



ELSEVIER

Contents lists available at ScienceDirect

Journal of Symbolic Computation

www.elsevier.com/locate/jsc



New refiners for permutation group search

Christopher Jefferson^a, Markus Pfeiffer^a, Rebecca Waldecker^b

^a University of St Andrews, School of Computer Science, North Haugh, St Andrews, KY16 9SX, Scotland, United Kingdom

^b Martin-Luther-Universität Halle-Wittenberg, Institut für Mathematik, 06099 Halle, Germany

ARTICLE INFO

Article history:

Received 26 May 2017

Accepted 11 December 2017

Available online 11 January 2018

Keywords:

Backtrack search

Refiners

Permutation groups

Algorithmic group theory

Computational algebra

Partition backtrack

ABSTRACT

Partition backtrack is the current generic state of the art algorithm to search for subgroups of a given permutation group. We describe an improvement of partition backtrack for set stabilizers and intersections of subgroups by using orbital graphs. With extensive experiments we demonstrate that our methods improve performance of partition backtrack – in some cases by several orders of magnitude.

© 2018 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Permutation groups are one of the most natural and convenient representations of finite groups. They have proved particularly useful for computational purposes, and systems such as GAP (2016) and Magma (Bosma et al., 1997) provide efficient implementations of many algorithms to solve a range of problems from membership testing to identification of the isomorphism type of a group.

Given a permutation group acting on a finite set Ω , some problems – for example checking whether or not a group contains a particular permutation, or computing the size of the group – can be solved in time polynomial in the size of Ω . There are problems for which no polynomial time algorithm is known, for example computing intersections, centralizers and normalizers of subgroups, and set and partition stabilizers. For these problems the best known algorithms perform a very sophisticated exhaustive search through the group in question. It is known (see for example Chapter 3

E-mail addresses: caj21@st-andrews.ac.uk (C. Jefferson), markus.pfeiffer@st-andrews.ac.uk (M. Pfeiffer), rebecca.waldecker@mathematik.uni-halle.de (R. Waldecker).

URLs: <http://caj.host.cs.st-andrews.ac.uk/> (C. Jefferson), <https://www.morphism.de/~markusp/> (M. Pfeiffer), <http://conway1.mathematik.uni-halle.de/~waldecker/index-english.html> (R. Waldecker).

<https://doi.org/10.1016/j.jsc.2017.12.003>

0747-7171/© 2018 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

in Seress, 2003) that these problems are at least as difficult as graph isomorphism, so it is unlikely that a polynomial time algorithm can be found.

The current state of the art algorithm is described in Leon (1991) and is commonly called partition backtrack. It extends ideas introduced in McKay (1980) to solve graph isomorphism problems. Implementations of partition backtrack are available in GAP (2016) and Magma (Bosma et al., 1997), and they solve the problems mentioned above very efficiently for a fair range of examples.

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 a range of examples, as will be demonstrated in Section 6. Partition backtrack should not be viewed as a monolithic algorithm, but rather as a combination of algorithms solving particular sub-problems. The concept of a **refiner** is one of its crucial components. They are used to detect and skip parts of the computation that would be superfluous: Better refiners lead to more superfluous computational steps to be skipped.

We describe a new class of refiners using orbital graphs. It can be used to improve performance of partition backtrack implementations, and we demonstrate the speed-up that can be achieved in our implementation of partition backtrack. This article is organized as follows:

Section 2 includes necessary notation and examples of permutation groups and ordered partitions. In Section 3 we give a brief description of backtrack search and the role of refiners, and we discuss some standard refiners. Section 4 introduces orbital graphs and gives a characterization of the cases where orbital graphs can benefit refinement, we briefly discuss the limits of refiners using orbital graphs: For example, for 2-transitive groups, orbital graphs do not provide any improvement. Section 5 then combines ideas introduced in Sections 3 and 4 to define new refiners using orbital graphs. Finally, we explain and present experiments and their outcomes in Section 6. In particular, there are cases in which the costs of calculating orbital graphs outweigh the advantages in the search, but for some search problems where the previously best techniques perform poorly, our methods prove to be very effective.

At the end of the paper we comment on related questions and further research.

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 $\{1, \dots, n\}$.

Definition 1 (Ordered partitions).

- An **ordered partition** of Ω is an ordered list of non-empty disjoint subsets (called **cells**) of Ω whose union is all of Ω . For all $i \in \Omega$ we write $\Delta_P(i)$ for the cell of P that contains the point i , and we let $\text{OPart}(\Omega)$ denote the set of **all ordered partitions of** Ω . The notation for ordered partitions will be explained in Example 2. If $P \in \text{OPart}(\Omega)$ and $i, j \in \Omega$, then we write $i \sim_P j$ if and only if $i \in \Delta_P(j)$.
- We say that Q is **finer than** P and write $Q \preceq P$ if and only if, for all $i, j \in \Omega$, it is true that $i \sim_Q j$ implies $i \sim_P j$. Conversely, we say that P is **coarser** than Q in this situation.

The relation \preceq defines a pre-order on $\text{OPart}(\Omega)$ and we point out that, for all $P \in \text{OPart}(\Omega)$, it is true 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 and only if Ω itself is its only cell.

- If P is an ordered partition of Ω and $g \in \text{Sym}(\Omega)$, then we write P^g for the ordered partition that we obtain by applying g to the elements in the cells of P .
- 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 $P \in \text{OPart}(\Omega)$, 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 . We point out 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 and Q are ordered orbit partitions of $\{1, \dots, n\}$, that $k \in \mathbb{N}$ and that $\Delta_1, \dots, \Delta_k$ are exactly the cells of P . 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 empty or a coset of $\text{Sym}(P)$. Moreover $\text{Co}(P, P) = \text{Sym}(P)$.

Example 2. Given the algorithmic background of our work, we view partitions as lists. For example $P := [1, 2, 3, 4 \mid 5, 6, 7]$ is an ordered partition of $\Omega := \{1, 2, 3, 4, 5, 6, 7\}$ with cells $\{1, 2, 3, 4\}$ and $\{5, 6, 7\}$. Then $P = [2, 1, 3, 4 \mid 5, 6, 7]$, because the ordering of elements within a cell is irrelevant. But $P \neq [5, 6, 7 \mid 1, 2, 3, 4]$, because the ordering of cells is relevant. However, we see that $P \preceq [5, 6, 7 \mid 1, 2, 3, 4]$. Let $Q := [1, 2, 3 \mid 4 \mid 5, 6, 7]$. Then $Q \preceq P$ and $Q \preceq [4, 5, 6, 7 \mid 3, 2, 1]$.

Next we consider the stabilizer in \mathcal{S}_7 of 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: $[1, 2, 3 \mid 4 \mid 5, 7 \mid 6]$ and $[4 \mid 6 \mid 5, 7 \mid 1, 2, 3]$.

Definition 3 (Meet). Let $P, Q \in \text{OPart}(\Omega)$. Then we define the **meet of P and Q** , denoted by $P \wedge Q$, as follows: $i, j \in \Omega$ are in the same cell of $P \wedge Q$ if and only if $i \sim_P j$ and $i \sim_Q j$. 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.

Remark 4. There are several reasons why ordered partitions are used in partition backtrack. One reason is that we can represent all elements of the group as coset representatives of some coset of the stabilizer of an ordered partition. This approach does not work with unordered partitions. Another reason is that, with ordered partitions, the relation “meet” is compatible with taking stabilizers of ordered partitions. This property no longer holds if unordered partitions are used. One final comment: The relation “meet” is not symmetric, as is illustrated in the following example.

