FINDING MINIMAL PERMUTATION REPRESENTATIONS OF FINITE GROUPS

BEN ELIAS, LIOR SILBERMAN, AND RAMIN TAKLOO-BIGHASH

ABSTRACT. A minimal permutation representation of a finite group G is a faithful G-set with the smallest possible cardinality. We study the structure of such representations and show that for most groups they may be obtained by a greedy construction. It follows that whenever the algorithm works (except when central involutions intervene) all minimal permutation representations have the same set of orbit cardinalities. Using the same ideas we also show that if the size $d(G)$ of a minimal faithful G-set is at least $c|G|$ for some $c > 0$ then $d(G) = |G|/m + O(1)$ for an integer m, with the implied constant depending at most on c.

1. INTRODUCTION

It is a classical theorem of Cayley's that a group G is isomorphic to a subgroup of a symmetric group. Accordingly we let the *degree* of the finite group G, denoted $d(G)$, be the least integer d such that G can be embedded in S_d, the symmetric group on d letters. More precisely, Cayley's discussion in [3] implicitly relies on the observation that the regular action of the group on itself gives an embedding of G into S_n, where $n = |G|$ is the order of G. It is then natural to ask to what extent the resulting bound $d(G) \leq n$ is sharp.

The problem of finding $d(G)$ was first studied by Johnson [7]. Among other things, he classified those groups for which $d(G) = n$. Except for a family of 2-groups, these groups are precisely the cyclic p-groups. A structure theorem for groups with $d(G) \geq cn$, c any fixed positive constant, was obtained by [1] (see Remark 4.3 below), while related results were obtained by Berkovich in [2].

Although easy to define, the degree is difficult to compute. It is more-or-less obvious that $d(G)$ can be computed by examining all subsets of the subgroup lattice of G. The main result of this note is that in some cases a "greedy" algorithm is also available, that is an algorithm that proceeds by making locally optimal choices rather than directly searching for the global minimum. This is hardly of practical application (the subgroup lattice of a group may be exponentially larger than the group itself), but it has surprising consequences for the structure of a minimal permutation representation. We note that whenever a group G acts on a set A, the sizes of the orbits of the action determine a partition of $|A|$. Our main application is:

Theorem 1.1. *Let G be a finite nilpotent group with at most one central involution. Then all minimal permutation representations define the same partition of the integer $d(G)$.*

This is a special case of a more general result. The statement of the more general result uses another invariant of finite groups, their *dimension*. We define the dimension $\dim G$ of a finite group G to be the

Date: April 21, 2008.

maximal integer t such that G contains a direct product of t normal subgroups. Then we have the following theorem which has Theorem 1.1 as a special case:

Theorem 1.2. *Let G be a finite group. Then all minimal permutation representations have at most $\dim G$ orbits. If G is nilpotent, then there exists a minimal permutation representation with exactly $\dim G$ orbits, and any two such representations define the same partition of the integer $d(G)$. Under the same hypothesis, if G has at most one central involution then all minimal permutation representations have $\dim G$ orbits.*

In the statements of these two theorems it is not necessary to assume that G is nilpotent. In fact, one only needs to assume that G satisfies the conclusion of Lemma 3.12 below. It would be interesting to get more information about the partition defined by a minimal permutation representation with $\dim G$ orbits. The algorithm presented here was first obtained for p-groups by the first author under the supervision of the third author ([5]). Here the main motivation was to understand the distribution of $\Delta(G)$ in the interval $[0, 1]$, where $\Delta(G) = d(G)/|G|$. For example, it was easy to show that every number of the form $\frac{1}{n}$, n a natural number, is a limit point of $\Delta(G)$ as $|G|$ tends to infinity. Clearly, zero is also a limit point. We show here (see Theorem 4.7 below) that these are the only limit points.

While one can obtain various results about gaps, limit points, and averages, it seems to us that the problem of understanding the distribution of $\Delta(G)$ is a deep problem. For example, numerical investigations using the algorithm developed here show that for an integer n at most five, the average value of $\Delta(G)$ over groups of order p^n varies *polynomially on residue classes*, ("PORC") as a funtion of $\frac{1}{p}$ (see [6] for the terminology). Further, for a fixed polynomial $Q(x)$ and integer $n \leq 5$, the number of groups of order p^n with $\Delta(G) = Q(\frac{1}{p})$ seems to be PORC in p (c.f. [6]).

Acknowledgements. We would like to acknowledge conversations with John Conway, William Kantor, Avinoam Mann and James Wilson. Our interest in minimal permutation representations was triggered by a question raised by Andre Kornell in a Princeton undergraduate Algebra class. During the initial investigation, the third author was assisted by Evan Hass, another student in that class. We wish to thank Neil Saunders for pointing out an error in an earlier draft of this paper (see Remark 3.27 for more on this). We also thank the referee for a very careful reading of the manuscript, and making many suggestions which led to various improvements of the style and presentation of the paper.

The second author's research was supported in part by a Porter Ogden Jacobus Fellowship at Princeton University, in part by a Clay Mathematics Foundation Liftoff Fellowship, and in part by the National Science Foundation under agreement No. DMS-0111298. The third author's research was partially funded by a Young Investigator Grant from the NSA and by the NSF.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF and the NSA.

2. Definitions

Let G be a finite group. By a *permutation representation* of G on the set X we simply mean an action of G on X, in other words a homomorphism $G \to S_X$ where S_X is the symmetric group on X. The cardinality of X is denoted by $|X|$ (similarly for other sets) and called the *degree* of the representation. The

representation is called *faithful* if no non-identity element of G acts trivially, that is if the corresponding homomorphism is injective.

It is a classical theorem of Cayley that every group G has a faithful permutation representation, afforded by the regular action of the group on itself. Since we know that a group G of order n is isomorphic to a subgroup of S_n, it is natural to ask whether G can be embedded in some S_d for $d < n$. Accordingly we define the *degree* of G, denoted $d(G)$, to be the least integer d such that G can be embedded in S_d, i.e. such that G has a faithful permutation representation on a set with d elements. We will call a faithful permutation representation of degree $d(G)$ *minimal*. In this language, Cayley's theorem states that $d(G) \leq n$. We will also consider the *relative degree*

$$\Delta(G) = \frac{d(G)}{|G|}.$$

We call a permutation representation an *orbit* if it is transitive. A transitive permutation representation of G is equivalent to the usual action of G on the set G/H of (right) cosets of H in G, where H is the stabilizer in G of any element in the orbit. The set of vertex stabilizers in an orbit is a conjugacy class of subgroups of G, and conversely a transitive permutation representation is determined up to isomorphism by a conjugacy class of subgroups. Any permutation representation is a disjoint union of orbits.

