting with the base-pair $(5, 6)$, we find:



If we want to obtain a connected digraph, then we could start with the pair $(2, 4)$:



We note that orbital graphs have been used in Theißen's PhD thesis (Theißen (1995)) for normalizer search.

3. Refining Ordered Partitions

3.1. Partition Backtrack

Partition backtrack is a technique for solving search problems on permutation groups. It is implemented in GAP (2016) and Magma (Bosma et al. (1997)) as the main technique for solving

a range of problems, including computing group and coset intersections, set and partition stabilizers, centralizers, and normalizers. This paper will not provide a full description of partition backtrack, instead we refer readers to Leon (1991).

In brief, partition backtrack searches for elements of $\text{Sym}(\Omega)$ that satisfy a list of properties. These properties will typically be of the form “element lies in a subgroup of $\text{Sym}(\Omega)$ ” or “element lies in a coset of a subgroup of $\text{Sym}(\Omega)$ ”. The subgroup can be expressed in a variety of ways, for example by a set of generators, as the normalizer of a group, as the automorphism group of a graph, or as the stabilizer of a set or ordered partition in $\text{Sym}(\Omega)$.

More complex problems are split up into a list of properties, for example the stabilizer of a set S in a subgroup G of $\text{Sym}(\Omega)$ can be expressed as the list of properties “stabilizes S in $\text{Sym}(\Omega)$ ” and “contained in G ”.

Partition backtrack takes a list of properties as input, then it starts from the pair (P_0, Q_0) of trivial ordered partitions and proceeds to perform search by repeatedly alternating between the following two phases, starting at $i = 0$:

1. **Refinement phase:** Start with the first property of the list and refine the pair of ordered partitions (P_i, Q_i) to a new pair (P, Q) such that the following holds: Every element of $\text{Sym}(\Omega)$ in $\text{Co}(P_i, Q_i)$ is also in $\text{Co}(P, Q)$. Hence, we have not removed any permutations which satisfy all properties on the list.

Repeat this process with the second property from the list and continue through all properties. Start again at the beginning of the list until the ordered partition does not change anymore. The resulting pair of ordered partitions will be called (P_{i+1}, Q_{i+1}) .

2. **Branching phase:** Take P_{i+1} and Q_{i+1} . At this point, three cases are possible:

- (a) There is no permutation in $\text{Sym}(\Omega)$ which maps P_{i+1} to Q_{i+1} . This means that this part of the search does not produce any permutation which satisfies all properties on the list.
- (b) Every cell in P_{i+1} and Q_{i+1} is of size one. Then there is exactly one permutation which maps P_{i+1} to Q_{i+1} . Perform a final check that this satisfies all properties on the list, and if it does, then record it as a solution.
- (c) Split the search, by producing a list of pairs of ordered partitions $(P'_1, Q'_1), \dots, (P'_k, Q'_k)$ where the $\text{Co}(P'_j, Q'_j)$ for $j \in \{1, \dots, k\}$ form a disjoint union of $\text{Co}(P_{i+1}, Q_{i+1})$.

The practical method that we choose for splitting is the following:

We choose $l \in \mathbb{N}$ such that the l -th cell c of P_{i+1} has size at least 2, and we choose a single element a from c . We make all the P'_j equal to P_{i+1} , except a is removed from the cell c and placed in a new cell at the end, by itself. For each element b of the l -th cell d of Q_{i+1} , create Q'_j which is identical to Q_{i+1} except b is removed from cell d and placed in a new cell at the end, by itself.

This gives new pairs of ordered partitions that are finer than (P_{i+1}, Q_{i+1}) . For each of these pairs, the search along this branch continues by going to the refinement phase.

This process will stop eventually because every branch of the search will either fail or produce a single permutation. In Leon (1991) the author explains how this algorithm can be implemented efficiently, and how we can find only generators of the group (or coset) for the elements of $\text{Sym}(\Omega)$ that have the desired properties.

This paper will concentrate on refinement because the quality of refiners is one of the main influences on performance. In general, refiners only operate on a single ordered partition – Lemma 7 explains how refiners are used on cosets.

Definition 6. Let $M \subseteq \text{Sym}(\Omega)$ be a set of elements with a given property. An M -refiner is a map $f_M : \text{OPart}(\Omega) \rightarrow \text{OPart}(\Omega)$ which satisfies the following conditions for any ordered partition P of Ω :

- $f(P) \preceq P$.
- For all $g \in M$, it is true that $f(P^g) = f(P)^g$.

The second property implies that every element in $\text{Sym}(P)$ that satisfies the property and stabilizes P also stabilizes $f(P)$. It turns out to be a very natural condition – our experience is that refiners tend to satisfy this property. It is not only a natural requirement for a refiner, but it is also one of the main reasons why partition backtrack is so efficient. For more details see Sections 6 and 7 in Leon (1991).

Lemma 7 shows how, using an M -refiner (which refers to only a single ordered partition), we can perform filtering on cosets, as required by our search defined above. This greatly simplifies our implementation.

Lemma 7. Let P, Q be a pair of ordered partitions of Ω and let f_M be an M -refiner for a property M . Then for all $g \in M$ such that $g \in \text{Co}(P, Q)$, it follows that $g \in \text{Co}(f_M(P), f_M(Q))$.