Example 5. Let $P := [1, 2, 3, 4 \mid 5, 6, 7]$ and $Q := [1, 2 \mid 5, 3 \mid 7, 4, 6]$ be ordered partitions of $\Omega := \{1, 2, 3, 4, 5, 6, 7\}$. We calculate $P \wedge Q$: For each element $i \in \Omega$, we find the indices of the cells $\Delta_P(i)$ and $\Delta_Q(i)$, in this order. This gives the following pairs:

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

Therefore $P \wedge Q = [1, 2 \mid 3 \mid 4 \mid 5 \mid 6, 7]$. Calculating $Q \wedge P$ instead gives the ordered partition $[1, 2 \mid 3 \mid 5 \mid 4 \mid 6, 7]$; it has the same cells as $P \wedge Q$, but in a different order.

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 6. Let $H \leq \text{Sym}(\Omega)$, let P and Q be ordered partitions of Ω , and let $h \in H$. Then $(P \wedge Q)^h = P^h \wedge Q^h$.

Proof. First we note: If $g \in \text{Sym}(\Omega)$, then $i \sim_{pg} j$ if and only if $i^{g^{-1}} \sim_p j^{g^{-1}}$. Let $T := P \wedge Q$. Then by definition $i \sim_T j$ if and only if $i \sim_P j$ and $i \sim_Q j$. Now, for all $h \in H$:

$$i \sim_{T^h} j \Leftrightarrow i^{h^{-1}} \sim_T j^{h^{-1}} \Leftrightarrow i^{h^{-1}} \sim_P j^{h^{-1}} \text{ and } i^{h^{-1}} \sim_Q j^{h^{-1}} \Leftrightarrow i \sim_{ph} j \text{ and } i \sim_{qh} j \Leftrightarrow i \sim_{ph \wedge qh} j.$$

So far we proved that $(P \wedge Q)^h \preceq P^h \wedge Q^h$ and $P^h \wedge Q^h \preceq (P \wedge Q)^h$, but this does not imply equality. We can describe the cell $\Delta_{P \wedge Q}(i)$ by the pair $[\Delta_P(i), \Delta_Q(i)]$. So we argue as follows:

$$\Delta_{(P \wedge Q)^h}(i) = \Delta_{(P \wedge Q)}(i^{h^{-1}}) = [\Delta_P(i^{h^{-1}}), \Delta_Q(i^{h^{-1}})] = [\Delta_{ph}(i), \Delta_{qh}(i)] = \Delta_{ph \wedge qh}(i).$$

This implies that $(P \wedge Q)^h = P^h \wedge Q^h$. \square

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 $\text{Sym}(\Omega)$ ”. For example, finding the stabilizer of a set S in a subgroup H of $\text{Sym}(\Omega)$ can be expressed as finding all elements that satisfy the list of properties “element stabilizes S in $\text{Sym}(\Omega)$ ” and “element is contained in the subgroup H ”. We will not specify “properties” further because we have many different applications in mind, but each property should be easy to verify with an algorithm. The groups that appear in the search 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 $\text{Sym}(\Omega)$.

To give a more precise description: 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 that 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 $\text{Sym}(\Omega)$ that maps P_{i+1} to Q_{i+1} . This means that this part of the search does not produce any permutation that 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 that 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 that 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} , we create Q'_j , which is identical to Q_{i+1} except that 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 continues by going to the refinement phase.

This process will stop eventually because there is a finite number of branches, they all have finite length, and the refinement stops if the partition does not change anymore. In [Leon \(1991\)](#) the author explains how this algorithm can be implemented efficiently.

This paper will focus on refinement, because the quality of refiners is one of the main influences on performance. Refiners are only defined on a single ordered partition – [Lemma 8](#) explains how refiners are used on cosets.

Definition 7 (*M-refiner*). Let $M \subseteq \text{Sym}(\Omega)$ be a set of elements with a given property. An *M-refiner* is a map $f_M : \text{OPart}(\Omega) \rightarrow \text{OPart}(\Omega)$ that satisfies the following conditions for any ordered partition P of Ω :

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

Part (b) of the definition implies that every element in $\text{Sym}(P)$ that satisfies the property (referring to M) and stabilizes P also stabilizes $f(P)$. This turns out to be a very natural condition – our experience is that refiners tend to satisfy (b). 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 8](#) 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 8. Let $P, Q \in \text{OPart}(\Omega)$ and let f_M be an *M-refiner* for $M \subseteq \text{Sym}(\Omega)$. 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 $Q = P^g$, and by definition of an *M-refiner* it holds 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

3.2. Refiners for permutation groups

We will now give a definition of the standard refiner for permutation groups given by a list of generators from [Leon \(1991\)](#). We will introduce refiners that use orbital graphs in Section 4.

Definition 9 (*Fixed_M*). Let M be a subgroup of $\text{Sym}(\Omega)$. Then 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 . Let M_0 denote the point-wise stabilizer of $\alpha_1, \dots, \alpha_k$ in M .
- Return the meet of P and an ordered orbit partition of M_0 .

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 resulting ordered orbit partition. In the implementation this is fixed by producing a list of orbits, outputting the cells in arbitrary order and then fixing this ordering for later instances. The details are given in the proof of the next lemma.

Lemma 10. *Let M be a subgroup of $\text{Sym}(\Omega)$. Then the map Fixed_M is a refiner.*

Proof. Let $P \in \text{OPart}(\Omega)$. 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 denote the point-wise stabilizer of $\alpha_1, \dots, \alpha_k$ in M and we note that $F := \langle M_0 \rangle$ stabilizes $\alpha_1, \dots, \alpha_k$. 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 some ordered orbit partition Q of F as described before the lemma. Then Q^g is an ordered orbit partition of F^g . Using Lemma 6 we deduce:

$$\text{Fixed}_M(P)^g = (P \wedge Q)^g = P^g \wedge Q^g = \text{Fixed}_M(P^g). \quad \square$$

The major limitation of Fixed is that it ignores non-singleton cells. More concretely, given a transitive group G and an ordered partition P that 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 that 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 that complement existing refiners for groups expressed by a list of generators. These refiners will provide useful refinement even for transitive groups and non-singleton cells. We are not the first ones to use this idea: Theißen (1995) uses orbital graphs as an ingredient for refiners for normalizer search. However, our work does not build on his – partly because our hypothesis is more general, and partly because his results have not been published except for in his PhD thesis. We show that the concept has not yet been fully exploited.

4. Orbital graphs

Here we introduce the graphs that we use in our new refiners – orbital graphs – and prove the properties that are necessary in order to decide whether or not a refinement by orbital graphs is computationally beneficial.

Definition 11 (Digraphs). 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) . An **isolated vertex** of a digraph is a vertex with no arcs going into it or coming out of it. A digraph is **complete** if and only if its set of arcs is exactly $\{(x, y) \mid x, y \in V, x \neq y\}$.

Γ is a **complete bipartite digraph** if and only if there exist disjoint subsets S, E of vertices such that V is the union of S (the “starting” vertices) and E (the “end” vertices) and the set of arcs is exactly $A = \{(x, y) \mid x \in S, y \in E\}$.

We refer the reader to Bang-Jensen and Gutin (2008) for standard notation and for the definitions of connected components, graph isomorphisms etc. We point out that by a **proper digraph** we mean a digraph that has at least one arc such that its reverse arc is not in the graph. All digraphs considered here have no multiple arcs and no loops. We also point out that whenever we refer to the **size of a connected component** we mean the number of vertices in the component.