Definition 2.1. The *core* $K(H)$ of a subgroup H of G is the maximal normal subgroup of G contained in H, i.e. $K(H) = \cap_{x \in G} H^x$. It is the kernel of the permutation representation of G on G/H. To a family $\mathcal{H}$ of subgroups of G we associate the core $K(\mathcal{H}) = \cap_{H \in \mathcal{H}} K(H)$, the kernel of the permutation representation $\rho_{\mathcal{H}}$ of G on $\amalg_{H \in \mathcal{H}} G/H$ (if $\mathcal{H}$ is empty set $K(\mathcal{H}) = G$). We call the collection $\mathcal{H}$ *faithful* if $K(\mathcal{H}) = \{1\}$, i.e. if the associated permutation representation is faithful.

For a faithful collection $\mathcal{H}$ of subgroups of G we set $\Delta(\mathcal{H}) = \sum_{H \in \mathcal{H}} \frac{1}{|H|}$. The discussion above shows that

$$\Delta(G) = \min \{\Delta(\mathcal{H}) \mid \mathcal{H} \text{ faithful}\}, \tag{2.1}$$

and we shall term a collection $\mathcal{H}$ such that $\Delta(\mathcal{H}) = \Delta(G)$ *minimal*. We observe that if H is a subgroup of G, and $x \in G$, then $K(H) = K(H^x)$. This observation implies that if $\mathcal{H}$ is a minimal collection, then no two elements of $\mathcal{H}$ can be equivalent under the conjugation action of G. Note that Equation (2.1) is, in fact, an algorithmic prescription for computing $d(G)$. In the next section we will show that much fewer collections need be considered to find a minimal one.

3. Determining $d(G)$

We now turn to the (algorithmic) problem of determining $d(G)$ *exactly*, given considerable information on the group in question. We have already remarked that knowing the subgroup lattice of G and, in addition, the order of each subgroup and whether it is normal in G is sufficient. Indeed, $d(G)$ can then be computed via Equation (2.1) in time exponential in the size of the input. In fact, this algorithm finds the minimal permutation representations themselves, not just their degrees:

Algorithm 3.1. *Test, in order, each collection $\mathcal{H}$ of subgroups of G. If $K(\mathcal{H}) = \{1\}$ evaluate $\Delta(\mathcal{H}) = \sum_{H\in\mathcal{H}} \frac{1}{|H|}$. Otherwise, ignore the collection. At the end output the smallest value seen and the collections which produced it.*

Note that the information we are giving ourselves (the subgroup lattice and in addition, for each subgroup, its order and whether it is normal) makes it is possible to compute the core $K(H)$ of a subgroup H in "reasonable time", at worst by scanning the entire lattice (here we use the fact that normal subgroups are marked). In the same way it is also possible to compute the intersection of any two subgroups.

In fact, at least for nilpotent groups, a minimal collection can be constructed in a time polynomial in the size of the input. Since determining the subgroup lattice is itself an intractable problem, the resulting algorithm is hardly practical. However, its analysis will expose considerable structure relating to a minimal permutation representation.

3.1. **The Socle.** For a general reference, we refer the reader to [9] or the more recent [4, Section 4.3]

Definition 3.2. Let G be a finite group. $\mathcal{M} = \mathcal{M}(G)$ will denote the set of minimal normal subgroups of G. $\mathcal{T} = \mathcal{T}(G)$ will denote the poset consisting of those subgroups of G generated by a subset of $\mathcal{M}$ ordered by inclusion. By definition, the subgroup generated by the empty set is $\{1\}$. $\mathcal{T}$ has a unique maximal element, the subgroup generated by all elements of $\mathcal{M}$. It is called the *socle* of G and will be denoted $M(G)$ or simple M.

Remark 3.3. Proposition 3.7(3) below shows $\mathcal{T}(G)$ is precisely the set of normal subgroups of G contained in M. In particular, it is a lattice.

Remark 3.4. A simple group is its own unique minimal normal subgroup. The discussion below then applies (rather trivially) to simple groups.

Since they are generated by a conjugation-invariant subset of G, the elements of $\mathcal{T}$ are normal subgroups. The main conclusion of this section is that $\mathcal{T}$ behaves essentially like the lattice of subspaces of a vector space.

Definition 3.5. The *dimension* of $T \in \mathcal{T}$, denoted $\dim_G T$, will be the cardinality of a minimal subset of $\mathcal{M}$ generating T. We define the *dimension of* G to be $\dim G = \dim_G M$ and let $\operatorname{codim}_G T = \dim G - \dim_G T$.

Lemma 3.6. *(quotients) Let G be a finite group and let T be a normal subgroup of G. Let $\bar{G} = G/T$ and set $\bar{K} = KT/T$ for any subgroup K of G. If $N \in \mathcal{M}(G)$ is not contained in T then $\bar{N} \in \mathcal{M}(\bar{G})$. Similarly, if $S \in \mathcal{T}(G)$ then $\bar{S} \in \mathcal{T}(\bar{G})$. Moreover, the corresponding map $\mathcal{T}(G) \to \mathcal{T}(\bar{G})$ is a map of posets, and it is injective on the subposet of normal subgroups of M containing T (if such subgroups exist).*

Proof. Suppose $N \in \mathcal{M}(G)$ is not contained in T. Then $N \cap T$ is a proper subgroup of N. Being a normal subgroup of G it is trivial. Thus $\bar{N} \simeq N$ is a non-trivial normal subgroup of $\bar{G}$. Suppose $\bar{N}$ is not minimal. Then there exists a non-trivial proper subgroup $L < N$ such that $\bar{L} \lhd \bar{G}$. Now let $x \in G$

and $l \in L$. By assumption we have $l^x = l't$ for some $l' \in L$ and $t \in T$. However, this means $t = l'^{-1}l^x$, i.e. $t \in N \cap T = \{1\}$. Thus L is a normal subgroup of G, a contradiction. Next suppose $S \in \mathcal{T}(G)$ is generated by $\{N_i\}_{i=1}^d \subset \mathcal{M}(G)$ for some $d \geq 1$. Then $\bar{S}$ is generated by $\left\{\bar{N}_i\right\}_{i=1}^d$, and each $\bar{N}_i$ is either trivial or belongs to $\mathcal{M}(\bar{G})$. Then the map $S \mapsto \bar{S}$ clearly preserves inclusions, and injectivity follows from the bijective correspondence between subgroups of G containing T and subgroups of $\bar{G}$. □

Proposition 3.7. *(Structure theory of $\mathcal{T}$ and M)*

(1) *Every $T \in \mathcal{T}$ is a direct product $\prod_{i=1}^k N_i$, where $N_i \in \mathcal{M}$.*
(2) *If $S \lneq T$ are elements of $\mathcal{T}$ then there exists $N \in \mathcal{M}$ with $N \cap S = \{1\}$ and $N < T$. In fact, there exists $U \in \mathcal{T}$ with $U \cap S = \{1\}$ and $SU = T$.*
(3) *We have*
$$\mathcal{T} = \{H \lhd G \mid H < M\}.$$
In particular, $\mathcal{T}$ is a lattice.
(4) *In (1), we have $k = \dim_G T$.*
(5) *Let $S, T \in \mathcal{T}$. Then $\dim_G ST = \dim_G S + \dim_G T - \dim_G S \cap T$.*
(6) *Let $S < T$ be elements of $\mathcal{T}$. Then $S = T$ if and only if $\dim_G S = \dim_G T$. If $\dim_G S \leq d \leq \dim G - 1$ we have for S fixed and T varying:*
$$\bigcap_{\substack{S < T \in \mathcal{T} \\ \dim_G T = d}} T = S.$$

Proof.