Proof. If $g \in \text{Co}(P, Q)$, then $Q = P^g$. Then the definition of an M -refiner implies that $f_M(Q) = f_M(P^g) = f_M(P)^g$, and therefore $\text{Co}(f_M(P), f_M(Q)) = \text{Co}(f_M(P), f_M(P)^g)$, and so $g \in \text{Co}(f_M(P), f_M(Q))$. \square

We now look at one of the standard refiners and then introduce a new refiner that uses orbital graphs.

Definition 8. Given $\text{Sym}(\Omega)$ and a subset M of elements that have a given property, the map $\text{Fixed}_M : \text{OPart}(\Omega) \rightarrow \text{OPart}(\Omega)$ is defined as follows:

- Given $P \in \text{OPart}(\Omega)$, let $k \in \mathbb{N}$ and $\alpha_1, \dots, \alpha_k \in \Omega$ be the elements in singleton cells of P . Then consider $M_0 := M \cap \text{Sym}(\Omega)_{\alpha_1, \dots, \alpha_k}$.
- Return the meet of P and an ordered orbit partition of $\langle M_0 \rangle$.

Note that elements of M_0 stabilize the singleton cells of P .

This map does not necessarily give a refiner as it is: the reason is that we do not define the ordering of the cells in the ordered orbit partition. In the implementation this is usually fixed by producing a list of orbits, outputting the cells in arbitrary order and then fixing this ordering for later instances. The details are in the proof of the next lemma.

Lemma 9. Let $M \subseteq \text{Sym}(\Omega)$ be a subset of elements that have a given property. Then the map Fixed_M is a refiner.

Proof. Let P be an ordered orbit partition of Ω . Then $\text{Fixed}_M(P) \preceq P$, because $\text{Fixed}_M(P)$ is the meet of P with another ordered partition and it is therefore finer than P .

Let $k \in \mathbb{N}$ and $\alpha_1, \dots, \alpha_k \in \Omega$ be such that these are precisely the elements in singleton cells of P . Now we let $M_0 := M \cap \text{Sym}(\Omega)_{\alpha_1, \dots, \alpha_k}$ and we note again that $F := \langle M_0 \rangle$ stabilizes all singleton cells of any ordered orbit partition of M .

Let $g \in G$. We need to show that $\text{Fixed}_M(P^g) = \text{Fixed}_M(P)^g$. First we note that F^g fixes $\alpha_1^g, \dots, \alpha_k^g$, which are exactly the singleton cells of P^g . Now we fix an ordered orbit partition

$O(F)$ of F as described before the lemma. This automatically fixes an ordered orbit partition $O(F^g) = O(F)^g$ of F^g .

Using Lemma 3 we deduce:

$$\begin{aligned} \text{Fixed}_M(P)^g &= (P \wedge O(F))^g \\ &= P^g \wedge O(F)^g \\ &= P^g \wedge O(F^g) \\ &= \text{Fixed}_M(P^g). \end{aligned}$$

□

The major limitation of Fixed is that it ignores non-singleton cells. More concretely, given a transitive group G and an ordered partition P which contains no singleton cells, $\text{Fixed}_G(P) = P$. We cannot easily use the same strategy as Fixed for non-singleton cells (finding the stabilizer of the non-singleton cells in G), because this would require solving the set stabilizer problem, which is exactly one of the problems which is solved via backtrack!

Instead, we look at other properties of groups we can use, which allow us to refine non-singleton cells efficiently. We will now show how orbital graphs can be used in refiners which complement existing refiners for groups expressed by a list of generators. These refiners will provide useful refinement even with transitive groups and non-singleton cells.

Theißen (1995) uses orbital graphs as an ingredient for refiners for normalizer search.

We will employ existing refiners for arbitrary graphs, as discussed in McKay (1980) and McKay and Piperno (2014), to create refiners for groups given as a set of generators.

Before we go into the details, we need some more definitions:

Definition 10. Let $\Gamma = (\Omega, E)$ be a directed graph.

For any ordered partition $P \in \text{OPart}(\Omega)$, we say that a cell Δ of P is Γ -**equitable** if, for all cells Δ' of P there is a $k \in \mathbb{N}$ such that for all elements i in the cell Δ it holds that $|\{(i, j) \in E \mid j \in \Delta'\}| = k$.

We note that in this definition the number of edges k depends on Δ and Δ' , but not on the individual vertices in Δ .

An ordered orbit partition P is called Γ -**equitable** if all its cells are Γ -equitable.

Next we suppose that P and P' are ordered partitions of Ω . We say that P' is a Γ -**equalizer for** P if and only if P' is Γ -equitable, moreover $P' \preceq P$ and P' is as coarse as possible with this property, meaning that whenever $Q \in \text{OPart}(\Omega)$ is also Γ -equitable and $Q \preceq P$, then $Q \preceq P'$.

Remark 11. Some remarks on the previous definition:

If $\Gamma = (\Omega, E)$ is a directed graph and $P \in \text{OPart}(\Omega)$, then two distinct Γ -equalizers for P can only differ by ordering of the cells.

We also note that, if P is Γ -equitable and g is an automorphism of Γ , then P^g is also Γ -equitable.

McKay and Piperno (2014) present algorithms for calculating Γ -equitable ordered partitions. We will present here a simple algorithm that computes a Γ -equalizer for P given Γ and P as input. The major simplification over the algorithm in McKay and Piperno (2014), and our implementation, is that there are many situations where we can skip some attempts to split \bar{P} by elements of T , as we know no splitting will occur.