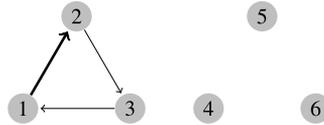
Definition 12 (Orbital graphs). Let Ω be a finite set, $G := \text{Sym}(\Omega)$ and $H \leq G$. For all vertices $\gamma \in \Gamma$ and all $h \in H$ we write γ^h for the image of γ under h in the original permutation action.

Now let $\alpha, \beta \in \Omega$ be distinct elements, chosen in this order. We define a digraph $\Gamma = (\Omega, A)$ where the set of arcs A is defined as $A := \{(\alpha^h, \beta^h) \mid h \in H\}$. This digraph is called the **orbital graph of H with base-pair** (α, β) , and is denoted by $\Gamma(H, \Omega, (\alpha, \beta))$.

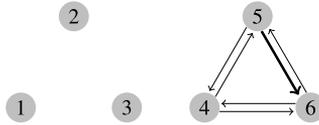
Following [Dixon and Mortimer \(1996\)](#) we say that an orbital graph is **self-paired** if and only if, for all $\gamma, \delta \in \Omega$, it is true that (γ, δ) is an arc if and only if (δ, γ) is an arc.

Example 13. If we build an orbital graph for \mathcal{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 $\alpha, \beta \in \{1, 2, 3\}$ such that (α, β) is a base-pair, and given any distinct $\gamma, \delta \in \{1, 2, 3\}$, there exists some $g \in \mathcal{S}_3$ such that $\alpha^g = \gamma$ and $\beta^g = \delta$. Therefore (γ, δ) 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 $H := \langle (123), (45), (46) \rangle \leq \mathcal{S}_6$.

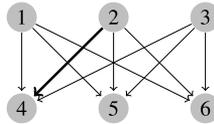
Then, starting with the base-pair $(1, 2)$, we obtain the following digraph:



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



For a connected digraph we start with the pair $(2, 4)$:



Some properties of orbital graphs can be found in [Cameron \(1999\)](#) and [Dixon and Mortimer \(1996\)](#), but we decided to include short proofs for the statements in the next lemma in order to make this article more self-contained.

Hypothesis 14. Let Ω be a finite set, let $H \leq G := \text{Sym}(\Omega)$ and let $\alpha, \beta \in \Omega$ be distinct. Let $\Gamma := \Gamma(H, \Omega, (\alpha, \beta))$ and let A denote the set of arcs of Γ .

Lemma 15. Suppose that [Hypothesis 14](#) holds. Then we have the following:

- (i) $\Gamma = \Gamma(H, \Omega, (\gamma, \delta))$ if and only if $(\gamma, \delta) \in A$.
- (ii) Γ is self-paired if and only if some $h \in H$ interchanges α and β .
- (iii) α^H is precisely the set of vertices of Γ that are the starting point of some arc.
- (iv) β^H is precisely the set of vertices of Γ that are the end point of some arc.
- (v) The number of arcs starting at α is $|\beta^{H\alpha}|$ and the number of arcs going into β is $|\alpha^{H\beta}|$.

Proof. (i) If $\Gamma = \Gamma(H, \Omega, (\gamma, \delta))$, then by definition (γ, δ) is an arc in Γ .

Conversely, suppose that (γ, δ) is an arc in Γ . Then there exists some $h \in H$ such that $(\alpha^h, \beta^h) = (\gamma, \delta)$. Hence the orbital graph with base-pair (α, β) is the same as the orbital graph with base-pair (γ, δ) .

(ii) By (i) Γ coincides with $\Gamma(H, \Omega, (\beta, \alpha))$ if and only if the arc (β, α) exists in Γ , which happens if and only if there exists some $h \in H$ such that $\alpha^h = \beta$ and $\beta^h = \alpha$.

(iii) Let $\gamma \in \alpha^H$ and $h \in H$ be such that $\gamma = \alpha^h$. Then (γ, β^h) is an arc with starting point γ . Conversely, if $\delta \in \Omega$ is such that (γ, δ) is an arc in Γ , then there exists some $h \in H$ such that $(\alpha^h, \beta^h) = (\gamma, \delta)$ and hence $\gamma = \alpha^h \in \alpha^H$. Similar arguments show (iv).

(v) The number of arcs starting at α is

$$|\{(\alpha, \gamma) \mid \gamma \in \Omega\}| = |\{(\alpha^h, \beta^h) \mid h \in H, \alpha^h = \alpha\}| = |\beta^{H\alpha}| \text{ and the number of arcs going into } \beta \text{ is}$$

$$|\{(\delta, \beta) \mid \delta \in \Omega\}| = |\{(\alpha^h, \beta^h) \mid h \in H, \beta^h = \beta\}| = |\alpha^{H\beta}|. \quad \square$$

Remark 16. Some comments:

(a) Parts (iii) and (v) of the lemma, together, give the total number of arcs in Γ . The number of arcs starting at α is exactly $|\beta^{H\alpha}|$, so we obtain $|A| = |\alpha^H| \cdot |\beta^{H\alpha}|$.

(b) In (ii) it is not true that H must contain the transposition (α, β) . A counterexample is provided by $H := \langle (12)(34) \rangle \leq S_4$ acting naturally on $\Omega := \{1, 2, 3, 4\}$ and its orbital graph with base-pair $(1, 2)$.

(c) Parts (iii) and (iv) of the lemma imply that, if H acts transitively on Ω , then Γ has no isolated vertices.

Lemma 17. Suppose that Hypothesis 14 holds. Then H acts on Γ as a group of graph automorphisms.

Proof. First we note that H acts faithfully on the set Ω . Now we let $\gamma, \delta \in \Omega$.

If (γ, δ) is an arc, then there exists some $h \in H$ such that $(\gamma, \delta) = (\alpha^h, \beta^h)$ by definition of Γ . Hence $(\gamma^g, \delta^g) = (\alpha^{hg}, \beta^{hg})$ is an arc. Conversely, if (γ^h, δ^h) is an arc, then there exists some $a \in H$ such that $(\gamma^h, \delta^h) = (\alpha^a, \beta^a)$ and hence $(\gamma, \delta) = (\alpha^{ah^{-1}}, \beta^{ah^{-1}})$ is an arc. As H is a group, the induced maps are bijective and hence every $h \in H$ induces a graph automorphism on Γ . \square

Lemma 18. Suppose that Hypothesis 14 holds and let Δ denote the connected component that contains (α, β) . Then every connected component of Γ that has size at least 2 is isomorphic to Δ .

Proof. Let Δ' denote an arbitrary connected component of Γ of size at least 2 and let (γ, δ) be an arc in Δ' .

From the definition of orbital graphs let $h \in H$ be such that $(\alpha^h, \beta^h) = (\gamma, \delta)$. Then h induces an automorphism on Γ by Lemma 17 and it moves all arcs from Δ to arcs in Δ' . Conversely, h^{-1} induces an automorphism on Γ that moves all arcs of Δ' into Δ . Thus it follows that Δ and Δ' are isomorphic as graphs. \square

Lemma 19 shows how to choose a set of base-pairs that determines all orbital graphs for a group H . Parts (i) and (ii) show to take a representative from each orbit of H as the first element of the base-pair, and then part (iii) shows that we must stabilize this representative in H , and take a representative from each orbit in this stabilizer for the second element of our base-pair. These base-pairs will allow us to analyze the set of orbital graphs of a group, before we construct any orbital graphs explicitly.