(1) Let T be generated by $\{N_i\}_{i=1}^k$, where none of the factors can be omitted. Then $T = \prod_{i=1}^k N_i$ since the N_i are normal, and the product is direct since each N_i is a minimal normal subgroup so that $N_i \cap \prod_{j \neq i} N_j$ is trivial for all i.
(2) Suppose $T = \prod_{i=1}^k N_i$. If all the N_i are contained in S then $S = T$. Otherwise, without loss of generality we may assume that $N_1 \cap S = \{1\}$ and then the product $S_1 = S \cdot N_1$ is direct. Continuing by induction, if $S_i \neq T$ we can find N_{i+1} such that $S_{i+1} = S_i \cdot N_{i+1}$ is a direct product. This process must terminate after some finite number r of steps, at which point we have $T = S \cdot U$ where $U = \prod_{i=1}^r N_i$.
(3) Certainly every element of $\mathcal{T}$ is a normal subgroup of G contained in M. Conversely, let H be such a subgroup, which we can assume to be non-trivial. We proceed by induction on the size of an element $T \in \mathcal{T}$ containing H. Bounding $|H|$ above by the size of such an element, we bound it below by the element $S \in \mathcal{T}$ generated by all minimal normal subgroups of G contained in H, a non-trivial group since H is non-trivial and hence must contain some minimal normal subgroup of G. Now if $S = H$ we are done, and otherwise both containments $S < H < T$ are proper. By part (2), we can write $T = S \cdot U$ for some $U \in \mathcal{T}$, and since S is non-trivial, U is a proper subgroup of T. Now H is a subgroup of the product SU which contains S and thus is of the form $S \cdot H_1$ for some subgroup H_1 of U. Moreover, H_1 is non-trivial, and is a normal subgroup of G since S and

H are. By induction, H_1 is generated by some elements of $\mathcal{M}$. By construction, S also has this property, and since $H = S \cdot H_1$ we are done.

(4) The case $\dim_G T = 0$ is clear, and we proceed by induction. Assume that $T = \prod_{i=1}^{d} N_i$ is of dimension d, but can also be written as the direct product $\prod_{j=1}^{k} N'_j$ with $k \geq d$. Passing to the quotient $\bar{G} = G/N_d$, we have $\bar{T} = \prod_{i=1}^{d-1} \bar{N}_i$, and in particular $\dim_{\bar{G}} \bar{T} \leq d-1$. Consider, now, the sequence of subgroups $T_m = \prod_{j=1}^{m} N'_j$ for $1 \leq m \leq k$ (and $T_0 = \{1\}$). Since $N_d < T_k = T$ we can let m_0 denote the smallest value of m such that $N_d < T_m$. We claim that $\bar{N}'_j \in \mathcal{M}(\bar{G})$ for $j \neq m_0$ and that the product $\bar{T} = \prod_{j \neq m_0} \bar{N}'_j$ is direct. We have thus written $\bar{T}$, which is of dimension at most $d-1$, as a strictly increasing product with $k-1$ factors. By the induction hypothesis we get $d - 1 \leq k - 1 = \dim_{\bar{G}} \bar{T} \leq d - 1$, and hence $d = k$.
To see the claim, note first that $\bar{T}_{m_0-1}$ is an isomorphic image of T_{m_0-1}, so that the product $\prod_{j<m_0} \bar{N}'_j$ is direct with non-trivial factors. Next, both T_{m_0-1} and N_d are contained in T_{m_0} by the choice of m_0, giving $T_{m_0-1}N_d \subset T_{m_0} = T_{m_0-1}N'_{m_0}$. Also, N_d is not contained in T_{m_0-1} (again by the choice of m_0) and neither is N'_{m_0} (since $T_{m_0} \neq T_{m_0-1}$). Thus their images in G/T_{m_0-1} are both minimal normal subgroups, which are contained in each other. Hence they are equal, i.e. $T_{m_0-1}N_d = T_{m_0}$, so that $\bar{T}_{m_0-1} = \bar{T}_{m_0}$. That the product $\bar{T}_{m_0-1} \cdot \prod_{j>m_0} \bar{N}'_j$ is direct with $\bar{N}'_j \in \mathcal{M}(\bar{G})$ now follows from the fact that $T_{m_0} \cdot \prod_{j>m_0} N'_j$ is a direct product with $N_d < T_{m_0}$.

(5) Suppose $S, T \in \mathcal{T}$. Set $d_S = dim_G S$ and $d_T = dim_G T$. $Write S \cap T = \prod_{i=1}^{d} N_i$ and extend to direct product representations $S = (S \cap T) \cdot \prod_{j=1}^{d_S-d} N_j^S$, $S = (S \cap T) \cdot \prod_{k=1}^{d_T-d} N_k^T$. Letting $\bar{G} = G/(S \cap T)$, both products $\bar{S} = \prod_{j=1}^{d_S-d} \bar{N}_j^S$ and $\bar{T} = \prod_{k=1}^{d_T-d} \bar{N}_k^T$ in $\mathcal{T}(\bar{G})$ are direct. If the product $\bar{S}\bar{T} = \prod_{j=1}^{d_S-d} \bar{N}_j^S \prod_{k=1}^{d_T-d} \bar{N}_k^T$ were not direct, we would have $\bar{S} \cap \bar{T} \neq \{1\}$, a contradiction. This means that the product

$$ST = \prod_{i=1}^{d} N_i \prod_{j=1}^{d_S-d} N_j^S \prod_{k=1}^{d_T-d} N_k^T$$

is direct.

(6) The case $S = \{1\}$ of the first claim is immediate, and the general case follows by writing $T = SU$. For the second assertion we may assume $d > \dim_G S$ and let S' be the intersection under consideration. Suppose $S' \neq S$. Then there must exist $N \in \mathcal{M}(G)$ contained in S' but not in S. We now choose $U \in \mathcal{T}'$ such that the product $M = S \cdot N \cdot U$ is direct and let $U' \in \mathcal{T}$ be a subgroup of U of dimension $d - \dim_G S$. Finally, we set $T = S \cdot U$. Then $\dim_G T = d$ but S' is not a subgroup of T since $N \cap T = \{1\}$.