Algorithm 1 Equitable Partitions

```
1: procedure EQUITABLE( $\Gamma, P$ )
2:    $\bar{P} := P$ 
3:    $T := P$ 
4:   while ( $T$  not empty) and ( $\bar{P}$  is not discrete) do
5:     Pick and remove some cell  $\Delta \in T$ 
6:     for  $\Delta' \in \bar{P}$  do
7:       Split  $\Delta'$  into  $\Delta'_1 \dots \Delta'_k$  equitably, according to edges starting at vertices in  $\Delta$ 
8:       if  $k > 1$  then
9:         Replace the cell  $\Delta'$  in  $\bar{P}$  with  $\Delta'_1 \dots \Delta'_k$ 
10:      Add  $\Delta'_1, \dots, \Delta'_k$  to  $T$ 
11:  return  $\bar{P}$ 
```

Using orbital graphs and ordered equitable partitions we define the following refiner:

Definition 12. Given $\text{Sym}(\Omega)$ and a subset M of elements that have a given property, the map $\text{Orb}_M : \text{OPart}(\Omega) \rightarrow \text{OPart}(\Omega)$ is defined as follows:

- Construct all orbital graphs of $\langle M \rangle$.
- Given an ordered orbit partition P of Ω , compute a Γ -equalizer for P for every orbital graph Γ from the previous step, using Algorithm 1.
- Return the meet of all refined ordered partitions from the previous step.

We use the notation Orb , without a subscript, for refiners for subgroup search.

We argue now that this map really gives a refiner, and a detailed example later in this section will explain how this refiner works in practice.

Lemma 13. Orb_M is a refiner.

Proof. For any ordered partition P of Ω , it holds that $\text{Orb}_M(P) \preceq P$, because Algorithm 1 splits cells of P .

Next, we need to show that for any $g \in M$ it holds that $\text{Orb}_M(P^g) = \text{Orb}_M(P)^g$, but this follows directly from the fact that g is an automorphism of any orbital graph, and that g commutes with taking meets of ordered partitions. \square

One obvious limitation of Orb is that, if the group is 2-transitive, then it does not perform any refinement, as the only orbital graph is the complete graph on Ω . We therefore introduce another orbital graph based refiner, which makes use of the fact that we can, like in Fixed , easily stabilize points in the group.

Definition 14. Given $\text{Sym}(\Omega)$ and a subset M of elements that have a given property, the map $\text{DeepOrb}_M : \text{OPart}(\Omega) \rightarrow \text{OPart}(\Omega)$ is defined as follows:

- Given $P \in \text{OPart}(\Omega)$, let $k \in \mathbb{N}$ and $\alpha_1, \dots, \alpha_k \in \Omega$ be the elements in singleton cells of P . Then consider $M_0 := M \cap \text{Sym}(\Omega)_{\alpha_1, \dots, \alpha_k}$.
- Construct all orbital graphs of $\langle M_0 \rangle$.

- Given an ordered orbit partition P of Ω , compute a Γ -equalizer for P for every orbital graph Γ from the previous step, using Algorithm 1.
- Return the meet of all refined ordered partitions from the previous step.

We use the notation DeepOrb , without a subscript, for this refiner.

DeepOrb can be seen as combining Orb and Fixed . The proof of correctness of DeepOrb is analogous to the proof for Orb . The major disadvantage of DeepOrb is that it requires calculating the orbital graphs at every position in search, rather than just once at the beginning. Our experiments will investigate the practical trade-offs between Orb and DeepOrb .

3.2. Examples of refinements via graphs

We will now present two examples of refinement. In each case, the Fixed refiner will perform no refinement at all, but orbital graphs provide useful refinement.

In the following two examples we let $G := \mathcal{S}_{10}$.

Example 15. We illustrate our method with the subgroups

$H_1 := \langle (1, 2, \dots, 10), (2, 10)(3, 9)(4, 8)(5, 7) \rangle$ and $H_2 := G_{\{1,5\}}$ of G . This problem is equivalent to calculating the set stabilizer of $\{1, 5\}$ in H_1 . While this problem is very simple, it allows us to show the algorithm working in detail.


We are interested in calculating $D := H_1 \cap H_2$ by looking at ordered partitions. Purely group-theoretically, if we take $x \in D$, then all we know is that x stabilizes the set $\{1, 5\}$ and so we obtain the ordered partition $P_1 := [1, 5 \mid 2, 3, 4, 6, 7, 8, 9, 10]$.

Looking at the orbits of H_1 produces no useful information, as H_1 is transitive and therefore $\text{Fixed}_{H_1}(P_1) = P_1$. Therefore the only information that we extract from reasoning about orbits alone is that D is contained in a subgroup of G which is isomorphic to $\text{Sym}(2, 8)$.

Now we use the orbital graph for H_1 that is obtained by starting with the edge $(1, 2)$. For simplicity we work with an undirected graph here because all edges exist in both directions. When calculating equitable graphs, we will only show steps where the algorithm causes a cell of an ordered orbit partition to split, skipping steps where no split occurs.

Step 1:

Using P_1 , we attach the colour 1 to the vertices 1 and 5 and colour 2 to all other vertices, to denote the cell of the ordered partition the vertices were contained in. So, as above, our first ordered partition is $[1, 5 \mid 2, 3, 4, 6, 7, 8, 9, 10]$ and again we see that D is isomorphic to a subgroup of $\text{Sym}(2, 8)$.



Step 2:

Running through the Equitable Partition algorithm, we look at the vertices of colour 2. They fall into two classes – those who have two neighbours of colour 2, and those who have a neighbour of colour 1 and a neighbour of colour 2. This splits the second cell of P_1 into two cells, with 2, 4, 6 and 10 in a new cell, with colour 3.