Lemma 19. Let Ω be a finite set and let $H \leq G := \text{Sym}(\Omega)$.

- (i) Suppose that $\alpha, \beta \in \Omega$ and $\alpha \in \beta^H$. Then the set of orbital graphs of H with base-pairs starting with α is equal to the set of orbital graphs of H with base-pairs starting with β .
- (ii) Suppose that $\alpha, \gamma \in \Omega$ and $\alpha \notin \gamma^H$. Then the set of orbital graphs of H with base-pairs starting with α is disjoint from the set of orbital graphs of H with base-pairs starting with γ .
- (iii) Suppose that $\alpha, \beta, \gamma \in \Omega$ and that $\alpha \neq \beta, \alpha \neq \gamma$. Let $\Gamma_1 := \Gamma(H, \Omega, (\alpha, \beta))$ and $\Gamma_2 := \Gamma(H, \Omega, (\alpha, \gamma))$. Then $\Gamma_1 = \Gamma_2$ if and only if $\gamma \in \beta^{H\alpha}$.

Proof. (i) Let $h \in H$ be such that $\alpha^h = \beta$. Then for all $\gamma \in \Omega$, it follows that $\Gamma(H, \Omega, (\alpha, \gamma)) = \Gamma(H, \Omega, (\alpha^h, \gamma^h)) = \Gamma(H, \Omega, (\beta, \gamma^h))$. Conversely $\Gamma(H, \Omega, (\beta, \gamma)) = \Gamma(H, \Omega, (\beta^{h^{-1}}, \gamma^{(h^{-1})})) = \Gamma(H, \Omega, (\alpha, \gamma^{(h^{-1})}))$.

- (ii) Suppose that the pairs (α, β) and (γ, δ) generate the same orbital graph of H . Then by [Lemma 15](#) (i) there is $h \in H$ such that $(\alpha^h, \beta^h) = (\gamma, \delta)$, which implies that $\alpha \in \gamma^H$. This proves the statement.
- (iii) If $\Gamma_1 = \Gamma_2$, then (α, β) and (α, γ) generate the same orbital graph. So by [Lemma 15](#) (i) there is $h \in H$ such that $(\alpha^h, \beta^h) = (\alpha, \gamma)$. This means that $\alpha^h = \alpha$ and therefore $h \in H_\alpha$, which implies that $\gamma \in \beta^{H_\alpha}$. Conversely, if $\gamma \in \beta^{H_\alpha}$ then there exists $h \in H_\alpha$ such that $(\alpha^h, \beta^h) = (\alpha, \gamma)$ and hence $\Gamma_1 = \Gamma_2$. \square

4.1. Futile orbital graphs

After the preparatory results above, we now characterize the situations where orbital graphs are beneficial in partition backtrack.

Definition 20 (*Futile orbital graph*). Suppose that [Hypothesis 14](#) holds and that P is an ordered orbit partition of H . We denote the stabilizer of P in G by $\text{Sym}(P)$, as we did in [Definition 1](#), and we emphasize that $\text{Sym}(P)$ stabilizes every H -orbit (i.e. every cell of the ordered partition P) as a set and that it acts as the full symmetric group on every orbit.

We say that the orbital graph Γ is **futile** if and only if $\text{Sym}(P)$, in its natural action on Ω , induces graph automorphisms on Γ .

Example 21. We refer to the graphs in [Example 13](#) for the group $H := \langle (123), (45), (46) \rangle \leq \mathcal{S}_6$, and we use the ordered orbit partition $P := [1, 2, 3|4, 5, 6]$. The first graph Γ_1 , with base-pair $(1, 2)$, is not futile. We observe that the transposition $(12) \in \mathcal{S}_6$ stabilizes the partition P , but it does not induce a graph automorphism on Γ_1 because $(1, 2)$ is an arc in Γ_1 and $(2, 1)$ is not.

The second graph Γ_2 , with base-pair $(5, 6)$, is futile.

For this we note that $\text{Sym}(P) = \langle (12), (23), (45), (46) \rangle$. Now if $g \in \text{Sym}(P)$, then g permutes the three isolated vertices in the graph and it permutes the set of vertices in the connected component containing 4, 5 and 6. Given that this connected component is a complete graph, it follows that g induces a graph automorphism on this component and hence on all of G .

Finally we look at the third graph Γ_3 with base-pair $(2, 4)$. This is a complete bipartite graph, so again we see that $\text{Sym}(P)$ induces graph automorphisms.

The intuition behind this definition is that we would like to be able to characterize orbital graphs where the graph structure does not give us any additional information compared to the orbit structure on Ω that comes from the action of H . In a futile orbital graph, all the information that could be gained from the graph structure can already be seen in an ordered orbit partition of Ω with respect to H . Building such a graph would be useless from a computational perspective. Therefore our main theoretical result on this topic classifies futile orbital graphs. We also discuss how to detect futile orbital graphs without even building them.

We note that the following result does not place any restrictions about the number of orbits of H on Ω . In particular there could be arbitrarily many isolated points in Γ .

Theorem 22. Suppose that [Hypothesis 14](#) holds. Then Γ is futile if and only if it has a unique connected component Δ of size at least 2 and moreover one of the following holds:

- (a) Δ is a complete bipartite digraph or
- (b) Δ is a complete digraph.

Proof. Let P be an ordered orbit partition of H . Then $\text{Sym}(P)$ acts on the set of orbits of H and it acts faithfully on the set of vertices of Γ . Hence to answer the question whether Γ is futile or not, we only have to consider arcs in Γ .

Throughout the proof let Δ denote a connected component of size at least 2, and without loss suppose that the arc (α, β) is contained in Δ .

We split our proof into two cases depending on whether or not Γ is a proper digraph. In both cases we begin by proving that the futility of Γ implies that Δ is the unique connected component of size at least 2 and that (a) or (b) holds, and then we discuss the converse.

Case 1: Γ is a proper digraph.

Then Γ is not self-paired and Lemma 15 (i) and (ii) imply that, for all $\omega_1, \omega_2 \in \Omega$, there is at most one arc between them. In the following arguments we will often refer to Lemma 15 (iii) and (iv) as well.

We suppose that Γ is futile and we prove in a series of little steps that Δ is the unique connected component of size at least 2 and that (a) is true.

(1) Suppose that $\gamma, \delta \in \Omega$ are distinct and in the same H -orbit. Then they are not on an arc. In particular $\alpha^H \neq \beta^H$.

Proof. As γ and δ are in the same H -orbit, they lie in the same cell of the partition P . It follows from the futility of Γ that the transposition $(\gamma, \delta) \in \text{Sym}(\Omega)$, which stabilizes P , induces a graph automorphism on Γ . Therefore neither (γ, δ) nor (δ, γ) is an arc. From this and the fact that $(\alpha, \beta) \in A$ it follows that $\alpha^H \neq \beta^H$. \square

(2) Suppose that $\omega \in \Omega$ is on an arc. Then it is either a starting point or an end point, but not both.

Proof. This follows from Lemma 15 (iii) and (1). \square

Let $S := \alpha^H$ and $E := \beta^H$, and let $I \subseteq \Omega$ denote the set of isolated vertices of Γ .

(3) $\Omega = S \dot{\cup} E \dot{\cup} I$. Moreover $S \cup E$ spans Δ , and Δ is a complete bipartite digraph.

