8

Algorithms for Sylow p-Subgroups and Solvable Groups

WILLIAM M. KANTOR

University of Oregon
Eugene, Oregon

1. INTRODUCTION

In the midst of all of the more practical papers presented at the conference, this one is injected in order to describe a more theoretical framework. Instead of having usability as our criterion for efficiency, we will employ polynomial time. This places an entirely different emphasis on what can or cannot be accomplished (as explained at the end of this section and in Section 3). On the other hand, this new emphasis occasionally leads to new methods that may have nontheoretical applicability (cf. Section 5).

Throughout the paper we will consider a subgroup $G = \langle \Gamma \rangle$ of S_n generated by a subset Γ which may be assumed to be "small" (say, of size $< n^2$; cf. (2.4iii)). After indicating some of the properties of G that can be obtained in polynomial time by using Sims' results or related algorithms (Section 2), we will turn to the more recent results due to Luks [L2], Rónyai [R], or the author [K3,K4] (Section 4). The remainder of the paper consists of examples of some of the methods in [K3,K4], specialized to solvable groups (Section 5)—including an unpublished result on Sylow normalizers—and the Replacement theorem of [K2] that switches from one primitive permutation representation to another one when G is simple or nearly simple (Section 6). We conclude with miscellaneous remarks (Section 7), including similar types of questions concerning polynomial-time Galois theory.

Unlike the situation in many of the talks at this conference (or in CAYLEY [C]), there are no probabilistic aspects of the algorithms presented here: no probabilistic algorithms are known in this area that provably run in polynomial time and are faster than any of the algorithms we discuss—although such algorithms would certainly be of interest. A more serious restriction is our avoidance of backtrack algorithms, as they are usually exponential. On the other hand, in the present context there is no difficulty with the use of the action of G on sets of

polynomial size, such as the set of all 2-element subsets of our n-set—whereas for practical purposes it may be undesirable to deal with such a "large" set (i.e., of size $O(n^2)$).

2. FUNDAMENTAL ALGORITHMS

Since we will be using recursion for subgroups of S_n, the following simple result is very useful.

Lemma 2.1

If $1 < H_1 < H_2 < \cdots < H_m \leq S_n$, then $m \leq n \log_2 n < n^2$.

Proof:

By Lagrange's theorem, $2^i \leq |H_i| \leq n!$ for each i.

Babai [B] has shown that, in fact, $m < 2n$.

Let $G = \langle \Gamma \rangle \leq S_n$, acting on the n-set $X = \{1,2,\ldots,n\}$. We begin with two simple results that give an indication of the meaning of "polynomial time." Note that in each situation we need an algorithm which, for <u>any</u> G, produces the desired information. We emphasize that <u>all</u> subgroups of S_n we mention are assumed to be specified by means of generating sets of permutations.

Proposition 2.2

In polynomial time all orbits of G can be found.

Proof:

Form the graph with vertex set X and edges $\{x,xg\}$ for $g \in \Gamma$ and $xg \neq x \in X$. This graph can be determined in time $O(n|\Gamma|)$. Its connected components are just the orbits of G.

Proposition 2.3 [A]

Assume that G is transitive on X.

(i) In polynomial time all minimal blocks of imprimitivity (of size > 1) can be determined.

(ii) In polynomial time a block system Σ of size > 1 can be found such that G^Σ is primitive.

Proof:

(i) Find the orbits of G in its natural action on X^2 (namely, $g : (x,y) \longmapsto (xg,yg)$). Each orbit other than the diagonal $\{(x,x) \mid x \in X\}$ determines a graph with vertex set X and edges $\{x,y\}$ for (x,y) or (y,x) in the orbit. The set of connected components of such a graph is a block system Σ, and each minimal block system arises in this manner.

(ii) Iterate (i), replacing X by Σ if $|\Sigma| < n$. (The number of iterations is $< \log_2 n$.)

Note that it is not possible (in polynomial time!) to find all block systems (consider the regular representation of an elementary abelian 2-group).

The main result concerning G is Sims' algorithm:

Theorem 2.4 [S1,S2,FHL]

In polynomial time the following can all be determined:

(i) $|G|$

(ii) For $i = 1,\ldots,n$, a set Δ_i such that $G_{12\ldots i} = \langle \Delta_i \rangle$ and $|\Delta_i| < n^2$

(iii) A set Γ' such that $G = \langle \Gamma' \rangle$ and $|\Gamma'| < n^2$.