Our second ordered partition is therefore $[1, 5 \mid 3, 7, 8, 9 \mid 2, 4, 6, 10]$ and this means that D is isomorphic to a subgroup of $\text{Sym}(2, 4, 4)$.

Step 3:

Continuing the algorithm, we see the vertices of colour 2 can be further divided into a single vertex which has two neighbours of colour 2 (vertex 3), those who have two neighbours of colour 3 (vertex 8) and those who have one neighbour of colour 2 and one of colour 3 (vertices 7 and 9). Using this information, we obtain the new ordered partition $[1, 5 \mid 7, 9 \mid 2, 4, 6, 10 \mid 3 \mid 8]$.

Now D is isomorphic to a subgroup of $\text{Sym}(2, 2, 4)$.

Step 4:

Our algorithm continues, looking at cell 3. Here the vertices can be divided into two categories, those which are a neighbour of colour 1 and a neighbour of colour 4 (vertices 2 and 4), and those with a neighbour of colour 1 and a neighbour of colour 2 (vertices 6 and 10). Our final ordered partition is therefore:

$$P_2 := [1, 5 \mid 8 \mid 6, 10 \mid 3 \mid 7, 9 \mid 2, 4].$$

There is no further information in the graph at the moment, so we conclude by stating that D is isomorphic to a subgroup of $\text{Sym}(2, 2, 2, 2)$, which has order 16. This is the finest the ordered partition can get, as this is actually the orbits of the group $H_1 \cap H_2$.

Since D also is a subgroup of H_1 which has order 20, we deduce that $|D| \leq 4$.

In this case, we were able to deduce the exact orbits of $H_1 \cap H_2$. This is not true in general as our next example will show – but we will still perform useful deductions which Fixed is unable to perform.

Example 16. Again $H_1 := \langle (1, 2, \dots, 10), (2, 10)(3, 9)(4, 8)(5, 7) \rangle$, but this time $H_2 := G_{\{1,6\}}$ and again we wish to calculate $D := H_1 \cap H_2$. Using reasoning from the orbits of H_2 we obtain the ordered partition $Q_1 = [1, 6 \mid 2, 3, 4, 5, 7, 8, 9, 10]$. We note that $\text{Fixed}_{H_2}(Q) = Q$ because H_2 is transitive and Q_1 has no singleton cells.

We consider the same orbital graph as in the previous example.

Step 1:

We create a graph where we attach the colour 1 to the vertices 1 and 6, and colour 2 to all other vertices. This gives the ordered partition $[1, 6 \mid 2, 3, 4, 5, 7, 8, 9, 10]$.

Step 2:

We split the vertices in the second cell into vertices which have different coloured neighbours (vertices 2, 7, 5 and 10) and those with two neighbours of colour 2 (vertices 3, 4, 8 and 9).

Our second ordered partition is therefore $Q_2 = [1, 6 \mid 3, 4, 8, 9 \mid 2, 7, 5, 10]$.

Step 3:

The algorithm will now run through the remaining reasoning, not producing any further refinements, as the ordered partition is already equitable.

We can say that D is isomorphic to a subgroup of $\text{Sym}(2, 4, 4)$ of order $2^7 \cdot 3^2$. Since D also is a subgroup of H_1 which has order 20, we deduce that $|D| \leq 4$.

Putting together both of our examples, using orbital graphs we know that any elements of H_1 in $\text{Co}(P_1, Q_1)$ are also in $\text{Co}(P_2, Q_2)$, by Lemma 7. However, $\text{Co}(P_2, Q_2) = \emptyset$, because P_2 and Q_2 have different numbers of cells, and therefore we have deduced there are no members of H_1 in $\text{Co}(P_1, Q_1)$.

4. Efficiently Creating Orbital Graphs

While orbital graphs can be very useful in reducing search, they are expensive to create. Therefore we want to only create them when they provide useful extra refinements.

Definition 17. Suppose that $H \leq \text{Sym}(\Omega)$ and that P is an ordered orbit partition of H . Let $\alpha, \beta \in \Omega$ and let Γ denote the orbital graph for H with base-pair (α, β) .

The orbital graph Γ is **futile** if and only if G_P , in its natural action on Ω , induces graph automorphisms on Γ .

We capture in this definition that some orbital graphs do not enable additional refinement when applying the orbital graph refiner from Definition 12 anywhere in the search, other than the orbit structure of the group which is already provided by the Fixed refiner.

Theorem 3.2 from Hähndel et al. (2017) provides a complete characterisation of futile graphs – they are either a complete digraph on some subset of the vertices, or a complete bipartite digraph between two disjoint subsets of the vertex set.

We will now describe our algorithm for efficiently finding all the orbital graphs of a group G .

Our algorithm for finding orbital graphs is given in Algorithm 2. This algorithm assumes the use of a computational group theory system, such as GAP, provides basic group theory algorithms. In particular, computing point stabilizers, orbits of points, and pairs of points.