Proof. The first statement follows from (2). Moreover there are no arcs between vertices in S or E , respectively, by (1). We show that all elements of E are on an arc with α :

For all $\gamma \in E$, we find the transposition $g := (\beta, \gamma) \in \text{Sym}(P)$, and it fixes α^H point-wise by (1). The futility of Γ implies that g maps the arc (α, β) to the arc (α, γ) . Now it follows that $A = S \times E$ and hence the digraph spanned by $S \cup E$ is a complete bipartite digraph. Then it must coincide with Δ and Δ is the unique connected component of size at least 2 of Γ . \square

Conversely, we suppose that Δ is the unique connected component of size at least 2 of Γ and that (a) holds. We prove that Γ is futile.

Let S and E denote the subsets of the vertex set of Γ such that all arcs start at S and end at E . Let I be the set of isolated vertices of Γ , so that $\Omega = S \dot{\cup} E \dot{\cup} I$.

Now $\alpha^H \subseteq S$ and the bipartite structure implies that even $\alpha^H = S$. Similarly $\beta^H = E$. Therefore $\text{Sym}(P)$ stabilizes the sets S, E and I . We already know that $\text{Sym}(P)$ permutes the vertices of Γ faithfully, so now we look at arcs.

Let $g \in \text{Sym}(P)$ and let $(\omega_1, \omega_2) \in A$. Then $\omega_1 \in S, \omega_2 \in E$ and there exists some $h \in H$ such that $(\alpha^h, \beta^h) = (\omega_1, \omega_2)$. Since $\text{Sym}(P)$ stabilizes the sets S and E , we see that $\omega_1^g \in S$ and $\omega_2^g \in E$. The completeness property then implies that $(\omega_1^g, \omega_2^g) \in A$.

Conversely, if $(\omega_1^g, \omega_2^g) \in A$, then there exists some $h \in H$ such that $(\alpha^h, \beta^h) = (\omega_1^g, \omega_2^g)$. Now $\omega_1 = \alpha^{hg^{-1}} \in S$ and $\omega_2 = \beta^{hg^{-1}} \in E$ whence $(\omega_1, \omega_2) \in A$ by completeness.

Hence Γ is futile.

Case 2: Γ is not a proper digraph, which means that it is self-paired.

We still have our connected component Δ and we begin, once more, with the hypothesis that Γ is futile. Let $\gamma \in \Omega$ be an arbitrary, non-isolated vertex.

We know that $\beta^H = \alpha^H$ by Lemma 15 (iii) and (iv), because Γ is self-paired. As γ was chosen to be a non-isolated vertex, there is some arc that starts or ends in γ . Therefore $\gamma \in \alpha^H$ and hence

α, β, γ are all in the same H -orbit and hence in a common cell of the partition P . In particular the transposition $g := (\beta, \gamma)$ is contained in $\text{Sym}(P)$ and, because of futility, it induces a graph automorphism on Γ .

Then $(\alpha, \beta) \in A$ implies that $(\alpha, \gamma) = (\alpha^g, \beta^g) \in A$. This argument shows that Δ is the only connected component of size at least 2 in Γ and that (b) is true.

We conversely suppose that Δ is the unique connected component of size at least 2 of Γ and that (b) holds. Together with the definition of orbital graphs (and the fact that arcs always go both ways in the present case) this implies that $\alpha^H = \beta^H$ spans Δ and that the isolated vertices, viewed as elements of Ω , are not contained in α^H .

We know that $\text{Sym}(P)$ acts faithfully on the vertex set of Γ . Now let $g \in \text{Sym}(P)$ and let $\omega_1, \omega_2 \in \Omega$. We recall that $\alpha^H = \beta^H$ is $\text{Sym}(P)$ -invariant.

Then it follows as in Case 1, using the completeness, that $(\omega_1, \omega_2) \in A$ if and only if $(\omega_1^g, \omega_2^g) \in A$. Consequently $\text{Sym}(P)$ acts as a group of graph automorphisms on Γ , i.e. Γ is futile. \square

We give an example in order to illustrate that futility of an orbital graph is not obvious and why further investigations into the computational usefulness of orbital graphs should be pursued.

Example 23. We let $G := S_9$ and we look at the subgroup $H := \langle (12), (13), (45), (46), (14)(25)(36), (789) \rangle$. Let Γ be the orbital graph for H with base-pair $(1, 2)$. Then Γ has the following shape:

On the vertices 1, 2, 3 and 4, 5, 6 we have a complete digraph, respectively, there is no arc between the sets $\{1, 2, 3\}$ and $\{4, 5, 6\}$, and the points 7, 8 and 9 are isolated. This might look like a futile graph, but according to the theorem it is not. Consider an ordered orbit partition $P := [1, 2, 3, 4, 5, 6 | 7, 8, 9]$ of H .

The group $\text{Sym}(P)$ contains the transposition $(24) \in G$. This element interchanges the vertices 2 and 4 of Γ and fixes 1, so this element does not induce an automorphism on Γ . (Otherwise the arc $(1, 2)$ would be mapped to the arc $(1, 4)$, which does not exist.) This graph can be used to deduce, for example, that any element which swaps 1 and 4 must also swap $\{2, 3\}$ with $\{5, 6\}$.

Hence $\text{Sym}(P)$ does not act as a group of automorphisms on Γ and we see that Γ is not futile.

It is important that we can detect futile graphs easily, without having to build them explicitly. We will now give a collection of lemmas that allow futile orbital graphs to be detected using only information about orbits and stabilizers of a group, without explicit construction of entire orbital graphs.

Lemma 24. Suppose that *Hypothesis 14* holds and that $\Omega = \alpha^H \dot{\cup} \beta^H \dot{\cup} I$, where $I \subseteq \Omega$ is the set of isolated vertices of Γ . Then Γ is futile if and only if H_α acts transitively on β^H .

Proof. Suppose that Γ is futile. Then *Theorem 22* and *Lemma 15* (iii) and (iv) imply that Γ is a complete bipartite digraph. In particular, for all $\delta \in \beta^H$ it follows that $(\alpha, \delta) \in A$ and so there exists some $h \in H$ such that $(\alpha, \delta) = (\alpha^h, \beta^h)$. In particular H_α is transitive on β^H . Conversely we suppose that H_α is transitive on β^H . It follows that for all $\beta' \in \beta^H$ there exists some $h \in H_\alpha$ such that $\beta^h = \beta'$.

We prove that H_β acts transitively on α^H , so we let $\alpha' \in \alpha^H$ and we choose $g \in H$ such that $\alpha' = \alpha^g$. Then, using the transitivity argument above, we let $h \in H_\alpha$ be such that $\beta^h = \beta^{g^{-1}}$, which implies $\beta^{hg} = \beta$ and $\alpha^{hg} = \alpha'$. Therefore H_β acts transitively on α^H . Now the definition of an orbital graph implies that Γ is a complete bipartite digraph and hence futile, by *Theorem 22*. \square

We finish this section by giving some concrete bounds on the number of edges in futile and non-futile orbital graphs.

Lemma 25. Suppose that *Hypothesis 14* holds. Let $n = |\alpha^H|$, $m = |\beta^H|$, and $I \subseteq \Omega$ be the set of isolated vertices of Γ . Then Γ is futile if one of the following holds.

- (i) $\beta \in \alpha^H$ and Γ has strictly more than $n(n - 2)$ arcs.
- (ii) $\Omega = \alpha^H \dot{\cup} \beta^H \dot{\cup} I$ and Γ has strictly more than $n(m - 1)$ or $m(n - 1)$ arcs.

