Quantifying symmetry

Jonathan A. Cohen

Abstract

Symmetry is all pervasive – from the day/night cycle to the rise and fall of the tides, from geometry to physics. It is thus natural that mathematicians should want to study symmetry, to quantify it, to *exploit* it. At its core, this article is about quantifying the amount of symmetry in an object and separating those that, in a certain well-defined sense, possess a *large* amount of symmetry from those that possess only a *small* amount.

1 A first attempt: size and lattice

How does one make precise the rather vague notion of the "amount of symmetry" in an object? The difficulty of this question is exacerbated by the many possible interpretations of "symmetry". We can, however, make the question slightly more manageable by restricting our attention to permutation groups, that is, to subgroups of symmetric groups. Let us rephrase our question in this context: How can one turn the *qualitative* notion that a permutation group has a "large" amount of symmetry into a *quantitative* notion?

Reflecting on this question, a reasonable expectation is that the symmetric group on n points, S_n, contains the "most" amount of symmetry for any group on n points and that the alternating group, A_n, contains the second most. At this level, one may make use of the group order - certainly, S_n is larger than any other group on n points and A_n is a maximal subgroup of index 2. Unfortunately, group order is too coarse a measure for our purposes. For instance, let us consider groups of order 8. In particular, we focus our attention on the dihedral group on 4 points, D_4, and the cyclic group on 8 points, C_8.

Now, D_4 arises naturally as the automorphism group of a square, whereas C_8 can be seen as the rotation group of a regular octagon. However, one may gainfully argue that D_4 possesses *more* symmetry, as it contains *every* symmetry of a square, whereas C_8 only contains half of the symmetries of a regular octagon. So, perhaps subgroup structure has something to do with quantifying symmetry? Looking at the nontrivial subgroups of D_4, we see that it possesses three subgroups of order 4 and five of order 2. One of the subgroups of order 4 corresponds to the rotation group of a square, whereas all but one of the subgroups of order 2 correspond to the group generated by reflecting the square about some axis (a diagonal, say). The key point, however, is that these subgroups behave *qualitatively* differently. On the other hand, C_8 possesses only one subgroup of order 4 and one of order 2, with the subgroup of order 2 contained in the subgroup of order 4. That is, the groups are *qualitatively* indistinguishable, both corresponding to rotations of the octagon.

2 A better attempt: bases

Unfortunately, looking at subgroup lattices is unlikely to solve our problem in general, as it becomes computationally intractable for groups of even moderate size. That is, the number

Jonathan Cohen is the winner of the 2004 B.H. Neumann prize. This contribution is a written version of his winning talk presented at the 2004 AustMS meeting. See also AustMS Gazette **31** (2004), 298–300.

of elements in the subgroup lattice grows to such an extent that it becomes very difficult to read off any useful information. So, we need to focus our attention on specific subgroups, which are more closely related to the actions of our groups.

Suppose that one has a permutation group G on a set Γ. Given some point $\alpha \in \Gamma$, a natural set to look at is all those permutations in G that do not move α. This is called the *pointwise stabiliser* of α in G (or stabiliser for short). It is not hard to see that this set forms a sub*group* of G and it is usually denoted by G_α. We can once more pick a point $\beta \in \Gamma$ and look at the stabiliser of β in G_α. Continuing in this manner, we can stabilise a sequence of points in Γ.

Let us now return to our previous example and see how stabilisers allow us to resolve our question in this context. Firstly, if we stabilise a vertex, α, of the square on which D_4 is acting, then there is still a nontrivial permutation contained in this stabiliser. This permutation reflects the square about the line passing through α and its diagonally opposite vertex. If we subsequently stabilise a vertex not on this line, then we obtain the trivial group. The situation is slightly different for C_8: every vertex on which it is acting is moved by every nontrivial permutation contained in C_8, so the stabiliser in C_8 of any vertex of the octagon is trivial. Success! We have managed to find a seemingly reasonable method for determining that D_4 has qualitatively more symmetry than C_8. Let us formalise our discovery in a definition.

Definition 1 (Base) Let G be a group acting on a set Γ. A *base* for G is a finite sequence of points of Γ whose stabiliser in G is trivial.

In the next few sections, we shall see a little more about how bases relate to the structure of the underlying group, why they are important and also some of the tools that have been used to exploit them.

3 A closer look at bases

Calling the sequence in Definition 1 a *base* may cause one to think about vector spaces and bases thereof. Indeed, there is a slight analogy in that, just as it is the case that a linear transformation on a vector space is uniquely determined by its image on a set of basis vectors, so it is that an element of a permutation group G is uniquely determined by the image of its action on a base for G. Moreover, only the trivial linear transformation does not alter any member of a basis for a vector space.