Algorithm 2 Find Orbital Graphs

```

1: procedure ORBITALBASE( $G, \Omega, SizeLimit$ )  $\triangleright$  Orbital graphs of  $G$ , a permutation group on  $\Omega$ 
2:   Graphs := []
3:   if  $G$  is  $k$ -transitive for  $k \geq 2$  then
4:     return Graphs
5:   for  $Orb \in \text{Orbits}(G)$  do
6:     if  $|Orb| > 1$  then
7:        $G' = \text{Stabilizer}(G, \text{Min}(Orb))$ 
8:       for  $InnerOrb \in \text{Orbits}(G')$  do
9:         if  $|Orb| \times |InnerOrb| \leq SizeLimit$  then
10:          if  $InnerOrb \notin \text{Orbits}(G)$  then
11:            if  $InnerOrb \in Orb$  and  $|InnerOrb| + 1 = |Orb|$  then
12:              Add(Graphs, Orbit( $G, (\text{Min}(Orb), \text{Min}(InnerOrb))$ ))
13:   return Graphs

```

We split the proof of Algorithm 2 into three pieces. First, we will prove some results about enumerating all orbital graphs, then we show how to identify futile orbital graphs, and finally we put these results together into a proof of the algorithm's correctness.

Theorem 18. *Algorithm 2 returns all orbital graphs, except those which are futile, or contain more than $SizeLimit$ edges.*

Proof. First, Line 3 performs an initial check if the group is k -transitive for $k \geq 2$. If it is, then we stop because, by Lemma 3.7 of Hähndel et al. (2017), all orbital graphs of G are futile in this case.

After this, Line 5 picks one member a from each orbit of G to be the first member of the ordered pair we will use as our base-pair. Then we apply Lemma 2.10 of Hähndel et al. (2017): Part (i) shows that it is enough to pick one element from each orbit as first member and that this will generate all orbital graphs, and Part (ii) shows that no orbital graph will be generated twice.

Lines 7 and 8 then pick one member from the stabilizer of each of our first base points. Part (iii) of the lemma mentioned above shows that this will give us a set of base-pairs which generate every orbital graph once.

We now move through the other lines, which skip orbital graphs we do not want to consider. Line 10 and Line 11 check the conditions of Lemmas 3.4 and 3.5 of Hähndel et al. (2017), rejecting all futile orbital graphs. Line 9 provides an extra check on the size of the orbital graph, allowing us to reject any orbital graph with a user-defined size limit.

Finally, Line 6 provides a fast early check. If some $a \in \Omega$ is already fixed in G , then $G = G_a$, which means that line 10 will always fail. Therefore we may as well reject such points immediately. \square

The runtime of Algorithm 2 is for most problems dominated by the calculation of the point stabilizers on Line 7. In GAP these are calculated using a randomized implementation of the Schreier-Sims algorithm. As our experiments will show, Algorithm 2 is highly efficient in practice.

5. Experiments

We will now demonstrate how our algorithm performs on a variety of problems. There are three main questions we want to answer:

1. How much can we reduce the size of search?
2. What are the worst-case slowdowns when using orbital graphs?
3. On which problems are orbital graphs effective?

Generating a set of “random” problems – taken from all possible group intersection and stabilizer problems – is not possible as there is no method of producing random permutation groups of large size. Also, we do not want to consider problems which are very simple – partition backtrack solves many problems almost instantly, with or without orbital graphs.

We consider two main categories of problem: Set stabilizers in groups we refer to as “grid groups”, which arise in A.I. in SAT solvers and Constraint Programming, and are used for symmetry-breaking, and taking intersections of primitive groups with wreath products.

We consider the following four algorithms:

1. **Fixed** : Fixed, the traditional algorithm of Leon.
2. **PreOrbital** : Use Fixed with Orb as described in Section 4.
3. **DeepOrbital** : Use Fixed with DeepOrb as described in Section 4.

4. **FirstOrbital**: Use Fixed and a variant of DeepOrb, which keeps regenerating orbital graphs until the first depth at which at least one non-trivial orbital graph is found. Then these orbital graphs are used at all further depths.

FirstOrbital is implemented as a variant **DeepOrbital**. **FirstOrbital** is difficult to cleanly specify theoretically, as it is tied to behaviour of the search. **FirstOrbital** is a valid refiner in practice, because down each branch of search we will find the first node with at least one non-trivial orbital graph at the same height – as the Fixed propagator ensures there is an element of the group which maps the fixed points down the first branch to the fixed points down every other branch at every level, else the search backtracks.

We expect that **Fixed** will always produce larger search trees than **PreOrbital**, which in turn produce larger search trees than **FirstOrbital**, which will produce larger search trees than **DeepOrbital**. In some rare cases the searches can be larger with a better refiner, as the partition backtrack algorithm may choose to branch on a different cell, or on the elements of a cell in a different order. In practice this occurred in less than 2% of our experiments.

The purpose of these experiments is to show when the decreased search size provided by **PreOrbital**, **DeepOrbital** and **FirstOrbital** outweighs the increased cost of calculating and filtering orbital graphs.

All of our experiments are performed in GAP (see GAP (2016)). We use the implementation of partition backtrack provided in the Ferret package (see Jefferson (2016)), which includes both an implementation of Leon’s original partition backtrack algorithms, and our new algorithms. In the experiments in this paper, Ferret’s implementations of Leon’s algorithms are always faster than the implementations in GAP, due to improved quality of implementation and the data structures used.

All of our experiments were performed on a machine with eight Intel Xeon E5520 cpus, running at 2.27GHz and 20GB RAM. Each experiment was given a five minute timeout and a limit of 1GB of RAM.