□

We make again the observation, used in the proof of part (3) above, that a non-trivial normal subgroup of G must intersect the socle, since it must contain a minimal normal subgroup. This easy implication is a key ingredient in the algorithm described below. To decide whether a permutation representation is faithful amounts to deciding whether its kernel is trivial. The point is that one only has to keep track of

the intersection of the kernel with the socle, which is an element of a well-behaved lattice. This motivates the following definition:

Definition 3.8. For a subgroup $H < G$ let $H_M \in \mathcal{T}$ denote the subgroup of H generated by all elements of $\mathcal{M}$ contained in H. We call this subgroup the *relative core* of H. If $\mathcal{H}$ is a family of subgroups of G, we similarly define its *relative core* to be the subgroup $\mathcal{H}_M = \cap_{H\in\mathcal{H}} H_M \in \mathcal{T}(G)$ (c.f. Definition 2.1). We ascribe the relative core M to the empty collection.

Remark 3.9. Part (3) of Proposition 3.7 shows that $\mathcal{H}_M = K(\mathcal{H}) \cap M$.

Definition 3.10. To a subgroup H of G we associate the numbers $\dim_G H = \dim_G H_M$ and $\operatorname{codim}_G H = \operatorname{codim}_G H_M$. Note that if $H \in \mathcal{T}(G)$ then this definition is compatible with the previous one, and that we have $\dim_G G = \dim G$. We also declare the trivial group to have dimension 0.

For our purposes, the most important feature of a subgroup H will be its relative core H_M. One construction we will use is the extension of H by an element $T \in \mathcal{T}$ to form the subgroup $H \cdot T$. If G is nilpotent then it is easy to compute the relative core of HT in terms of H_M and T. The key feature of nilpotent groups is the following observation:

Lemma 3.11. *Suppose G is a nilpotent group. Then its socle $M(G)$ is contained in its center $Z(G)$, and every minimal normal subgroup is a cyclic group of prime order. In fact, $\mathcal{M}(G)$ consists precisely of the central cyclic subgroups of prime order.*

Proof. We start with the well-known fact that a nilpotent group is the direct product of its Sylow subgroups, which are normal and hence characteristic. On the other hand, if N is a minimal normal subgroups of a group G then N is *characteristically simple* (has no non-trivial characteristic subgroups) since characteristic subgroups of a normal subgroup are normal in the ambient group.

Thus, let N be a minimal normal subgroup of the non-trivial nilpotent finite group G. Then N itself is a non-trivial nilpotent group. Let P be a non-trivial Sylow subgroup of N. Then P is normal in G, hence equal to N. It follows that N is a p-group for some prime p, and hence contained in the (unique) p-Sylow subgroup of G, which we denote Q. Since N is a subgroup of the p-group Q it intersects the center $Z(Q)$, which is a normal subgroup of G (it is characteristic in the normal subgroup Q). It follows that N is contained in $Z(Q)$, which is central in G since Q is a direct factor of G.

Finally, since N is a central subgroup, every subgroup of N is normal in G. It follows that N has no non-trivial subgroups, and hence is a cyclic p-group. Conversely, it is clear that a central cyclic subgroup of prime order is a minimal normal subgroup. □

Lemma 3.12. *Suppose G is a nilpotent group. Let $H < G$, $T \in \mathcal{T}$. Then $(H \cdot T)_M = H_M T$.*

Proof. Writing $T = (T \cap H_M) \cdot S$ with $S \cap H_M = \{1\}$ we have $HT = HS$ and $H_M T = H_M S$ so we may assume $H \cap T = \{1\}$. Clearly $H_M T \subset (H \cdot T)_M$. Conversely, let $N < HT$ be a minimal normal subgroup of G. If $N < T$ there is nothing to prove, so we may assume $T \cap N = \{1\}$. Since H and T are disjoint, every $n \in N$ can be uniquely written in the form $n = h_n t_n$ for some $h_n \in H$ and $t_n \in T$. Note that the map $n \mapsto h_n$ is a group homomorphism (it is the restriction to N of the quotient map $HT/T \simeq H$), and since N and T are disjoint it is an isomorphism onto its image N'.

Since N and T are central subgroups (here we use the nilpotence of G via the previous Lemma), it follows that N' is a central subgroup as well, and since N was a cyclic group of prime order so is N'. It follows that N' is a minimal normal subgroup of G, contained in H. We conclude that $N \subset N'T \subset H_M T$. □

Remark 3.13. There are non-nilpotent groups for which this lemma fails (c.f. Remark 3.27 below).

3.2. **Minimal faithful collections and codimension one subgroups.** Throughout this section G is any finite group satisfying the conclusion to Lemma 3.12. Recall that any finite nilpotent group has this property. We are interested in constructing a minimal faithful collection of subgroups of G, and a natural way to do so is step-by-step, incrementally adding subgroups to our collection until it is faithful. Rather than keeping track of $K(\mathcal{H})$, we note that $\mathcal{H}_M$ carries sufficient information to decide whether $K(\mathcal{H})$ is trivial. Moreover, while the cores $K(\mathcal{H})$ decrease through the lattice of all normal subgroups of G, the relative cores $\mathcal{H}_M$ decrease through the lattice $\mathcal{T}(G)$ which is much easier to work with.

We now turn to the "minimality" property of a collection, which appears to push in the opposite direction to "faithfulness". The first favors selecting large subgroups, and having few of them. The second seems to suggest choosing small subgroups, or else many large ones will be needed. The multiplicative property of orders of subgroups actually implies that choosing many large subgroups is the right way (in fact, usually a necessary approach). The analysis is very similar to that of Johnson [7]. In both cases it is shown that the elements of a minimal faithful collection may be (and in some cases, must be) drawn from a particular class of subgroups, using the same trick. However, the class of subgroup we employ seems more useful in practice. The reader should compare our next result with [7, Lemma 1]

Lemma 3.14. *("replacement lemma") Let $H < G$ be of codimension at least 2. Then there exist subgroups H^1 and H^2 of G containing H such that $H^1_M \cap H^2_M = H_M$ and $\frac{1}{|H_1|} + \frac{1}{|H_2|} \leq \frac{1}{|H|}$. Moreover, this inequality is strict unless G contains at least two central involutions.*