Sadly, the analogy with vector spaces starts to break down once one begins to look slightly deeper. Noting that stabilising a sequence of points is a recursive process such that each step of the recursion leads to a new subgroup, we come to the following definition.

Definition 2 (Stabiliser Chain) Let G be a permutation group and $\mathcal{B} = [\alpha_1, \alpha_2, \ldots, \alpha_k]$ be a base for G. Then, the *stabiliser chain for G relative to $\mathcal{B}$* is

$$G = G^{(1)} \geq G^{(2)} \geq \cdots \geq G^{(k+1)} = \{1\}, \tag{1}$$

where for each i with $2 \leq i \leq k+1$, the group $G^{(i)}$ is the stabiliser in G of the first $i-1$ points of $\mathcal{B}$.

This is where the analogy with vector spaces begins to go awry. Perhaps the most fundamental fact about a set of basis vectors for a vector space is that its cardinality is uniquely determined by the vector space. Unfortunately, the same does not ring true for permutation groups. For instance, there is nothing in Definition 1 that requires the elements of a base to be unique. Thus, any group with a base of size k has a base of size m, where m is any integer bigger than k (since the elements of a base need not be unique). This may seem

like a sly way to obtain a counterexample: the cardinality of a basis for a vector space ceases to be unique if one drops the requirement that the vectors be linearly independent. Very well then, let us call a base *irredundant* if all of the inclusions in the stabiliser chain that it induces are strict and *redundant* otherwise. Surely, then, the length of an irredundant base for a permutation group must be invariant?! Alas, this is not the case and counterexamples are not hard to find. For instance, one may take a cyclic group of order 4 acting irregularly, say the group, G, generated by the permutation $(1,2)(3,4,5,6)$. Upon a little scrutinising, it is evident that both $\langle 3 \rangle$ and $\langle 1,3 \rangle$ are irredundant bases for G. So, not only is it the case that the length of an irredundant base is not an invariant of a group, but one must also take heed of the *order* in which points appear in the base.

Thus, it seems to be the case that the thing we really ought to be interested in is the *minimal* length of a base for a group. Earlier, we saw that the minimal length of a base for D_4 is 2 and for C_8 is 1. In fact, our observations easily generalise to any dihedral or cyclic group. That is, the minimal length of a base for D_n is 2 and for C_n is 1. The following lemma treats slightly more interesting examples.

Lemma 1 *The minimal length of a base for the symmetric group, S_n, is $n-1$ and for the alternating group, A_n, is $n-2$.*

Proof. For S_n, suppose that we have stabilised fewer than $n-1$ points. Then, there are at least 2 points which have not been stabilised and, since S_n contains *all* permutations of an n element set, the transposition which interchanges these two points is contained within the stabiliser, so the stabiliser is nontrivial. If, on the other hand, we stabilise $n-1$ distinct points, then there is only one point left which has not been explicitly stabilised, so the stabiliser contains only the trivial permutation, hence is trivial.

The argument for A_n is similar, except we only need to stabilise $n-2$ distinct points, as A_n contains no transpositions. □

Henceforth, we use the notation $b(G)$ to denote the minimal length of a base for a permutation group G. The question, then, is how much information is conveyed by $b(G)$ for a given group G. For instance, it follows immediately from the above lemma that $b(A_n) = b(S_{n-1}) = n-2$. Surely, though, a symmetric group has slightly more symmetry than an alternating group, given the presence of transpositions? We can alleviate this objection by considering not individual groups, but *infinite families* of groups. If we compare the family of *all* alternating groups with the family of *all* symmetric groups then, for any given n, the value of $b(S_n)$ just edges out that of $b(A_n)$, as it ought to do. Well, alternating and symmetric groups are still very large compared to, say, dihedral groups, so we still need to dream up a criterion for when we call a family of groups "large". The following definition turns out to be very reasonable, as we discover in the next section. For the moment, however, a little trust on the reader's part is required.

Definition 3 (Small/Large Base) Let $\mathcal{G}$ be an infinite family of permutation groups. We call $\mathcal{G}$ *small base* if there are positive integers a and b such that each $G \in \mathcal{G}$ acting on a set of cardinality n admits a base of size $a \log^b(n)$. Otherwise, we call $\mathcal{G}$ *large base*.

We shall sometimes abuse Definition 3 slightly by referring to individual groups as small/large base when it is clear what infinite family they belong to. Notice that this definition neatly separates the groups we have studied thus far: by Lemma 1, both the alternating groups and symmetric groups form large base families. However, from our previous discussion, we know that both the dihedral and cyclic groups have constant length minimal base size, so are certainly both small base.