5.1. Set stabilizers in Grid Groups

In this experiment, we consider finding stabilizer of sets of integers in the **grid group**:

Definition 19. Let $m, n \in \mathbb{N}$. The direct product $\mathcal{S}_m \times \mathcal{S}_n$ acts on the set $\{1, \dots, m\} \times \{1, \dots, n\}$ of pairs in the following way:

For all $(i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}$ and all $(\sigma, \tau) \in \mathcal{S}_m \times \mathcal{S}_n$ we define

$$(i, j)^{(\sigma, \tau)} := (i^\sigma, j^\tau).$$

The subgroup $G \leq \text{Sym}(\{1, \dots, m\} \times \{1, \dots, n\})$ defined by this action is called the $m \times n$ **grid group**.

While the construction of the grid group is done by starting with an m by n grid of points and permuting rows or columns independently of each other, we represent this group as a subgroup of $\mathcal{S}_{m \cdot n}$, and we do not assume prior knowledge of the grid structure of the action. We considered two different variants of set stabilizers: stabilizing a random subset of $\{1, \dots, m^2\}$ of size $\lfloor \frac{m^2}{2} \rfloor$, and stabilizing random subsets of $\{1, \dots, m^2\}$ containing exactly $\lfloor \frac{m}{2} \rfloor$ points in each row.

Firstly, we generate a random set of size $\lfloor \frac{m^2}{2} \rfloor$ of $\{1 \dots m^2\}$. The results are shown in Figure 1. We show the average of the number of nodes taken from instances which finished (we show N/A where no instance was completed). For times, we use the timeout time (5 minutes) for instances

Grid Size	Fixed		PreOrbital		DeepOrbital	
	#	Time	#	Time	#	Time
10	10	0.1	10	0.1	10	0.5
20	10	2.0	10	0.5	10	10.8
30	10	15.6	10	2.9	10	79.4
40	10	108.6	10	10.5	10	262.9
50	7	1,506.5	10	40.2	10	643.1
60	3	2,554.5	10	101.9	6	2,253.7
70	0	3,000	9	490.6	2	2,862.1
80	0	3,000	10	497.8	4	1,951.4
90	1	2,772.3	10	1,052.7	0	3,000
100	0	3,000	9	1,195.2	4	2,009.1
110	0	3,000	9	1,554.3	5	1,869.7
120	0	3,000	0	3,000	0	3,000

Figure 1: Time taken to find set stabilizers in square grid graphs (average of 10 runs). “#” is the number of problems solved. “Time” is the total time in seconds to complete (or timeout) the 10 instances.

Grid Size	Fixed		PreOrbital		DeepOrbital	
	#	Time	#	Time	#	Time
10	10	0.1	10	0.1	10	0.3
20	4	1,892.8	10	0.4	10	0.9
30	3	2,107.3	10	2.1	10	5.1
40	0	3,000	10	8.1	10	8.2
50	0	3,000	10	30.3	10	30.3
60	0	3,000	10	87.3	10	84.5
70	0	3,000	10	235.0	10	236.2
80	0	3,000	10	542.0	9	802.2
90	0	3,000	10	1,327.4	9	1,393.1
100	0	3,000	3	2,589.7	4	2,553.4
110	0	3,000	2	2,919.0	1	2,941.8
120	0	3,000	1	2,880.9	1	2,890.2

Figure 2: Time taken to find row-balanced set stabilizers in square grid graphs (average of 10 runs). “#” is the number of problems solved. “Time” is the total time in seconds to complete (or timeout) the 10 instances.

Problem	Refiner	Building Stab Chains	Building Orbital Graphs	Refining Orbital Graphs
Set Stab	Fixed	39.4	0	0
	PreOrbital	95.2	1.1	0.1
	DeepOrbital	25.5	58.5	0.4
Row Balanced Set Stab	Fixed	1.3	0	0
	PreOrbital	97.8	0.9	0.1
	DeepOrbital	97.5	1.0	0.1

Figure 3: Percentage of runtime spent on stabilizer chains and orbital graphs in Grid set stabilizer experiments.

which ran out of either time or memory. We do not show results for **FirstOrbital**, as they are identical to **PreOrbital**.

As the problems are randomly generated some are simpler than others, which is why we see that **Fixed** is able to solve one problem on a grid of size 90. However, for all sizes of grids we see a substantial improvement from building orbital graphs. As **DeepOrbital** keeps regenerating the graphs whenever a new point is fixed in the ordered partition, the time taken is generally longer, as the extra graphs do not pay back their cost for most of these problems.

In the second experiment, we generate row-balanced sets. These include exactly $\lfloor \frac{m}{2} \rfloor$ points in each for row of the grid. We intuitively expect these instances to be more difficult, as it is more difficult to show that that the stabilizer does not permute the rows of the grid.

The results of this experiment are shown in Figure 2. The major difference here is that no instances of size bigger than 30 were solved by **Fixed**, while **PreOrbital** and **DeepOrbital** solved almost all instances. **FirstOrbital** once again had identical results to **PreOrbital**, so we skip this step.

In both experiments, we also recorded the amount of time spent performing different parts of the search. These results are shown in Table 3. We see that **PreOrbital** spends around 1% of the total time both building and refining orbital graphs. While a much higher proportion of time is spent building stabilizer chains, these are chains are a subset of the chains which had to be built by **Fixed**. Further, we investigated the size of the searches produced, and found that in almost every problem **DeepOrbital** and **PreOrbital** were able to find the stabilizer (which was the trivial group) without performing any branching. Therefore, the building of these stabilizer chains is now the limiting factor to any further improvement.

We see that **DeepOrbital** behaves very strangely, spending much more time building graphs with non row-balanced sets than with row-balanced sets. These results occur because during refinement the non row-balanced problems take more refinement steps to reach a fixed ordered partition, which results in more graphs being created. This shows that the performance of **DeepOrbital** can be very unpredictable.