Moreover, if G also acts on a second set X' of polynomially-bounded size, then the kernel of the action of G on X' can also be found in polynomial time.

The original algorithm used by Sims [S1,S2] finds a base and strong generating set, as discussed at other talks at this conference. It was modified in [FHL] so as to be visibly polynomial-time. For $|\Gamma| = O(n^2)$ the algorithm in [FHL] runs in time $O(n^6)$; but, as pointed out by Babai [B] and Finkelstein [F], the original algorithm of Sims can be written so that it runs in time $O(n^5)$ (also see [J]). The last statement in (2.4) is essentially a special case of (ii). It has the effect of allowing several permutation representations to be dealt with simultaneously (for an example, see (4.1)).

The fact that $O(n^5)$ is presently the best time available for (2.4) influences all estimates of running times of later algorithms, making them seem much less practical than they may in fact be.

There are many useful consequences of (2.4). The simplest is

Corollary 2.5

If $f \in S_n$, then in polynomial time one can decide whether or not $f \in G$.

Proof:

Test whether $|\langle \Gamma \cup \{f\} \rangle| = |G|$.

Proposition 2.6 [FHL]

Given $\Delta \subset G$, in polynomial time the normal closure $\langle \Delta^G \rangle$ can be found.

Proof:

Successively test each $g \in \Gamma$ to see whether $\langle \Delta \rangle = \langle \Delta g \rangle$ (using (2.4i)). If this fails for some g replace $\langle \Delta \rangle$ by $\langle \Delta \cup \Delta g \rangle$, otherwise output $\langle \Delta \rangle$.

Corollary 2.7 [FHL]

The derived series and descending central series of G can be found in polynomial time. (Hence, solvability and nilpotence can be tested in polynomial time.)

Proposition 2.8

Given $A,B \leq S_n$ such that A normalizes B, the intersection $A \cap B$ can be found in polynomial time.

The proof follows from a straightforward extension of Sim's algorithm (cf. [FHL]). It should be noted, however, that in CAYLEY [C] all intersections are handled in the same manner, using a backtrack search. For further comments concerning intersections, see Section 3. For now we note one further situation in which intersections can be found in our context:

Theorem 2.9 [L1]

Given an integer b, there is a polynomial-time algorithm which, when given $G,H \leq S_n$ such that all noncyclic composition factors of G have order $\leq b$, finds $G \cap H$.

The algorithm for (2.9) has the disadvantage of requiring time $O(n^{f(b)})$ with $f(b) \to \infty$ as $b \to \infty$. Thus, while polynomial-time, for large b it is perhaps unreasonably theoretical.

Finally, we note that it is easy to find the center Z(G) in polynomial time. More generally, if $A \trianglelefteq G$ is given then $C_G(A)$ can be found in polynomial time [L3].

For an exposition of most of the above results, see [H].

3. GRAPH ISOMORPHISM

There are probably severe restrictions on what can be accomplished in polynomial time. Namely, consider the following four problems (where $G \leq S_n$ is as usual):

1. Given $G,H \leq S_n$, find $G \cap H$.
2. Given a p-subgroup P of G, find $N_G(P)$.
3. Given an involution $t \in G$, find $C_G(t)$.
4. Given $Y \subset X$ find the setwise stabilizer of Y in G.

Theorem 3.1

If any of the problems 1–4 can be solved in polynomial time, then so can the GRAPH ISOMORPHISM problem.

Here, GRAPH ISOMORPHISM is the following: Given two n-vertex graphs, decide whether or not they are isomorphic. The above somewhat surprising-looking result is due to Luks [L1,L3]. Parts of the theorem and other similar results of Luks are described in [H].

It is generally believed that there is no polynomial-time algorithm for GRAPH ISOMORPHISM. If that turns out to be the case then none of 1–4 can be accomplished in polynomial time. In any event, it should be evident that 1–4 must be avoided in the context of the present subject—except, of course, for the unlikely possibility that the study of

polynomial-time group-theoretic algorithms might produce a solution to the GRAPH ISOMORPHISM problem.

Finally, we note that (2.9) provides polynomial-time algorithms for 1, 3, and 4 when G is bounded as in (2.9) (e.g., if G is solvable). However, 2 remains open even for solvable groups G.