4 A little history

Let us now briefly consider one of the original motivations for studying base size. Recall our off-the-cuff remark that that an element of the permutation group G is uniquely determined by its action on a base, $\mathcal{B}$, for G. The proof of this fact proceeds as follows. Consider the stabiliser chain

$$G = G^{(1)} \geq G^{(2)} \geq \cdots \geq G^{(k+1)} = \{1\}$$

for G relative to $\mathcal{B}$. For each i, with $1 \leq i \leq n$, let R_i be a set of coset representatives for $G^{(i)}/G^{(i+1)}$. Let $g \in G^{(1)}$. Then, it follows from the Orbit-Stabiliser Theorem that there is a unique $r_1 \in R_1$ such that $g \in G^{(2)}r_1$. Noting that $gr_1^{-1} \in G^{(2)}$, we may continue by induction and write g uniquely as a product $r_1r_2 \dots r_k$, where each $r_i \in R_i$. Conversely, each such product is a member of G. Computational group theorists call this procedure *sifting* and it provides an efficient method for testing membership in a permutation group. Along with other useful algorithms that make essential use of bases, this method was originally proposed in two seminal papers by Charles Sims [1, 2] and variants of them are still in wide use today [3].

Thus, being able to construct a base is of vital importance for computing with permutation groups. It is for this reason that small base groups are especially attractive, since the rate of growth of their bases is such that the classical methods are computationally tractable. Indeed, for much of the brief history of computational permutation group theory, small base groups were the only class that was treated effectively when the degree of the group is large. Only relatively recently has there been progress on algorithms for computing with "big" large base groups (see, e.g., [3, Chapter 10]). So, it is a prudent endeavour to attempt to identify those families of groups that are large base, so that we can develop specific algorithms for handling them. We embark on such a search in the following section.

5 Towards a characterisation of large base groups

In the last section, we caught a glimpse of why it is important to characterise permutation groups according to base size. In this section, we set about this task. As it turns out, the journey is perilous and we need to make use of some very heavy duty machinery in order to make much progress.

From the onset, however, looking at *all* permutation groups does not give us much to grab hold of, since an arbitrary permutation group has very little combinatorial structure. So, we look at certain subclasses of permutation groups. We call a permutation group G on a set Γ *transitive* if we can get from any point of Γ to any other point of Γ via the action of a group element. This property can also be put slightly differently. For $\alpha \in \Gamma$, we call the set of images of α under members of G the *orbit* of α under G and the collection of all such orbits partition Γ. Then, calling G transitive amounts to saying that it has only one orbit on Γ.

Unfortunately, the class of all transitive groups is still rather broad, so we place a further restriction on the groups that we look at. We say that a permutation group leaves a partition *invariant* if it leaves it unchanged as a collection of sets. Every permutation group G acting on a set Γ leaves the singleton partition $\{\{\alpha\} \colon \alpha \in \Gamma\}$ invariant, since any permutation must map a singleton to another singleton. Dually, every such group leaves the partition $\{\Gamma\}$ invariant. If G is transitive and leaves no other partitions invariant, then we say that G is *primitive*.

Primitive groups turn out to be a good class to look at, since any intransitive group can be embedded as a "subdirect product" of its transitive constituents (which correspond in a

well-defined sense to its orbits) and a transitive but imprimitive group can be embedded in an "iterated wreath product" of primitive groups. As we shall not actually need these constructions, we refer the reader to [4] for details. Unfortunately, however, the decomposition into primitive groups can not, in general, be carried out in a unique way. That is, there may be more than one way to "build" up a group from primitive "pieces"; see [5] for further details. Nevertheless, it is often possible to "lift" a result about primitive permutation groups to more general classes.

For this reason, the theory of primitive permutation groups is very highly developed. Typically, one is led to analyse the subgroup structure of a primitive group. The subgroup which turns out to be most fruitful to look at is the so-called *socle*, which for an arbitrary group G is defined to be the subgroup generated by the minimal normal subgroups[1] of G. It turns out that the socle of a primitive permutation group has a particularly perspicuous structure, being a direct product of isomorphic simple groups [6, Corollary 4.3B].

It is, of course, possible to go through all of the possibilities for a finite simple group, by using the classification of finite simple groups [7]. Although, since the proof of this theorem currently weighs in at a whopping 15,000 or so pages, one loses any intuition as to *why* the result holds. Nevertheless, we push on with this approach. In order to do this, one needs slightly more information about the possible structures for the socle of a primitive permutation group. This is the content of the celebrated O'Nan-Scott theorem [8, 9], which can be equivalently stated as classifying the maximal subgroups of the symmetric and alternating groups. Of course, the maximal subgroups of the "largest" groups of any degree is a prudent place to look for other large base groups. This was the technique utilised by Cameron [5] and later extended by Liebeck [10] in order to characterise the large base primitive permutation groups.