Proof. By assumption, there exists a subgroup $T \in \mathcal{T}$ of codimension 1 such that H_M is a proper subgroup of T. By Proposition 3.7 we choose $N_1, N_2 \in \mathcal{M}$ with $N_1 \subset T$ such that the inclusions $H_M < H_M N_1$ and $T < TN_2$ are proper, and set $H^1 = H \cdot N_1$, $H^2 = H \cdot N_2$ (semi-direct products as the N_i are minimal normal subgroups). Now let $x \in H^1_M \cap H^2_M$. By Lemma 3.12 we can write $x = h_1 n_1 = h_2 n_2$ with $h_1, h_2 \in H_M$, $n_i \in N_i$. This means $n_2 = h_2^{-1} h_1 n_1 \in H_M N_1 < T$, i.e. $n_2 \in N_2 \cap T = \{1\}$, and hence $x = h_2 \in H_M$. Finally, since H is a proper subgroup of both H^1, H^2 its index in both subgroups is at least 2, and we have

$$\frac{1}{|H^1|} + \frac{1}{|H^2|} \leq \left(\frac{1}{2} + \frac{1}{2}\right) \frac{1}{|H|} = \frac{1}{|H|}.$$

Equality can only happen if both N_1 and N_2 are of order 2, in which case the non-trivial elements of N_i are both central involutions. □

Definition 3.15. Let $\mathcal{A} = \mathcal{A}(G)$ denote the set of subgroups of G of codimension 1.

The reader should compare the next theorem with [7, Cor. 1].

Theorem 3.16. *There exist minimal faithful collections contained in $\mathcal{A}$, and these are the ones of maximal size. If G has at most one central involution then every minimal faithful collection is contained in $\mathcal{A}$.*

Proof. Let $\mathcal{H}$ be a faithful collection, and let $H \in \mathcal{H} \setminus \mathcal{A}$. If H is of codimension 0 (i.e. $H_M = M$) we have

$$\{1\} = \mathcal{H}_M = (\mathcal{H} \setminus \{H\})_M \cap H_M = (\mathcal{H} \setminus \{H\})_M .$$

In particular, $\mathcal{H} \setminus \{H\}$ is also faithful. Otherwise, let H^1, H^2 be the subgroups constructed in Lemma 3.14, and let $\mathcal{H}' = (\mathcal{H} \setminus \{H\}) \cup \{H^1, H^2\}$. By construction we have $\mathcal{H}'_M = \mathcal{H}_M = \{1\}$ so that $\mathcal{H}'$ is faithful. In addition, Lemma 3.14 yields $\Delta(\mathcal{H}') \leq \Delta(\mathcal{H})$, and this inequality is strict if G has at most one central involution. In general we note that $\mathcal{H}'$ has more elements than $\mathcal{H}$. In particular, a minimal faithful collection of maximal size must consist of codimension-one subgroups. □

Definition 3.17. A collection $\mathcal{H} \subset \mathcal{A}$ is said to be *independent* if its relative core is strictly contained in that of any proper sub-collection $\mathcal{H}' \subsetneq \mathcal{H}$. The empty collection is also assumed to be independent.

A minimal faithful collection $\mathcal{H} \subset \mathcal{A}$ is certainly independent – otherwise it would have a faithful proper sub-collection.

Proposition 3.18. *The set of independent subsets of $\mathcal{A}$ forms a matroid:*

1. *A subcollection of an independent collection is independent.*
2. *$\mathcal{H} \subset A$ is independent if and only if $\operatorname{codim}_G \mathcal{H}_M = |\mathcal{H}|$.*
3. *If $\mathcal{H}$, $\mathcal{H}'$ are independent collections with $|\mathcal{H}'| > |\mathcal{H}|$ then there exists $H' \in \mathcal{H}'$ such that $\mathcal{H} \cup \{H'\}$ is independent.*

Proof.

(1) Let $\mathcal{H} \subset \mathcal{A}$ be independent, and suppose $\mathcal{H}''$ us a proper subcollection of $\mathcal{H}' \subset \mathcal{H}$ such that $\mathcal{H}''_M = \mathcal{H}'_M$. Letting $\bar{\mathcal{H}} = \mathcal{H} \setminus \mathcal{H}'$, we have

$$\left(\bar{\mathcal{H}} \cup \mathcal{H}''\right)_M = \bar{\mathcal{H}}_M \cap \mathcal{H}''_M = \bar{\mathcal{H}}_M \cap \mathcal{H}'_M = \mathcal{H}_M,$$

contradicting the independence of $\mathcal{H}$.

(2) Let $S, T \in \mathcal{T}(G)$ with $\operatorname{codim}_G T = 1$. Then ST either equals T or M, and we have $\dim_G S \cap T = \dim_G S$ or $\dim_G S - 1$, respectively, by the inclusion-exclusion formula of Lemma 3.7(5). By induction on the size of any collection $\mathcal{H} = \{H^i\}_{i=1}^k \subset \mathcal{A}$ we see that $\operatorname{codim}_G \mathcal{H}_M \leq |\mathcal{H}|$, with equality if and only if the sequence of intersections $\cap_{i=1}^m H^i_M$ is strictly decreasing with m, $1 \leq m \leq k$.

(3) We have $\dim_G \mathcal{H}'_M < \dim_G \mathcal{H}_M$, and hence $\mathcal{H}'_M$ does not contain $\mathcal{H}_M$. It follows that we can find $H' \in \mathcal{H}'$ such that H'_M does not contain $\mathcal{H}_M$. Then $\dim_G(\mathcal{H}_M \cap H'_M) = \dim_G \mathcal{H}_M - 1$ (equality is not possible by the choice of H'). By part (2) we see that that $\mathcal{H} \cup \{H'\}$ is independent.

□

Corollary 3.19. *Let $\mathcal{H} \subset A$ be independent. Then the following are equivalent:*

1. *$|\mathcal{H}| = \dim G$;*
2. *$\mathcal{H}$ is faithful;*

(3) *$\mathcal{H}$ is a maximal independent subset of $\mathcal{A}$. Here,* maximal *means maximal with respect to inclusion.*

Proof. The equivalence of (1) and (2) is contained in part (2) of Proposition 3.18. An independent collection with $\mathcal{H}_M = \{1\}$ is certainly maximal. An independent collection with $\mathcal{H}_M \neq \{1\}$ is not maximal since in that case there exists some $T \in \mathcal{T}$ of codimension 1 which does not contain $\mathcal{H}_M$, and we can add it to $\mathcal{H}$ to form a larger independent collection. □

Corollary 3.20. *A subset $\mathcal{H} \subset \mathcal{A}$ is a minimal faithful collection if and only if it is independent and maximizes*

$$w(\mathcal{H}) = \sum_{H \in \mathcal{H}} \left(2 - \frac{1}{|H|}\right)$$

among the independent subsets.