4. NONSOLVABLE GROUPS

There are a number of other elementary consequences of the results in Section 2. However, we will now move to more recent and more complicated results due to Luks [L2,BKL], Rónyai [R], or the author [K3,K4].

Let $G = \langle \Gamma \rangle$ be as usual.

Theorem 4.1 (Luks [L2]; cf. [BKL])

A composition series of G can be found in polynomial time.

Note that it is not possible to find all composition series, as there may be too many of them (once again, consider the regular representation of an elementary abelian 2-group).

Sketch:

First an auxiliary procedure PRIM is needed: if G acts transitively on a set X^* (not necessarily our original X) of size >1, PRIM finds a block system Σ of size >1 on which G acts primitively (using (2.3ii)), finds the kernel of this action (using (2.4)), and outputs that kernel if it is nontrivial or Σ if the kernel is trivial but $|\Sigma| < |X| = n$.

Next, note that all we need is to find either a smaller set on which G acts faithfully or a proper normal subgroup of G or to determine that no proper normal subgroup exists. For then recursion can be applied.

With this in mind, Luks' algorithm proceeds as follows. If, at any stage, we produce either a smaller set on which G acts faithfully or a nontrivial normal subgroup of G, then we can apply recursion.

1. Call PRIM for one nontrivial orbit of G.

WLOG G is primitive on X.

2. Test whether $G \neq G'$.
3. Pick any distinct $x,y \in X$, and find the set Z of fixed points of $G_{x,y}$. For each $z \in Z$,

   Test whether $(y,z) \in (x,y)^G$.
   If so find $g \in G$ with $(y,z) \in (x,y)g$, form the orbit $x^{\langle g \rangle}$, and test whether $|(x^{\langle g \rangle})^G| < n^2$.
   If so call PRIM for G on $Y = (x^{\langle g \rangle})^G$.

4. Pick $x \in X$. For each $y \in X - \{x\}$ call PRIM for $\{x,y\}^G$.
5. Pick $x \in X$. For each $y, z, w \in X$, let $H = \langle G_{xy}, G_{zw} \rangle$, and if $H \neq G$ then call PRIM for G/H.
6. If G passes all of the above steps then G is simple.

Comments:

3: If G has a regular normal subgroup then $x^{\langle g \rangle}$ will be a line of an n-point affine space; such a space has $< n^2$ lines. Note that we are not looking at the action of G on all p-element subsets of X, but rather on a severely limited collection of such subsets.

4-6: These are motivated by the O'Nan-Scott theorem (see, e.g., [AS Appendix]). Namely, if the socle N of G is not regular in its action on X, then N is the direct product of a certain number k of isomorphic nonabelian simple groups. If $k > 1$ and the action of G on X is the wreathed product action, then it is not difficult to show that 4 will produce a smaller set or a proper normal subgroup. If $k > 1$ and the action arises from a diagonal action on the direct product N (see [AS]), then the algorithm stops at 5. Finally, if $k = 1$ then, since 2 has been passed, $G = N$ by the correctness of Schreier's conjecture.

The above algorithm runs in time $O(n^8)$—with the exponent due, to a large extent, to the $O(n^5)$ for (2.4).

Corollary 4.2

(i) [L2] Simplicity of G can be tested in polynomial time.

(ii) For each successive pair $A \lhd B$ in the composition series, in polynomial time a set of size $\leq n$ can be found on which B/A acts faithfully.

Proof:

(i) is clear, so consider (ii). WLOG $G = B$. Let Y be the set of orbits of A on X. If $G^Y \neq 1$ output Y. WLOG $G^Y = 1$. Since G acts nontrivially on some member of Y, WLOG A is transitive on X. Now $G = AG_x$ for $x \in X$, so that $G/A \cong G_x/A_x$ and we can apply recursion to the pair G_x, $X - \{x\}$. (For a somewhat different argument, see [L1, (3.2)].)

Theorem 4.3 (Rónyai [R])

A chief series of G can be found in polynomial time.

Outline:

Using (4.1) it is easy to find a normal series for G each of whose factors is either (i) elementary abelian or (ii) the direct product of nonabelian simple groups permuted transitively by G. (Namely, consider the normal closures (2.6) of all the terms in (4.1).) Therefore, (4.3) can be viewed as a special case of the following situation. Given a set Δ of linear transformations of a finite vector space, find a Δ-irreducible subspace. Rónyai considers this latter problem in terms of the algebra of linear transformations generated by Δ. This is dealt with by an ingenious use of classical ideas concerning finite-dimensional algebras.