Proof. To prove (i) suppose that $\gamma, \delta \in \alpha^H$ are distinct and such that $(\gamma, \delta) \notin A$. Let r be the number of arcs starting in γ . Now (γ, γ) and (γ, δ) are not in A , so it follows that $r \leq n - 2$. We recall that H is transitive on α^H , and hence all connected components of Γ have size at least 2, by Remark 16 (c). In particular γ is contained in a connected component of Γ of size at least two, so we deduce from Lemma 15 (iii) and (iv) and Lemma 18 that for every vertex of Γ , the number of arcs starting there is r . Consequently $|A| = n \cdot r \leq n \cdot (n - 2)$. This means, conversely, that Γ is a complete digraph on α^H as soon as it has strictly more than $n \cdot (n - 2)$ arcs.

To show (ii) suppose that there are $\gamma \in \alpha^H$ and $\delta \in \beta^H$ such that $(\gamma, \delta) \notin A$. Let r be the number of arcs starting in γ . As all arcs starting in γ end in a vertex of $\beta^G \setminus \{\delta\}$ it follows that $r \leq m - 1$. Let $\omega \in \alpha^H$. Then it follows from Lemma 15 (iii) that the number of arcs starting in ω is $|\{(\omega_1, \beta^g) \mid g \in H, \alpha^g = \omega_1\}|$. Hence $|A| = n \cdot r \leq n \cdot (m - 1)$.

By counting the number of arcs ending in some vertex we obtain, in a similar way, that $|A| \leq m \cdot (n - 1)$ as well. Hence if Γ has strictly more than $n \cdot (m - 1)$ or $m \cdot (n - 1)$ arcs, then Γ is a complete bipartite digraph. \square

In practice we use Corollary 26, which combines Lemma 25 with Remark 16 to efficiently identify futile orbital graphs before they are constructed.

Corollary 26. *Suppose that Hypothesis 14 holds. Then Γ is futile if and only if one of the following conditions is true:*

- (i) $\beta \in \alpha^H$ and $|\beta^{H_\alpha}| = |\alpha^H|$.
- (ii) $\beta \notin \alpha^H$ and $|\beta^{H_\alpha}| = |\beta^H|$.

Proof. (i) We are in Case (i) of Lemma 25. By Remark 16 the orbital graph has size $|\alpha^H| \cdot |\beta^{H_\alpha}|$. The only way this can be larger than $|\alpha^H|(|\alpha^H| - 2)$ is if $|\beta^{H_\alpha}| + 1 \geq |\alpha^H|$. As $\beta \in \alpha^H$, we see that β^{H_α} is a proper subset of α^H (the subset is proper because it does not contain α). Therefore $|\beta^{H_\alpha}| + 1 = |\alpha^H|$.
 (ii) We are in Case (ii) of Lemma 25. Again by Remark 16 the orbital graph has size $|\alpha^H| \cdot |\beta^{H_\alpha}|$. The only way this can be larger than $|\alpha^H|(|\beta^H| - 1)$ is if $|\beta^H| \leq |\beta^{H_\alpha}|$. As β^H contains β^{H_α} , this implies that $|\beta^H| = |\beta^{H_\alpha}|$. \square

Lemma 27. *Suppose that Hypothesis 14 holds and that H acts transitively on Ω .*

- (i) *If H acts 2-transitively on Ω , then Γ is futile.*
- (ii) *If Γ is futile, then H acts 2-transitively on Ω (and hence all orbital graphs are futile).*

Proof. For (i) we suppose that H acts 2-transitively on Ω . Then whenever $\gamma, \delta \in \Omega$ are distinct, there exists some $h \in H$ such that $(\alpha^h, \beta^h) = (\gamma, \delta)$ and hence Γ is a complete digraph. By Theorem 22 it follows that Γ is futile.

For (ii) we suppose that Γ is futile and we deduce, again by Theorem 22, that Γ is a complete digraph or a complete bipartite digraph. The second case is impossible because H is transitive on Ω . So Γ is a complete digraph and for any two distinct elements $\gamma, \delta \in \Omega$, we deduce that $(\gamma, \delta) \in A$. Then by definition of an orbital graph, there is $h \in H$ such that $(\alpha^h, \beta^h) = (\gamma, \delta)$. Hence H acts 2-transitively on Ω and the last statement follows from (i). \square

So we see that for transitive groups if one orbital graph is futile, then all of them are. Lemma 27 lets us quickly detect this, as the level of transitivity of a group can be efficiently calculated.

4.2. Efficiently creating orbital graphs

While orbital graphs can be very useful in reducing search, they are expensive to create, so we only want to compute them when they provide extra refinements. We use [Algorithm 1](#) to compute orbital graphs, which assumes the use of a computational group theory system, such as GAP, that provides basic algorithms to compute point stabilizers, orbits of points, and orbits of pairs of points.

Algorithm 1 Find orbital graphs.

```

1: procedure ORBITALBASE( $G, \Omega, SizeLimit$ ) ▷ 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

```

The correctness of [Algorithm 1](#) is proven by applying results from the preceding sections, in particular [Lemma 19](#), [Lemma 27](#), and [Theorem 22](#).

Theorem 28. *Algorithm 1 returns all orbital graphs, except those that are futile or that contain more than $SizeLimit$ edges.*

Proof. First, Line 3 performs an initial check whether the group is k -transitive for $k \geq 2$. If it is, then we stop because, by [Lemma 27](#)(i), 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 a base-pair. Here we apply [Lemma 19](#): In order to generate all orbital graphs, it is enough to pick one element from each orbit as first member, by Part (i) of the lemma. Then Part (ii) shows that no orbital graph will be generated twice. Lines 7 and 8 pick one member from each of the orbits of the stabilizer of each of our first base-pair points. Part (iii) of the [Lemma 19](#) shows that this will give us a set of base-pairs from which every orbital graph arises exactly 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 [Theorem 22](#), rejecting all futile orbital graphs. Line 9 allows us to reject any orbital graph that exceeds a user-defined limit on the number of edges.

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

The runtime of [Algorithm 1](#) 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 1](#) performs well in practice.

5. Graph refiners

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 29 (Equalizer). Let $\Gamma = (\Omega, A)$ be a digraph.

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 some $k \in \mathbb{N}$ such that for all elements i in the cell Δ it holds that $|\{j \in \Delta' \mid (i, j) \in A \text{ or } (j, i) \in A\}| = k$.

We note that in this definition the number of arcs 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 30. Some remarks on the previous definition:

If $\Gamma = (\Omega, A)$ is a digraph and $P \in \text{OPart}(\Omega)$, then two distinct Γ -equalizers for P can only differ by the ordering of their 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 discuss a simple algorithm that computes a Γ -equalizer for P , given a digraph Γ and an ordered orbit partition P as input. The main difference between the algorithm in McKay and Piperno (2014), and our implementation, is that McKay and Piperno optimize this algorithm by showing several cases where they can skip some attempts to split because they know that no splitting will occur. We omit these improvements because the total time taken by the refinement algorithm in our partition back-tracker is usually negligible.

Algorithm 2 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 our new refiner.

Definition 31 (Orb_M). Given M a subgroup of $\text{Sym}(\Omega)$, the map $\text{Orb}_M : \text{OPart}(\Omega) \rightarrow \text{OPart}(\Omega)$ is defined as follows:

- Construct all orbital graphs of M .
- 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 2.
- 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 Orb_M is in fact a refiner. A detailed example later in this section will illustrate how this refiner works.