Most instances which do not finish run out of time. Of the 240 experiments, for **Fixed**, 68 instances finish, 151 timeout and 21 run out of memory. For **PreOrbital**, 203 instances finish, 37 run out of time and no instance runs out of memory. For **DeepOrbital**, 165 instances finish, 51 run out of time and 24 run out of memory.

The main observation of this experiment is that the overhead of **PreOrbital** is very small while there is an exponential decrease in search size. While **DeepOrbital** did not take much more time, it did take much more memory. Further, its behaviour is unpredictable. The behaviour

of **FirstOrbital** is identical to **PreOrbital** on this problem. The limiting factor is now the time taken finding stabilizer chains, which is outside the scope of this article.

5.2. Intersection of Primitive Groups

In our second experiment, we consider intersection of groups. We consider intersecting primitive groups, as the primitive group library in GAP provides easy access to a large set of groups.

First we tried intersecting pairs of primitive groups, but found this problem is usually very simple, with and without orbital graphs. Therefore we intersect a primitive group with a wreath product of symmetric groups, as we found this produced challenging problems.

For each $n \in \mathbb{N}$ we take the primitive groups except the symmetric groups S_n , the alternating groups A_n , the cyclic groups C_n and the dihedral groups D_n in their natural action on n points, where these groups are primitive, as the intersections involving these groups are simple to construct.

As 2-transitive groups have futile orbital graphs we consider them separately.

Given a primitive group G acting on n points, we create the wreath products $S_{n/x} \wr S_x$, where $x \in \{2, \dots, 7\}$, in the case where n/x is an integer, then we conjugate each of these wreath products by a random permutation and intersect the result with G . We experimented with other wreath products and found similar results.

We consider the primitive groups on up to 600 points. This produces a total of 2,752 experiments on primitive groups which are not 2-transitive and 1,140 experiments for primitive groups which are 2-transitive. We ran each experiment for a maximum of five minutes.

Note that when considering 2-transitive groups, we still build an orbital graph for the wreath product and these orbital graphs can be very large. One of the orbital graphs of the wreath product of $S_a \wr S_b$ consists of a cliques of size b , and the second all the other edges of the graph.

Looking firstly at groups which are primitive but not 2-transitive, **Fixed** solved 1,794 problems within the timeout, **PreOrbital** solved 2,410, **DeepOrbital** solved 2,377 and **FirstOrbital** solved 2,411. We show the cumulative time taken to solve (or timeout) all problems in Figure 5. Our result show that all techniques involving orbital graphs solve substantially more problems within the timeout, and the problems solved were solved much faster with orbital graphs. **DeepOrbital** is slightly slower but not significantly so, and **PreOrbital** and **FirstOrbital** are almost identical, which we would expect.

For 2-transitive groups **Fixed** solved 992 problems, **PreOrbital** solved 982, **DeepOrbital** solved 1,083 and **FirstOrbital** solved 1,071. We show the cumulative time taken to solve (or timeout) all problems in Figure 6. Here we find orbital graphs are less useful, but still allow us to solve more problems within the timeout, and to solve problems within the timeout faster. **PreOrbital** does not pay off the cost of building orbital graphs, taking significantly longer. **DeepOrbital** solves more problems, and solves them faster. **FirstOrbital** solves the most problems, taking only slightly longer than **DeepOrbital**.

Figure 4 shows how long was spent in different parts of the algorithm during the search. As with the grid experiments, the majority of time was spent finding stabilizer chains. The **PreOrbital** and **FirstOrbital** algorithms both spend very little time building and refining with orbital graphs. The search sizes are significantly reduced, for groups of size over 500 **Fixed** averages around 3.7 million nodes on problems it can solve within the timeout, while **FirstOrbital** averages 21,000 nodes on the same set of problems. Similar to the earlier grid experiments, the main limiting factor in further performance improvements is the building of stabilizer chains.

Problem	Refiner	Build Stab Chains	Build Orbital Graphs	Refine Orbital Graphs
Not 2-trans	Fixed	99.7	0	0
	PreOrbital	97.5	0.5	0.1
	DeepOrbital	77.5	7.2	10.2
	FirstOrbital	97.5	0.5	0.1
2-trans	Fixed	80.8	0	0
	PreOrbital	78.3	0.1	4.7
	DeepOrbital	86.7	3.6	6.1
	FirstOrbital	88.9	0.6	6.6

Figure 4: Percentage of runtime spent building stabilizer chains and orbital graphs in Primitive Intersection experiments.

5.3. Conclusions

Our experiments show that using orbital graphs in partition backtrack can lead to substantial performance improvements – in the case of grid groups we can easily solve much larger problems than before, and become limited only by how quickly we can calculate stabilizer chains. While **DeepOrbital** is still a huge improvement over **Fixed** alone, the overhead of calculating orbital graphs for many groups is not recovered. Further improvements to the performance of set stabilizers in larger grids would require either new methods of finding stabilizer chains, or fundamental changes to partition backtrack to remove the need to calculate so many stabilizer chains.

For primitive groups, the results are more mixed. Here we can see that the cost of orbital graphs is larger – wreath products of symmetry groups in particular produce large orbital graphs. For non-2 transitive groups, all our orbital graph methods perform similarly. For 2-transitive groups, we see **PreOrbital**, as expected, performs slightly slower than **Fixed**, but **FirstOrbital** and **DeepOrbital** perform well. There are a small set of problems which only **DeepOrbital** is able to solve within the time limit.