We briefly digress in order to present the following elementary variant of (4.3) that will be needed later:

Lemma 4.4 (Rónyai)

Given an m-dimensional vector space V over GF(p) and a set Γ of linear transformations, there is a polynomial (in m,p, and $|\Gamma|$) time algorithm that finds the space of fixed vectors of Γ.

Proof:

Use elementary linear algebra in order first to find the space of fixed vectors of each member of Γ and then to intersect these subspaces.

Simplicity is one of the standard and most basic questions concerning a finite group. Almost as basic are Cauchy's and Sylow's theorems. All of the standard proofs for these theorems either clearly do not produce polynomial-time algorithms or probably do not. For example, the proof of Cauchy via the "class equation" is purely existential. Similarly, the most standard proofs of the existence of Sylow subgroups involve—in addition to Cauchy's theorem—the use of normalizers or centralizers of p-subgroups of G, and these must be avoided by Section 3. (On the other hand, the algorithm in CAYLEY builds up a Sylow subgroup by using centralizers.) Other proofs of Cauchy or Sylow involve the examination of potentially exponential-size subsets of G. Finally, the conjugacy part of Sylow's theorem is standardly proved by a purely existential argument. Consequently, new techniques were required in order to obtain polynomial-time algorithms.

Cauchy's theorem was dealt with in [K2]. Once again, the classification of finite simple groups was involved! However, unlike the situation with (4.1), detailed information was needed concerning such groups (cf. Section 6).

In [KT], polynomial-time algorithms were obtained for special cases of Sylow's theorem, such as for solvable groups—in which case Hall's theorem was also dealt with. These solvable group algorithms were later modified in [K3] (cf. Section 5, and Section 6, Remark 1) in the process of obtaining methods that led to the general case:

Theorem 4.5 [K3]

If p is a prime then the following can be found in polynomial time:

(i) Given a p-subgroup of G, a Sylow p-subgroup of G containing it;

(ii) Given two Sylow p-subgroups of G, an element of G conjugating the first one to the second.

More recently, the set of all conjugating elements in (4.5ii) has, in effect, been specified:

Proposition 4.6 [K4]

Given a Sylow p-subgroup P of G, $N_G(P)$ can be found in polynomial time.

All of these results depend on detailed information concerning all finite simple groups—not, for example, just on the finiteness of the

number of sporadic simple groups. On the other hand, the main ideas of the proofs can be seen in two diametrically opposite situations: solvable groups, and switching permutation representations of simple groups. These will be discussed in the next two sections.

5. SOLVABLE GROUPS

In this section we will prove the following special cases of (4.5) and (4.6).

Theorem 5.1

Given a solvable subgroup G of S_n and a prime p, the following can be found in polynomial time:

(i) A Sylow p-subgroup P of G;

(ii) Given a Sylow p-subgroup P_0 of G, an element $g \in G$ such that $P^g = P_0$; and

(iii) $N_G(P)$.

Proof of (i) and (ii) [K3 Appendix]:

We will proceed by means of two reductions: from (i) to (ii), and then from (ii) to a third situation (which is, in fact, just a special case of (ii)). In each reduction we will also use recursion.

Reduction from (i) to (ii). Assume that we have available a polynomial-time algorithm for (ii); we will present a polynomial-time algorithm for (i).

Find $M \lhd G$ with $|G/M|$ prime [use G/G' (2.7)].
Let $g \in \Gamma - M$.
Recursively find a Sylow p-subgroup P of M.
Use (ii) to find $m \in M$ such that $(P^g)^m = P$.
Find a Sylow p-subgroup $\langle g' \rangle$ of the cyclic group $\langle gm \rangle$.
Then $\langle P, g' \rangle$ is Sylow in G.
[Note that, in effect, we have used the Frattini argument—cf. (5.2).]

Reduction from (ii) to the following statement (ii*):

(ii*) Given a solvable subgroup G of S_n, $M \lhd G$ with $|G/M| = p$, and Sylow p-subgroups P and P_0 of G such that $P \cap M = P_0 \cap M \trianglelefteq G$, find $g \in G$ such that $P^g = P_0$.