Lemma 32. Orb_M is a refiner.

Proof. For all $P \in \text{OPart}(\Omega)$, it holds that $\text{Orb}_M(P) \preceq P$, because Algorithm 2 splits cells of P . Next, we need to show that for all $g \in M$ 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 by Lemma 6. \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 (see Lemma 27). We therefore introduce another orbital graph based refiner that makes use of the fact that we can, like in Fixed, easily stabilize points in the group.

Definition 33 (*DeepOrb_M*). Given M a subgroup of $\text{Sym}(\Omega)$, 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 let M_0 denote the point-wise stabilizer of $\alpha_1, \dots, \alpha_k$ in M .
- Construct all orbital graphs of M_0 .
- 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 2.
- 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 level in the search, rather than just once at the beginning. Our experiments will investigate the practical trade-offs between Orb and DeepOrb .

5.1. Examples of refinements via graphs

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

We follow McKay and Piperno (2014), referring to partitions as colourings. We will refer interchangeably to colouring vertex j of the graph with colour δ and placing value j into cell δ of the ordered partition.

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

Example 34. Let $H_1 := \langle (1, 2, \dots, 10), (2, 10)(3, 9)(4, 8)(5, 7) \rangle$ and let H_2 denote the stabilizer of the set $\{1, 5\}$ in G . We are interested in calculating $D := H_1 \cap H_2$, which is equivalent to calculating the stabilizer of $\{1, 5\}$ in H_1 . While this problem is very simple, it allows us to show in detail how the algorithm works.

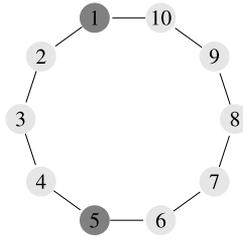
Purely group-theoretically, if we take $x \in D$, then all we know without further calculation is that x stabilizes the set $\{1, 5\}$ and so we obtain the ordered partition $P_1 := [1, 5 \mid 2, 3, 4, 6, 7, 8, 9, 10]$.

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 that is isomorphic to $\text{Sym}(P_1)$.

Now we use the orbital graph for H_1 with base-pair $(1, 2)$. For simplicity we work with an undirected graph here because all arcs exist in both directions. When calculating equitable partitions, 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 (first cell of the partition) and colour 2 to all other vertices (second cell of the partition).



Step 2:

Using Algorithm 2, 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 $P_2 := [1, 5 \mid 3, 7, 8, 9 \mid 2, 4, 6, 10]$ and D is isomorphic to a subgroup of $\text{Sym}(P_2)$.

Step 3:

Continuing the algorithm, we see that the vertices of colour 2 can be further divided into a single vertex that has two neighbours of colour 3 (vertex 3), those that have two neighbours of colour 2 (vertex 8) and those that have one neighbour of colour 2 and one of colour 3 (vertices 7 and 9). With this information we obtain the new ordered partition $P_3 := [1, 5 \mid 7, 9 \mid 2, 4, 6, 10 \mid 3 \mid 8]$, and D is isomorphic to a subgroup of $\text{Sym}(P_3)$.

Step 4:

Our algorithm continues, looking at cell 3. Here the vertices can be divided into two categories, namely those with 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 $P_4 := [1, 5 \mid 8 \mid 6, 10 \mid 3 \mid 7, 9 \mid 2, 4]$.

There is no further information in the graph at the moment, so we conclude by stating that D is isomorphic to a subgroup of $\text{Sym}(P_4)$, which has order 16.

Given that P_4 is an ordered orbit partition for $H_1 \cap H_2$, it is no surprise that no further refinement is possible. We also know that $D \leq H_1$ and $|H_1| = 20$, so $|D| \leq 4$.

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 using the orbital graph that cannot be performed by Fixed.

Example 35. Again $H_1 := \langle (1, 2, \dots, 10), (2, 10)(3, 9)(4, 8)(5, 7) \rangle$, but this time H_2 is the stabilizer of the set $\{1, 6\}$ in G . 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 $\text{Fixed}_{H_2}(Q_1) = Q_1$ 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 that 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 $\text{Sym}(Q_2)$ 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 8. 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 $\text{Co}(P_1, Q_1)$.

6. Experiments

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

1. How much can we speed up the search?
2. What are the worst-case slowdowns when using orbital graphs?

Generating a set of “random” problems – taken from all possible group intersection and stabilizer problems – is not possible because there is no method of producing random permutation groups of large size. Also, we do not want to consider problems that are very simple – partition backtrack solves many problems almost instantly, with or without orbital graphs. We consider two main categories of problems: The first one is calculating set stabilizers in “grid groups” that arise in A.I. in SAT solvers (the Boolean satisfiability problem) and Constraint Programming and are used for symmetry-breaking, and the second one is taking intersections of primitive groups with wreath products. We discuss 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 that keeps building orbital graphs until the first refinement phase at which at least one non-futile orbital graph is found. Then these orbital graphs are used in all later refinement phases.

FirstOrbital is implemented as a variant of **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-futile orbital graph at the same height. This is a consequence of the definition of the Fixed refiner – it ensures that there is an element of the group that maps the fixed points down the first branch to the fixed points down every other branch at every level.

We expect that **Fixed** will always produce larger search trees than **PreOrbital**, which in turn produces 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 2.27 GHz and 20GB RAM. Each experiment was given a five minute timeout and a limit of 1GB of RAM.

Table 1

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

m	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

6.1. Set stabilizers in grid groups

A typical example for a problem that involves grid groups is shift scheduling. If we have n workers and m time-slots, where the workers are interchangeable (in terms of their qualification) and the time-slots are equally important, then this symmetry in the system can be expressed via a grid group.

Definition 36 (Grid group). Let $n \in \mathbb{N}$. The direct product $\mathcal{S}_m \times \mathcal{S}_m$ acts on the set $\{1, \dots, m\} \times \{1, \dots, m\}$ of pairs in the following way:

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

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

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

While the construction of the grid group is done by starting with an m by m grid of points and permuting rows or columns independently of each other, we represent this group as a subgroup of \mathcal{S}_{m^2} , 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. The specific choice of sets in the second variant leads to more difficult problem, intuitively, because it is harder to prove that the stabilizer does not permute the rows of the grid.

First, we generate a random set of size $\lfloor \frac{m^2}{2} \rfloor$ of $\{1 \dots m^2\}$. The results are shown in [Table 1](#). For times, we use the timeout time (5 minutes) for instances that ran out of either time or memory. We do not display results for **FirstOrbital**, because they are identical to **PreOrbital**. Our problems are randomly generated and therefore 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 rebuilding 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. As mentioned earlier, we expect these instances to be more difficult. The results of this experiment are shown in [Table 2](#). The major difference here is that no instances of size bigger than 30 were solved by **Fixed**, while **PreOrbital** and

Table 2

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

m	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

Table 3

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

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

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 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. This results in more graphs being created. We conclude that the performance of **DeepOrbital** can be very unpredictable.