Proof. We have already noted that a minimal faithful collection contained in $\mathcal{A}$ is independent and maximal (with respect to inclusion), and that a maximal (with respect to inclusion) independent set is a faithful collection. It is clear that a subset maximizing this weight function is maximal independent, since $2 - \frac{1}{|H|} > 0$ for all subgroups H. Finally, we note that a maximal independent set H satisfies

$$w(\mathcal{H}) = 2 \dim G - \Delta(\mathcal{H}).$$

□

Corollary 3.21. *There exist minimal faithful collections of cardinality $\dim G$. If G has more than one central involution, there may also exist minimal faithful collections of smaller cardinality.*

Proof. We have seen that there exist minimal faithful collections contained in $\mathcal{A}$, that these are independent sets, and that every independent set has $\dim G$ elements. □

Example 3.22. Let G be a p-group for a prime p, and let $Z = Z(G)$ be its center. It is well-known (and follows from the class formula) that every normal subgroup of G intersects the center non-trivially. Since every subgroup of the center is normal, it follows that $\mathcal{M}(G) = \mathcal{M}(Z)$, and in particular $\dim G = \dim Z(G)$. This observation recovers [7, Thm. 3]:

Theorem 3.23. *Let G be a p-group with center Z. Then there exists a minimal faithful collection for G of cardinality $\dim Z$. If p is odd this holds for all minimal faithful collections.*

3.3. **The Algorithm.** Again we assume that G is a finite group satisfying the conclusion of Lemma 3.12. We have reduced the problem of finding a minimal faithful collection to maximizing an additive weight function on a matroid. This is a problem which is solvable by a greedy algorithm. Before we give the algorithm we record a simplifying Lemma:

Lemma 3.24. *Let $\mathcal{H} \subset \mathcal{A}$ be independent, and suppose $H' < G$ has the largest cardinality possible such that H'_M does not contain $\mathcal{H}_M$. Then $H' \in \mathcal{A}$, $\mathcal{H} \cup \{H'\}$ is independent, and H' maximizes the function $w(H) = 2 - \frac{1}{|H|}$ among all $H \in \mathcal{A}$ such that $\mathcal{H} \cup \{H\}$ is independent.*

Proof. By Lemma 3.7(6) we can find $T \in \mathcal{T}$ of codimension 1 containing H'_M but not containing $\mathcal{H}_M$. Setting $H = H'T$ we have $H_M = H'_M T = T$, which does not contain $\mathcal{H}_M$. By the maximality of H' we have $H = H'$ implying $H'_M = T$, so that H' is of codimension 1 and $\mathcal{H} \cup \{H'\}$ is independent. Finally H' was chosen to maximize $w(H)$ in an even larger family than needed. □

Algorithm 3.25. *("Greedy Algorithm") Let G be a non-trivial group. We assume we are given the subgroup lattice of G, and that normal subgroups are marked as such.*

(1) *Fine the socle M of G.*
(2) *Initialize $\mathcal{H} = \emptyset$, $T = M$, $\Delta = 0$.*
(3) *Repeat until $T = \{1\}$*
   (a) *Find a subgroup H of maximal cardinality not containing T.*
   (b) *Add H to $\mathcal{H}$, $\frac{1}{|H|}$ to Δ.*
   (c) *Set $T = T \cap K(H)$.*
(4) *Output Δ, $\mathcal{H}$.*

Theorem 3.26. *The algorithm will repeat step (3) exactly $\dim G$ times, after which $\mathcal{H}$ will contain a minimal faithful collection of size $\dim G$ and Δ will equal $\Delta(G)$.*

Proof. From Lemma 3.24 it is clear that the independence of $\mathcal{H}$ and the equality $T = \mathcal{H}_M$ are invariants of the loop, and that $\dim T$ decreases by 1 after each iteration. In particular the loop terminates after exactly $\dim G$ steps.

We show by induction that after k iterations of step (3), $\sum_{H\in\mathcal{H}} \frac{1}{|H|}$ is minimal among independent collections of size k. This is certainly the case before the loop begins. Thus let $\mathcal{H}$ be the set at the *beginning* of the kth iteration (hence of size $k-1$), and let H_k be the subgroup chosen at that iteration. Suppose there is an independent collection $\mathcal{H}' \subset \mathcal{A}$ of size k such that $\sum_{H'\in\mathcal{H}'} \frac{1}{|H'|} < \frac{1}{|H_k|} + \sum_{H\in\mathcal{H}} \frac{1}{|H|}$. We may then write $\mathcal{H}' = \mathcal{H}'' \cup \{H'_k\}$ where H'_k is a member of minimal cardinality. By the inductive hypothesis, $\sum_{H\in\mathcal{H}} \frac{1}{|H|} \leq \sum_{H'\in\mathcal{H}''} \frac{1}{|H'|}$, and hence we must have $|H_k| < |H'_k|$. By the choice of H'_k, we actually have $|H_k| < |H'|$ for all $H' \in \mathcal{H}'$. We now use the matroid property of the independent subcollections of $\mathcal{A}$ shown in Proposition 3.18(3): since $\mathcal{H}'$ is of size k, while $\mathcal{H}$ is of size $k-1$, there exists some $H' \in \mathcal{H}'$ such that $\mathcal{H} \cup \{H'\}$ is independent. In particular this implies that $(\mathcal{H} \cup \{H'\})_M$ is strictly contained in $\mathcal{H}_M$, and as $|H'| > |H_k|$ we have a contradiction to the existence of $\mathcal{H}'$. □

Remark 3.27. In a preliminary version of this paper, Lemma 3.12 was stated for an arbitrary finite group G. That the lemma fails in general was pointed out by Neil Saunders [10] who also constructed an interesting counterexample. Saunders' counterexample can be described in the following fashion. The cyclic group C_7 has two non-isomorphic three-dimensional irreducible representations over $\mathbb{F}_2$. Let V be the direct sum of these two representations V_1, V_2 and let G be the semidirect product of C_7 and V. The minimal normal subgroups of G are V_1, V_2; the socle is thus V.

Let $W \subset V$ be any codimension-1 subspace (in the sense of linear algebra) not containing either of the V_i. Then W is a core-free subgroup of G of index 14; clearly there do not exist smaller faithful transitive G-sets. To see that this representation is minimal we need to rule out representations with two orbits,

which here must be the ones produced by our algorithm (since the only possible relative cores are V, V_i and the empty set). Since the largest subgroup not containing V_i is $C_7 \ltimes V_{3-i}$ which is of index 8, our algorithm produces a non-minimal permutation representation of degree 16.