This time we are assuming the availability of a polynomial-time algorithm for (ii*). Using it, we will present a polynomial-time algorithm for (ii).

Find $M \lhd G$ with $|G/M|$ prime. WLOG this prime is p.
Find $P \cap M$ and $P_0 \cap M$. [Use (2.8)].
Recursively find $m \in M$ such that $P \cap M = (P_0 \cap M)^m$.
Let $G^* = \langle P, (P_0)^m \rangle$ and $M^* = G^* \cap M$ [using (2.8)].

Apply (ii*) to the quadruple G^*, M^*, P, $(P_0)^m$ in order to conjugate $(P_0)^m$ to P. [Note that $P \cap M$ is normal in P, $(P_0)^m$, and hence G^*, so that (ii*) applies.]

[The idea here was to "move P and P_0 closer together." Note that $G^*/P \cap M$ has Sylow subgroups of order p. This is essentially what (ii*) is all about.]

Algorithm for (ii*). We are given G, M, P and P_0. As above, the goal will be to move P and P_0 closer together.

WLOG $P \cap M < P$.
Let $h \in P - M$.
Let $h_0 \in P_0 - M$ with $h^{-1}h_0 \in M$.
Let $t = h^{-1}h_0$.
Let $u \in \langle t \rangle$ with $u^p t \in P \cap M$. [Note that $M/P \cap M$ is a p'-group.]
Let $m = u^h(u^2)^{h^2} \cdots (u^{p-1})^{h^{p-1}}$
Let $G^* = \langle P, (P_0)^m \rangle$ and $M^* = G^* \cap M$ [using (2.8)].
Apply recursion to the quadruple G^*, M^*, P, $(P_0)^m$.

[The fact that $G^* < G$ can be seen by taking a normal subgroup N of G such that $P \cap M \leq N < M$ and M/N is a G-chief factor and passing mod N. That is, we may assume that $N = 1$. Then G is a semidirect product of an elementary abelian q-group M (for some prime $q \neq p$) with a group P of order p acting irreducibly on M. Now there is exactly one element of M conjugating P_0 to P, which elementary linear algebra shows to be just m.]

This completes the proof of (i) and (ii).

Remark:

The algorithms just presented seem simple enough to be efficiently programmed. No visible use was made of the n-set X—although X was certainly used in essential ways since (2.4) and its consequences were employed often. It also seems as if it should be possible to effectively use some of the ideas in the above algorithms (or the next one or others in [K3 Appendix]). In fact, the Frattini approach and the linear algebra trick at the end of (ii*) have been used in [G].

Before we turn to (iii), we need the following

Corollary 5.2 (Frattini argument [KT])

If M is a normal subgroup of the solvable subgroup $G = \langle \Gamma \rangle$ of S_n, and if P is a Sylow p-subgroup of M, then in polynomial time a subgroup $D \leq N_G(P)$ can be found such that $G = DM$.

Proof:

For each $g \in \Gamma$ find $m \in M$ such that $P^{gm} = P$, and let D be the group generated by all of the resulting elements gm.

In (5.4) we will see that $N_G(P)$ itself can be found.

Corollary 5.3 [KT]

Given a solvable subgroup G of S_n and a prime p, the largest normal p-subgroup $O_p(G)$ of G can be found in polynomial time.

Proof:

If P is an intersection of some Sylow p-subgroups of G then successively test the elements $g \in \Gamma$ to see whether $P^g = P$. If this fails for some g, replace P by $P \cap P^g$, otherwise output P.

Proof of (5.1iii):

1. WLOG P is not normal in G.
2. Find $O_p(G)$ using (5.3). [Then $O_p(G) < P$.]
3. Find normal subgroups $K < L < M$ of G such that

   $O_p(G) \leq K$,
   L/K is an elementary abelian q-group for some $q \neq p$, and
   M/L is a p-group.

[Use (2.7): since $G/O_p(G)$ is not a p-group, K, L, and M exist and can be found by using the derived series (2.7) together with (2.4i).]

4. Find $R = P \cap M$ using (2.8).

[Then R is a Sylow p-subgroup of M, $R > O_p(G)$, and $M = RL$.]

5. If $RK \trianglelefteq G$ then use (2.6) and (2.7) to find K_1 with $O_p(G) \leq K_1 < K$, $K_1 \trianglelefteq G$, and K/K_1 elementary abelian; then replace the triple (M, L, K) by (RK, K, K_1), and return to 4.