The most important result to take away from our experiments is that **FirstOrbital** is a balanced algorithm with good practical behaviour. In all our experiments it has a low overhead, and performs close to the performance of whichever of the other algorithms is best. This suggests it would be a good option to always enable, as it is always much better than **Fixed**, and close to the best method for all problem classes we considered.

6. Final comments and future work

Our experiments show that orbital graphs can be very useful for refining ordered partitions. On a large range of problems they provide significant performance improvements, and the worst-case slowdowns they cause are not significantly harmful. In particular, **FirstOrbital** provides a simple and cheap method of significantly improving the state of the art. However, we cannot theoretically predict in advance when orbital graphs will be useful, and exactly how useful they will be: How can we measure the usefulness of a graph for refinements? Can we always find a graph that is useful in terms of refining ordered partitions?

The results in Section 4 suggest that it is worth investigating alternative refiners in more detail. For example, if it is possible to feed theoretical information about other structural in-

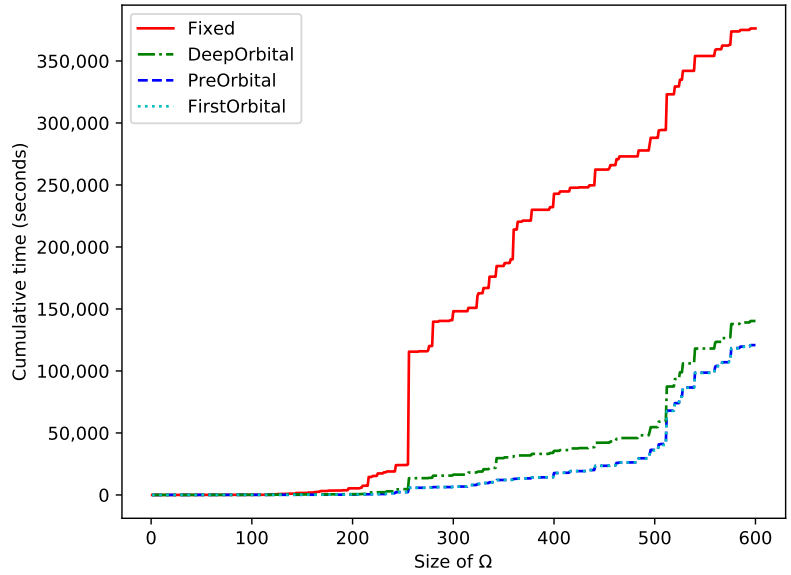


Figure 5: Cumulative time taken to intersect primitive and not 2-transitive primitive groups with a wreath product




Figure 6: Cumulative time taken to intersect 2-transitive groups with a wreath product

formation of groups into a refiner, then we may be able to improve the runtime of the search algorithm even more.

In cases where using Orb and DeepOrb still leads to long runtimes in practice, we intend to investigate other types of graphs and combinatorial structures that can be used for refinements in future work. Also we plan to improve other parts of the partition backtrack framework, as for many problems we are reaching the limit of improvements which better refiners alone can provide. In particular, to further improve the problems in our experiments it appears we must either speed up the calculation of stabilizer chains, or reduce the number of stabilizer chains which must be found.

References

- Bang-Jensen, J., Gutin, G. Z., 2008. Digraphs: Theory, Algorithms and Applications, 2nd Edition. Springer Publishing Company, Incorporated.
- Bosma, W., Cannon, J., Playoust, C., 1997. The Magma algebra system. I. The user language. *J. Symbolic Comput.* 24 (3-4), 235–265, computational algebra and number theory (London, 1993).
URL <http://dx.doi.org/10.1006/jsc.1996.0125>
- Butler, G., 1983. Computing normalizers in permutation groups. *J. Algorithms* 4 (2), 163–175.
URL [http://dx.doi.org/10.1016/0196-6774\(83\)90043-3](http://dx.doi.org/10.1016/0196-6774(83)90043-3)
- Dixon, J., Mortimer, B., 1996. *Permutation Groups*. Graduate Texts in Mathematics. Springer New York.
- GAP, 2016. GAP – Groups, Algorithms, and Programming, Version 4.8.4. The GAP Group.
URL <http://www.gap-system.org>
- Hähndel, P., Jefferson, C., Pfeiffer, M., Waldecker, R., 2017. Computational aspects of orbital graphs. submitted.
- Jefferson, C., 2016. The Ferret Package.
URL <https://github.com/gap-packages/ferret>
- Leon, J. S., 1991. Permutation group algorithms based on partitions, I: Theory and algorithms. *Journal of Symbolic Computation* 12 (4), 533–583.
URL <http://www.sciencedirect.com/science/article/pii/S0747717108801034>
- McKay, B. D., 1980. Practical graph isomorphism. *Congr. Numer.* (30), 45–87.
- McKay, B. D., Piperno, A., 2014. Practical graph isomorphism, II. *Journal of Symbolic Computation* 60, 94–112.
URL <http://www.sciencedirect.com/science/article/pii/S0747717113001193>
- Seress, Á., 2003. *Permutation Group Algorithms*. Cambridge Tracts in Mathematics. Cambridge University Press.
URL https://books.google.co.uk/books?id=hxFqdbfc_CMC
- Theißen, H., 1995. Eine Methode zur Normalisatorberechnung in Permutationsgruppen mit Anwendungen in der Konstruktion primitiver Gruppen. Ph.D. thesis, Lehrstuhl D für Mathematik, RWTH Aachen.