To construct V explicitly note that the cyclotomic polynomial $p(x) = (x^7-1)/(x-1)$ decomposes over $\mathbb{F}_2$ as the product of two co-prime irreducible cubic polynomials p_1, p_2. Take any matrix $A \in M_6(\mathbb{F}_2)$ with characteristic polynomial p. Then A is of order 7 (the eigenvalues are roots of unity of order 7) and hence defines a representation of C_7 on $V \simeq \mathbb{F}_2^6$. By Galois invariance the subspaces $V_i(\overline{\mathbb{F}}_2)$ of $V(\overline{\mathbb{F}}_2)$ spanned by eigenvectors of A with eigenvalues which are roots of the p_i are defined over $\mathbb{F}_2$; they are obviously A-invariant. The resulting representations are not isomorphic since the eigenvalues of A are distinct.

4. Applications

4.1. **The proof of Theorem 1.2.** Let G be a finite group satisfying Lemma 3.12. We show here that all minimal permutation representations with $\dim G$ orbits have the same (multi-)set of orbit sizes.

It is an easy corollary of the proof of Theorem 3.26 that a unique set of orbit sizes is associated to all the minimal permutation representations that may be constructed by Algorithm 3.25.

Conversely, we show that every minimal permutation representation with $\dim G$ orbits may be constructed by the algorithm.

Proposition 4.1. *Let $\mathcal{H} = \{H^i\}_{i=1}^{\dim G}$ be a minimal faithful collection, w.l.g. ordered such that*

$$\left|H^1\right| \geq \left|H^2\right| \geq \cdots \geq \left|H^{\dim G}\right|.$$

Then each H^k has maximal cardinality among all subgroups H' of G such that $\left(\{H^i\}_{i=1}^{k-1} \cup \{H\}\right)_M$ is a proper subgroup of $\left(\{H^i\}_{i=1}^{k-1}\right)_M$.

Proof. By induction, it suffices to check that if a subgroup $H' < G$ is independent of $\{H^i\}_{i=1}^{k-1}$ then there exists $l \geq k$ such that $\mathcal{H} \cup \{H'\} \setminus \{H^l\}$ is independent. For this we set $S_j = \cap_{i=1}^{j} H^i_M$. It is then easy to see that we may take l to be the first j such that $H'_M \cap S_j = S_j$. □

4.2. **Accumulation points of $\Delta(G)$.** Let $n, p \in \mathbb{N}$ with $p > n$ a prime. Then $\Delta(C_n \times C_p) = \frac{1}{n} + \frac{\Delta(C_n)}{p} = \frac{1}{n} + O(\frac{1}{p})$. In particular, $\lim_{p\to\infty} \Delta(C_n \times C_p) = \frac{1}{n}$. This means that for each positive integer n, the point $\frac{1}{n}$ is an accumulation point of the set $\{\Delta(G)$;G finite group$\}$ in the interval $[0, 1]$. In fact, in Theorem 4.7 below we show that these points are the only non-zero accumulation points. We begin with some preliminary lemmas.

Lemma 4.2. *Let $H < G$ be a subgroup. Then $d(H) \leq d(G)$ and $\Delta(G) \leq \Delta(H)$.*

Proof. The first claim is obvious. For the second, let $\mathcal{H}'$ be a faithful collection of subgroups of H and note that $\Delta(\mathcal{H})$ is independent of the ambient group. Then $K_G(H_i) \subset K_H(H_i)$ (larger intersection). In particular, $K_G(\mathcal{H}) = \{1\}$. Choosing $\mathcal{H}$ minimal for H we deduce that $\Delta(G) \leq \Delta(\mathcal{H}) = \Delta(H)$. □

Remark 4.3. A cyclic p-group has relative degree 1. In particular, if $P < G$ is a cyclic p-group then

$$\Delta(G) \geq \frac{d(P)}{|G|} = \frac{1}{[G:P]}.$$

Conversely, Babai-Goodman-Pyber [1] give an explicit function $f\colon [0,1] \to \mathbb{R}$ such that if $\Delta(G) \geq \Delta$ then G has a cyclic p-subgroup of index at most $f(\Delta)$. In other words, as $|G|$ grows with $\Delta(G) \geq \Delta$, the degree of G is controlled (up to bounded multiplicative error) by the size of the largest cyclic p-subgroup of G. Specifically, they show that when G does not possess a large cyclic group of prime-power order it has a pair of reasonably large subgroups with trivial intersection.

Note that the above bound on $\Delta(G)$ is derived from a faithful collection of size 2. In Lemma 4.4 we show that when $\Delta(G) \geq \Delta$ there exists k depending only on Δ such that a minimal permutation representation of G has at most k orbits. The case of groups of prime exponent and nilpotence class two, studied in [1, Thm. 3.6] as well as [8] shows that we need $k > 2$ in general.

Lemma 4.4. *Let $k = \dim G$. Then $\Delta(G) \leq \frac{k}{2^{k-1}}$.*

Proof. Let $M = M(G)$ be the socle of G and write M as the direct product of k minimal normal subgroups $\{S_i\}_{i=1}^k$. For $1 \leq i \leq k$ let $H_i = \prod_{j \neq i} S_j$. It is clear that $\{H_i\}$ is a faithful collection of size k and each of its elements has size at least 2^{k-1}. □

Lemma 4.5. *Let P be a cyclic p-subgroup of G. Then $P_M < M(P)$ ($M(P)$ is the socle of P). If $|G|$ is large enough compared to $[G:P]$ then equality holds.*

Proof. Let $N < P$ be non-trivial and normal in G. Then $M(P)$ is a characteristic subgroup of N, so P_M is either trivial, or equal to $M(P)$. In any case, we have $\dim_G P \leq 1$.

Finally, the core of P has index at most $([G:P])!$ (it is the kernel of a homomorphism into $S_{[G:P]}$). If $|G| > ([G:P])!$ then $K_G(P)$ is a non-trivial normal subgroup of G contained in P, hence $M(P)$ is normal in G and thus $P_M = M(P)$. □

In fact, if G has a large cyclic p-subgroup then a permutation representation with two orbits is almost optimal:

Corollary 4.6. *Let P be a cyclic p-subgroup of G, and let $l(G)$ be the order of the smallest point stabilizer in an orbit in a minimal permutation representation of G. Then*

$$\frac{1}{l(G)} \leq \Delta(G) \leq \frac{1}{l(G)} + \frac{1}{|P|}.$$

Proof. Let $\mathcal{H}$ be a minimal faithful collection for G, chosen so that it contains an element H_1 of smallest possible order (denoted above by $l(G)$). Clearly $\Delta(G) = \Delta(\mathcal{H}) \geq \frac{1}{l(G)}$. For the other assertion, we may as well assume $M(P) \in \mathcal{M}(G)$, otherwise $K_G(P) = \{1\}$ and the claim is clear. Then $\mathcal{H}$, being faithful, must contain an element H_2 disjoint from $M(P)$, hence $\{P, H2\}$ is a faithful collection. □