[Clearly $K < RK$. If $RK \trianglelefteq G$ then K cannot be a p-group, so that $O_p(G) < K$. Hence K_1 exists, and can easily be found using K'. Since we are decreasing K, this loop eventually leads to the situation that RK is not normal in G.]

6. Find $D \leq N_G(R)$ with $G = DM$ and $D \geq R$. [Use (5.2). Note that $G = DRL = DL$.]
7. Find A with $K \leq A \leq L$ and $C_{L/K}(R) = A/K$, using (4.4).

[Since L/K is an elementary abelian r-group for some prime $r \neq p$, it can be viewed as a vector space, so that (4.4) applies.]

8. Recursively find and output $N_{\langle D,A\rangle}(P)$.

[We must show that $N_{\langle D,A\rangle}(P) = N_G(P)$ and that $\langle D,A\rangle \neq G$. If $^-$ is the natural homomorphism $G \to G/K$, then $[N_{\bar{L}}(\bar{R}),\bar{R}] \leq \bar{L} \cap \bar{R} \leq \bar{K} = 1$ (since $\bar{L}$ is a p'-group and $\bar{R}$ is a p-group), so that $N_{\bar{L}}(\bar{R}) \leq C_{\bar{L}}(\bar{R}) = \bar{A}$. Also, $N_G(P) \leq N_G(R)$, while $N_{\bar{G}}(\bar{R}) = N_{\bar{D}\bar{M}}(\bar{R}) = N_{\bar{D}\bar{R}\bar{L}}(\bar{R}) = \bar{D}N_{\bar{L}}(\bar{R}) \leq \bar{D}\bar{A} < \bar{G}$ since $\bar{R} = RK/K$ is not normal in $\bar{G} = G/K$.]

Corollary 5.4

Given a Sylow p-subgroup P of a normal subgroup M of a solvable group $G \leq S_n$, $N_G(P)$ can be found in polynomial time.

Proof:

Find $N_M(P)$ using (5.1iii). Find $D \leq N_G(P)$ such that $G = DM$ using (5.2). Then $DN_M(P)$ is the desired normalizer.

It should be evident that all of the above arguments used in the proof of (5.1) are very different from the standard ones. It is also clear that they are far less transparent than the purely existential proofs found in textbooks.

It is natural to ask exactly where solvability was used in this section, and how the classification of finite simple groups might enter into the nonsolvable case. Note that, in the proof of (5.1i,ii), we began with a normal subgroup M such that G/M was cyclic of prime order. In the general case of (4.5), this quotient group might be a nonabelian simple group. In [K3] that situation was reduced to the case of a simple group G, and then Sylow subgroups of G were found and conjugated by using (4.2ii) and the Replacement Theorem of the next section.

6. THE REPLACEMENT THEOREM

We begin with a result that is a fairly straightforward consequence of the results in [K1], together with some geometry of the classical groups. Let G be a group acting on an n-set X. Let $x \in X$. For each $y \in X$ let $\mathfrak{A}_X(x,y)$ consist of $G_{\{x,y\}}$ together with the set of all proper subgroups of G of which $G_{\{x,y\}}$ is a maximal subgroup, and let $\mathfrak{A}_X(x,y)^*$ consist of $\mathfrak{A}_X(x,y)$ together with all proper subgroups of G of which some member of $\mathfrak{A}_X(x,y)$ is a maximal subgroup. Let $\mathfrak{B}(G,X)$ be any set of maximal subgroups of G such that each member of $\cup \{\mathfrak{A}_X(x,y)^* \mid y \in X\}$ is contained in some member of $\mathfrak{B}(G,X)$. Finally, let

$$\mathfrak{B}(G) = \cup \{\mathfrak{B}(G,G/M) \mid M \in \mathfrak{B}(G,X)\}$$

and

$$b(G) = \min\{|G/H| \mid H \in \mathfrak{B}(G)\}$$

Note that $b(G) \leq n$ since $y = x$ was allowed. Using all of this notation, we have the following useful result:

Proposition 6.1

Let T be a simple group, Let $T \leq G \leq \mathrm{Aut}(T)$, and suppose that G acts primitively on an n-set X. Then one of the following holds:

(a) $|G| < n^8$; or

(b) If $M \in \mathfrak{B}(G)$ is such that $b(G) = |G : M|$, then either

(i) M is a proper normal subgroup of G, or

(ii) T restricted to G/M is equivalent to the faithful permutation representation of T of smallest degree.