Before giving a statement of the characterisation, we need to cover a certain construction, called the *wreath product in its product action*[2]. Given a permutation group H acting on a set Δ, as well as a permutation group G acting on a set Γ, where $|\Gamma| = n$, the construction proceeds as follows. Take the cartesian product Δ^n and, for each component, create a new copy of H and allow it to act on that component naturally. Then, allow G to permute the indices of Δ^n in the same way as it permutes Γ. We denote this construction[3] by $H \wr G$. We are now in a position to state Liebeck's result.

Theorem 1 (Liebeck [10]) *If G is a primitive group of degree n then one of the following holds.*

(i) *G is a subgroup of $S_m \wr S_r$ containing $(A_m)^r$, where the action of S_m is on k-element subsets of $\{1, 2, \ldots, m\}$ and the wreath product is in its product action of degree m^r.*
(ii) *$b(G) < 9 \log(n)$.*

This theorem makes sense intuitively, since it effectively says that the only large base primitive permutation groups are those that contain a "big" alternating group. It also says slightly more, namely that the gap in minimal base size between the large and small base groups is quite large, since the latter is logarithmically bounded. By setting $r = k = 1$ in Theorem 1, we obtain our previous result that the symmetric and alternating groups are large base.

Lifting the last theorem to larger classes of permutation groups has proven to be rather difficult. Praeger and Shalev [11] have lifted it to those permutation groups all of whose

[1] That is, they do not themselves contain a nontrivial normal subgroup of G.
[2] This is different from the more commonly used, or *imprimitive*, action of a wreath product.
[3] Algebraically, this is the semidirect product $H^n \rtimes G$.

minimal normal subgroups are transitive (all minimal normal subgroups of a primitive permutation group are transitive, but the converse need not hold). Subsequently, Bamberg [12] has managed to lift the result to groups that contain at least one minimal normal subgroup. Lifting it all the way to the class of transitive groups seems out of reach at the moment. Indeed, it is unlikely that the characterisation is the same at this level of generality.

6 Acknowledgments

I learnt almost all that I know about group theory while I was an honours student at The University of Western Australia. Thanks go to my supervisors Alice Niemeyer and John Bamberg as well as to the other members of the Groups and Combinatorics research group.

References

[1] C.C. Sims, *Computational methods in the study of permutation groups*, in: J. Leech (ed.), *Computational Problems in Abstract Algebra*, (Pergamon Oxford 1970), 169–183.
[2] C.C. Sims, *Computation with permutation groups*, in: *Proceedings of the second ACM symposium on Symbolic and algebraic manipulation*, (ACM Press Los Angeles 1971), 23–28.
[3] A. Seress, *Permutation group algorithms*, (Cambridge University Press Cambridge 2003).
[4] P.J. Cameron, *Permutation groups*, (Cambridge University Press Cambridge 1999).
[5] P.J. Cameron, *Finite permutation groups and finite simple groups*, Bull. London Math. Soc. **13** (1981), 1–22.
[6] J.D. Dixon, and B. Mortimer, *Permutation groups*, (Springer-Verlag New York 1996).
[7] D. Gorenstein, R. Lyons, and R. Solomon, *The classification of the finite simple groups*, (American Mathematical Society Providence 1994).
[8] M.W. Liebeck, C.E. Praeger, and J. Saxl, *On the O'Nan-Scott theorem for finite primitive permutation groups*, J. Austral. Math. Soc. Ser. A **44** (1988), 389–396.
[9] L.L. Scott, *Representations in characteristic p*, in: B. Cooperstein and G. Mason (eds.), *The Santa Cruz Conference on Finite Groups*, (American Mathematical Society Providence 1980), 319–331.
[10] M.W. Liebeck, *On minimal degrees and base sizes of primitive permutation groups*, Arch. Math. (Basel) **43** (1984), 11–15.
[11] C.E. Praeger, and A. Shalev, *Bounds on finite quasiprimitive permutation groups*, J. Aust. Math. Soc. **71** (2001), 243–258.
[12] J. Bamberg, *Bounds and quotient actions of innately transitive groups*, to appear in J. Aust. Math. Soc..

Computer Sciences Laboratory, Research School of Information Sciences and Engineering, The Australian National University, Canberra ACT 0200
E-mail: Jonathan.Cohen@anu.edu.au

Received 11 April 2005, accepted 20 April 2005.