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

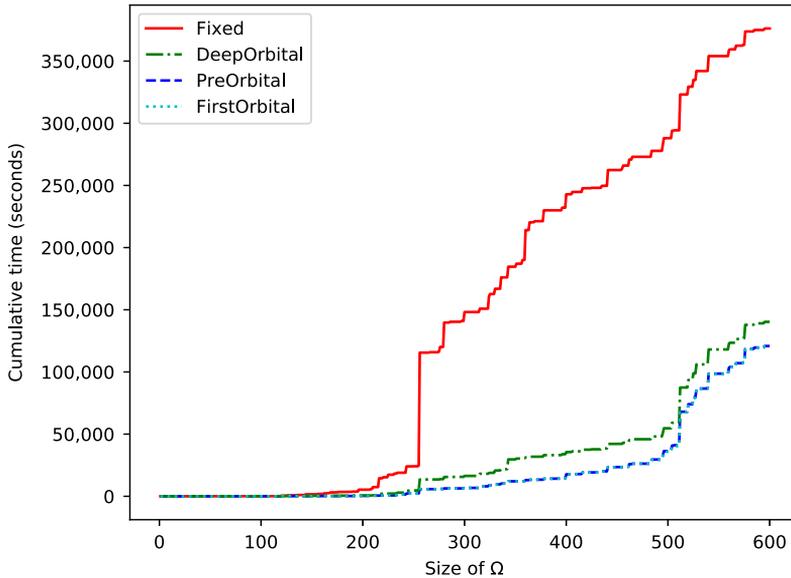


Fig. 1. Cumulative time taken to intersect primitive and not 2-transitive primitive groups with a wreath product

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.

6.2. Intersection of primitive groups

In our second experiment, we consider intersection of groups. We concentrate on primitive groups because the primitive group library in GAP provides easy access to a large set of groups. First we intersected pairs of primitive groups, but then it turned out that this problem is usually extremely fast to solve, with or without orbital graphs. Therefore we instead intersect a primitive group with a wreath product of symmetric groups, and we found that this produced challenging problems.

For each $n \in \mathbb{N}$ we take all 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. We remove these groups, because GAP handles them as a special case, and intersections involving these groups are simple to construct. All 2-transitive groups are primitive and therefore we will encounter many 2-transitive groups in this experiment. However, we already know from Lemma 27 that they have futile orbital graphs. Therefore we consider 2-transitive groups separately.

Given a primitive group G acting on n points, we create the wreath products $S_{n/x} \wr S_x$, where $x \in \{2, \dots, 7\}$ and n/x is an integer. We conjugate each of these wreath products by a random element of S_n and intersect the result with G . We experimented with other wreath products and found similar results. With primitive groups on up to 600 points, we produce a total of 2,752 experiments on primitive groups that are not 2-transitive and 1,140 experiments for 2-transitive groups. We ran each experiment for a maximum of five minutes. Note that we build the orbital graphs for both the primitive groups and the wreath products. The orbital graphs on wreath products are often very large. There are two orbital graphs of the wreath product of $S_a \wr S_b$, one 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 Fig. 1. Here we can see that all the orbital techniques solve problems on average much faster. In particular, the many primitive groups on 256 points cause severe difficulties for the **Fixed** refiner. There is not a similar spike for 512 points, because we only include problems at least one al-

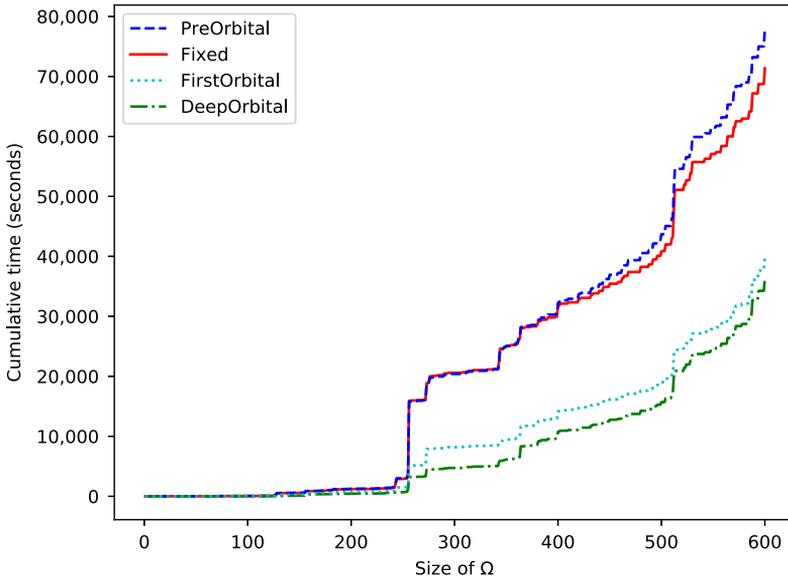


Fig. 2. Cumulative time taken to intersect 2-transitive groups with a wreath product

Table 4

The relative performance, compared to **Fixed**, on Primitive Intersection experiments.

		Orbital much slower	Orbital slower	Orbital faster	Orbital much faster
Not 2-trans	PreOrbital	13	157	149	1519
	DeepOrbital	297	191	77	1358
	FirstOrbital	12	146	162	1519
2-trans	PreOrbital	114	526	63	0
	DeepOrbital	38	197	206	351
	FirstOrbital	33	176	219	346

gorithm solved within the timeout. Our results 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 Fig. 2. 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**.

In Table 4 we compare the algorithms **PreOrbital**, **DeepOrbital** and **FirstOrbital** to **Fixed**. The column “much faster” captures instances where the algorithms involving orbital graphs are at least two times faster than **Fixed**, and similarly the column “much slower” contains cases where orbital graphs are at least two times slower. Table 5 shows how much time 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 acting on over 500

Table 5
Percentage of runtime spent building stabilizer chains and orbital graphs in Primitive Intersection experiments.

Problem	Refiner	Build Stab Chains	Build Orbital Graphs	Refine Orbital Graphs
Not	Fixed	99.7	0	0
2-trans	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

points **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.

6.3. 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. Here the only limit is 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 building an orbital graph is larger – wreath products of symmetric groups in particular produce large orbital graphs. For groups that are not 2-transitive, all our orbital graph methods perform similarly. For 2-transitive groups however, we see that **PreOrbital**, as expected, is slightly slower than **Fixed**, but **FirstOrbital** and **DeepOrbital** perform well. There is a small set of problems that 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. We suggest always using **FirstOrbital**, because often it is much better than **Fixed**, and it is close to the best of the algorithms on all our problem classes.

7. 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 – if it is possible to efficiently feed structural properties of groups or their action under consideration, 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 run-times 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, because for many problems we are reaching the limit of improvements that can be provided by better refiners alone. This is

because the search involves almost no branching anymore. In particular, our experiments suggest that for more substantial progress we must either speed up the calculation of stabilizer chains, or reduce the number of stabilizer chains that must be found.

References

- Bang-Jensen, J., Gutin, G.Z., 2008. *Digraphs: Theory, Algorithms and Applications*, 2nd edition. Springer Publishing Company.
- Bosma, W., Cannon, J., Playoust, C., 1997. The Magma algebra system. I. The user language. In: *Computational Algebra and Number Theory*, London, 1993. *J. Symb. Comput.* 24 (3–4), 235–265. <https://doi.org/10.1006/jsco.1996.0125>.
- Cameron, P., 1999. *Permutation Groups*. London Math. Soc. Student Texts. Cambridge University Press.
- 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. <http://www.gap-system.org>.
- Jefferson, C., 2016. The Ferret Package. <https://github.com/gap-packages/ferret>.
- Leon, J.S., 1991. Permutation group algorithms based on partitions, I: theory and algorithms. *J. Symb. Comput.* 12 (4), 533–583. <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. *J. Symb. Comput.* 60, 94–112. <http://www.sciencedirect.com/science/article/pii/S0747717113001193>.
- Seress, Á., 2003. *Permutation Group Algorithms*. Cambridge Tracts in Mathematics. Cambridge University Press. 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.