The representation in (ii) is either the usual action on an r-set if $T \cong A_r$, or the action on the unique shortest orbit of 1-spaces or hyperplanes of the underlying vector space if T is a classical group.

Proof:

This is essentially the same as [K2, Theorem 6.1]. If we assume that (a) does not hold then the simple normal subgroup T of G is alternating

or classical, and the permutation representation of G on X is very restricted [K1]. (Namely, either $T \cong A_r$ and G_X is the stabilizer of a subset or a partition of the r-set into subsets of equal size, or T is classical and G_X is the stabilizer of a subspace.) The proof of the aforementioned theorem goes through with one minor change: it is conceivable that some $M \in \mathfrak{B}(G)$ contains T, in which case (i) can occur (for $T = PSL(d,q)$ or $P\Omega^+(2d,q)$ and suitable G).

Note that (6.1bi) produces a normal subgroup of G to which (6.1) can again be applied. Moreover, when that is done, (6.1bi) does not occur a second time. Also, the bound in (6.1a) (and the next result) can be improved, with n^4 in place of n^8.

Theorem 6.2 (Replacement Theorem) [K2]

Given a simple subgroup G of S_n of order $\geq n^8$, there is a polynomial-time algorithm that finds the natural permutation representation of G (and that permutation representation has degree $< 2n$).

Proof:

Use (2.2) and (2.3) in order to reduce to the case in which G acts primitively. By (2.3), all of the sets $\mathfrak{A}_X(x,y)$, $\mathfrak{A}_X(x,y)^*$, $\mathfrak{B}(G,X)$, and $\mathfrak{B}(G)$ can be found in polynomial time.

In fact, much more is proved in [K2,K3] in the case of classical groups: in polynomial time the underlying vector space can be found, as can a group of linear transformations inducing G. Moreover, in the case of a symplectic, orthogonal, or unitary group, a suitable form is constructed on the vector space. This has the effect of replacing permutation group considerations by linear algebra. In view of this, it should come as no surprise that the simple group case of Theorems 4.4 and 4.5 can be deduced from the Replacement Theorem (by means of some rather tedious work).

7. CONCLUDING REMARKS

Remark 1:

There are also polynomial-time algorithms for Hall's theorem and related results [KT; K3 Appendix]; Carter subgroups of solvable groups can be found and conjugated to one another (in polynomial time, as usual), as can system normalizers; and the Fitting subgroup and the generalized Fitting subgroup can be found, as can $O_\pi(G)$ for any set π of primes. (In the case of Hall's theorem, the arguments in Section 4 go through almost verbatim.)

Remark 2:

It is not known whether the Frattini subgroup or the ascending central series can be found in polynomial time (though they probably can). Of course, more technical group-theoretic subgroups such as the Thompson

subgroup J(P) of a p-group probably cannot be found since their definitions involve knowledge of potentially exponential-size sets (the set of all maximal-size abelian subgroups of P). Moreover, many of the more familiar types of group-theoretic methods involve normalizers of p-subgroups, and hence must be avoided (cf. Section 3).

Remark 3:

Some of the methods described here have some applicability to a somewhat similar but much harder subject: polynomial-time Galois theory. There, one is given a finite extension K of $\mathbb{Q}$ and a polynomial $f \in K[x]$ and asked to find the Galois group of f. For simplicity, assume that $f \in \mathbb{Z}[x]$ (although extensions of $\mathbb{Q}$ must eventually be considered as well). Then the problem is to determine Gal(f) in time that is polynomial in the number of (binary or decimal) digits required to write f. Analogs of (2.2) and (2.3) exist: f can be factored into irreducibles [LLL]; an extension $L = \mathbb{Q}(\alpha)$ of $\mathbb{Q}$ can be obtained with $f(\alpha) = 0$ (described as a vector space over $\mathbb{Q}$ with a distinguished basis and a multiplication rule for that basis); and, when f is irreducible, subfields of L can be specified in polynomial time that correspond to the blocks involved in (2.3) (specified as the sets of roots of explicitly constructed polynomials) [LM].