Theorem 4.7. *Let G_n be a sequence of groups with orders increasing to infinity such that $\lim_{n\to\infty} \Delta(G_n) > 0$. Then this limit is of the form $1/l$ for some $l \in \mathbb{N}$.*

Proof. For n large enough we have $\Delta(G_n) > \Delta > 0$. The main result of [1], already quoted above, is that G_n has a cyclic p_n-subgroup P_n of index at most $f(\Delta)$ for some $f\colon [0,1] \to \mathbb{N}$. It follows that

$$\left|\Delta(G_n) - \frac{1}{l(G_n)}\right| \leq \frac{f(\Delta)}{|G_n|}.$$

Here $l(G_n)$ is as in the statement of Corollary 4.6. As $|G_n| \to \infty$, we see that $\frac{1}{l(G_n)}$ tends to a positive limit. The sequence of integers $l(G_n)$ must then be eventually constant. □

Note that we have shown more, that if $\Delta(G) \geq \Delta > 0$ then any minimal permutation representation consists of one large orbit of size essentially $|G|\,\Delta(G)$, and several other orbits of size and number bounded in terms of Δ. Indeed, the number of orbits is bounded by Lemma 4.4. We have an obvious bound $l(G) \leq (\Delta(G) - f(\Delta)/\,|G|)^{-1}$. Next, as soon as $|G|$ is large enough so that $\frac{1}{l(G)+1} + \frac{f(\Delta)}{|G|} < \frac{1}{l(G)}$, the subgroups H_1, H_2 of Lemma 4.5 must have the same cardinality. We conclude that if $\Delta(G) > \Delta$ and $|G|$ is large enough (depending on Δ), G has a cyclic p-subgroup P of index at most $f(\Delta)$ such that $M(P)$ is normal in G and a subgroup H of order $l(G)$ belonging to a minimal faithful collection and disjoint from $M(P)$. Then every other member of that minimal faithful collection may be replaced with P keeping the collection faithful. Hence all other orbits in the representation must have size at most $f(\Delta)$.

4.3. **Some numerical results.** The thesis [5] contains an implementation of Algorithm 3.25 in the algebraic programming language MAGMA [11]. Using the limited computing power of a personal computer, p-groups of order p^n with $n \leq 6$ with small p were examined. Any such group can be found in the MAGMA database. Let us summarize the findings.

There is only one group G of order p, and for this group $\Delta(G) = 1$. There are two groups of order p^2, namely $\mathbb{Z}_p \times \mathbb{Z}_p$ and $\mathbb{Z}_{p^2}$. Here $\Delta(\mathbb{Z}_p \times \mathbb{Z}_p) = \frac{2}{p}$ and $\Delta(\mathbb{Z}_{p^2}) = 1$. Consequently $\sum_{|G|=p^2} \Delta(G) = 1 + \frac{2}{p}$. There are five groups of order p^3: one cyclic with $\Delta = 1$; one elementary abelian with $\Delta = \frac{3}{p^2}$; one abelian with a generator of order p^2, having $\Delta = \frac{1}{p} + \frac{1}{p^2}$; and two non-abelian groups both having $\Delta = \frac{1}{p}$. Observe that $\sum_{|G|=p^3} \Delta(G) = 1 + \frac{3}{p} + \frac{4}{p^2}$. For groups of order p^4 and p^5 we state the following conjecture:

Conjecture 4.1. *For $p > 3$*

$$\sum_{|G|=p^4} \Delta(G) = 1 + \frac{5}{p} + \frac{11}{p^2} + \frac{9}{p^3},$$

$$\sum_{|G|=p^5} \Delta(G) = 1 + \frac{7}{p} + \frac{34 + 2\gcd(p-1,3) + \gcd(p-1,4)}{p^2} + \frac{54}{p^3} + \frac{24}{p^4}.$$

For any prime $p \geq 3$, there are exactly fifteen groups of order p^4, and these can be enumerated and described. So the proof of the first part of the conjecture should be straightforward. We have computationally verified the conjecture for groups of order p^4 for every prime p in the range $3 < p < 50$ and several larger values of p (≈ 1000). We considered the groups of order p^5 for $p \leq 19$. Note that the number of groups of order p^5 is $61 + 2p + 2\gcd(p-1,3) + \gcd(p-1,4)$. For groups of order p^6, we did not have enough data points to be able to guess a formula.

REFERENCES

1. László Babai, Albert J. Goodman, and László Pyber, *On faithful permutation representations of small degree*, Comm. Algebra **21** (1993), no. 5, 1587–1602. MR MR1213975 (94c:20004)
2. Yakov Berkovich, *The degree and index of a finite group*, J. Algebra **214** (1999), no. 2, 740–761. MR MR1680536 (2000c:20040)
3. Arthur Cayley, *Desiderate and suggestions: No. 1. the theory of groups*, Amer. J. Math. **1** (1878), no. 1, 50–52.
4. John D. Dixon and Brian Mortimer, *Permutation groups*, Graduate Texts in Mathematics, vol. 163, Springer-Verlag, New York, 1996. MR MR1409812 (98m:20003)
5. Benjamin Elias, *Minimally faithful group actions and p-groups*, Princeton University Senior Thesis, 2005.
6. Graham Higman, *Enumerating p-groups. II. Problems whose solution is PORC*, Proc. London Math. Soc. (3) **10** (1960), 566–582. MR MR0123605 (23 #A930)
7. D. L. Johnson, *Minimal permutation representations of finite groups*, Amer. J. Math. **93** (1971), 857–866. MR MR0316540 (47 #5087)
8. Peter M. Neumann, *Some algorithms for computing with finite permutation groups*, Proceedings of groups—St. Andrews 1985 (Cambridge), London Math. Soc. Lecture Note Ser., vol. 121, Cambridge Univ. Press, 1986, pp. 59–92. MR MR896501 (89b:20004)
9. Robert Remak, *Über minimale invariante untergruppen in der theorie der endlichen gruppen*, J. Reine Angew. Math. **162** (1930), 1–16.
10. Neil Saunders, *Private communication*, 09/01/2007.
11. MAGMA Computational Algebra System, *http://magma.maths.usyd.edu.au/*.

BEN ELIAS, COLUMBIA UNIVERSITY DEPARTMENT OF MATHEMATICS, NEW YORK, NY 10027
E-mail address: belias@math.columbia.edu

LIOR SILBERMAN, HARVARD UNIVERSITY DEPARTMENT OF MATHEMATICS, CAMBRIDGE, MA 02138
E-mail address: lior@math.harvard.edu

RAMIN TAKLOO-BIGHASH, DEPARTMENT OF MATH, STAT, AND COMP SCI, UNIVERSITY OF ILLINOIS AT CHICAGO, CHICAGO, IL 60607
E-mail address: rtakloo@math.uic.edu