This situation is harder than the one in this paper because (i) Galois groups are determined only up to conjugacy in symmetric groups; (ii) no efficient way is known for finding a nontrivial element of G = Gal(f) (except possibly for complex conjugation); and (iii) a splitting field of f generally has nonpolynomial degree over $\mathbb{Q}$, and hence cannot be written (as a vector space over $\mathbb{Q}$) in polynomial time. In view of these difficulties, it is not surprising that no polynomial-time algorithm is known for determining $|G|$.

Only the following have been proved: in polynomial time it can be decided whether or not G is solvable [LM] (or a p-group); and weak analogs of (4.1) and (4.2) have been obtained [KL,K5] based on the fact that all of the sets appearing in (6.1) can be found in polynomial time.

ACKNOWLEDGMENT

The preparation of this paper took place while the author was a member at the Institute for Advanced Study, whose hospitality is gratefully acknowledged. Partial support was provided by National Science Foundation DMS-8320149 (University of Oregon) and MCS-8108814(A04) (Institute for Advanced Study).

REFERENCES

[A] M. D. Atkinson, An algorithm for finding the blocks of a permutation group, Math. Comput. 29 (1975), 911–913.

[AS] M. Aschbacher and L. L. Scott, Maximal subgroups of finite groups, J. Algebra 92 (1985), 44–80.

[B] L. Babai, On the length of subgroup chains in the symmetric group, Comm. Algebra, 14 (1986), 1729–1736.

[BKL] L. Babai, W. M. Kantor, and E. M. Luks, Computational complexity and the classification of finite simple groups, Proc. 24th IEEE Symposium Foundations of Computer Science (1983), 162–171.

[C] J. J. Cannon, An introduction to the group theory language CAYLEY, Computational Group Theory, M. D. Atkinson (ed.), Academic Press, New York (1984), 145–183.

[F] L. Finkelstein, personal communication.

[FHL] M. Furst, J. Hopcroft, and E. Luks, Polynomial-time algorithms for permutation groups, Proc. 21st IEEE Symposium Foundations of Computer Science (1980), 36–41.

[G] S. P. Glasby, Constructing normalizers in finite soluble groups, in preparation.

[H] C. M. Hoffmann, Group-theoretic algorithms and graph isomorphism, Lecture Notes in Computer Science, 136, Springer, Berlin (1982).

[J] M. R. Jerrum, A compact representation for permutation groups, Proc. 23rd IEEE Symposium Foundations of Computer Science (1982), 126–133.

[K1] W. M. Kantor, Permutation representations of the finite classical groups of small degree or rank, J. Algebra 60 (1979), 158–168.

[K2] W. M. Kantor, Polynomial-time algorithms for finding elements of prime order and Sylow subgroups, J. Algorithms, 6 (1985), 478–514.

[K3] W. M. Kantor, Sylow's theorem in polynomial time, J. Comput. Syst. Sci. 30 (1985), 359–394.

[K4] W. M. Kantor, Finding Sylow normalizers in polynomial time, to appear in J. Algorithms.

[K5] W. M. Kantor, in preparation.

[KL] W. M. Kantor and E. Lander, in preparation.

[KT] W. M. Kantor and D. E. Taylor, Polynomial-time versions of Sylow's theorem, J. Algorithms, 9 (1988), 1–17.

[L1] E. M. Luks, Isomorphism of graphs of bounded valence can be tested in polynomial time, J. Comput. Syst. Sci. 25 (1982), 42–65.

[L2] E. M. Luks, Computing the composition factors of a permutation group in polynomial time, Combinatorica, 7 (1987), 87–99.

[L3] E. M. Luks, unpublished.

[LLL] A. K. Lenstra, H. W. Lenstra, and L. Lovász, Factoring polynomials with rational coefficients, Math. Ann. 261 (1982), 513–534.

[LM] S. Landau and G. L. Miller, Solvability by radicals is in polynomial time, J. Comput. Syst. Sci. 30 (1985), 179–208.

[R] L. Rónyai, Zero divisors and invariant subspaces, to appear.

[S1] C. C. Sims, Computational methods in the study of permutation groups, Computational Problems in Abstract Algebra, J. Leech (ed.), Pergamon, Elmsford, N.Y. (1970), 169–183.

[S2] C. C. Sims, Some group-theoretic algorithms, Lecture Notes in Math., 697, Springer, Berlin (1978), 108–124.