DIMACS Series in Discrete Mathematics
and Theoretical Computer Science
Volume 11, 1993

Computing Composition Series in Primitive Groups

LÁSZLÓ BABAI, EUGENE M. LUKS, AND ÁKOS SERESS

ABSTRACT. We give a fast polynomial-time algorithm for computing a composition series in a primitive permutation group given by a list of generators. Permutation representations for the composition factors are also obtained.

The algorithm will be a key procedure in an $O(sn^3 \log^c n)$ algorithm that will the solve the same problems for an arbitrary subgroup of S_n given by s generators. The general algorithm will be described in a forthcoming paper.

The first polynomial-time algorithm for this problem was given by Luks. Our procedure follows the overall architecture of that original algorithm, while replacing the subroutines by much faster ones. New combinatorial estimates of suborbit sizes in primitive groups guarantee the improved performance.

1. Introduction

In recent years, polynomial-time algorithms have been found for a great number of permutation-group problems. These range from constructing a strong generating set (SGS) [24], which is the fundamental data structure for membership testing and other basic tasks (finding the order, constructing normal closures, etc.), to obtaining structural information, such as composition factors [20] and Sylow subgroups [14]. (See [15] and [21] for other examples.) In this paper, we develop a more efficient procedure for the Composition-Series Problem: *Given generators for a permutation group G, find generators for the subgroups in a composition series $G = N_0 \rhd N_1 \rhd \ldots \rhd N_m = 1$ and permutation representations for the composition factors N_i/N_{i+1}.*

1991 *Mathematics Subject Classification.* Primary 20B40,68Q40; Secondary 20B15.

The first author was supported in part by NSF Grant CCR-9014562 and "OTKA" Grant (Hungary) #2581. The second author was supported in part by NSF Grant CCR-9013410. The third author was supported in part by NSF Grant CCR-9201303 and NSA Grant MDA904-92-H-3046.

This is the final version of this paper.

Our main result is

THEOREM 1.1. *Given s permutations generating a primitive group $G \leq S_n$, a composition series for G, together with permutation representations for the composition factors, can be found in $O^{\sim}(n^3 + sn^2)$ time.*

In the above, as in other stated results, we use the "soft" version ([5]) of the big-O notation: for two functions $f(n), g(n)$, we write $f(n) = O^{\sim}(g(n))$ (read "soft-oh") if, for some constants $c, C > 0$, $f(n) \leq Cg(n)\log^c n$. Note, in this notation, that we ignore the possible dependency of f and g on s. Indeed, the main impact of our present algorithms is in the (most typical) situation where s is small relative to n.

For possible implementations, it may cause concern that the exponent of $\log n$, hidden in the tilde notation, is rather large (greater than 20 in some lower order terms $n\log^c n$). However, for the present, our objective is to establish an asymptotic performance guarantee that is several orders of magnitude better than those previously known and we do not attempt to minimize the exponent of $\log n$. Nevertheless, we do point out that, in the "practical" range of n, say $n < 10^6$, one can eliminate the procedures that contribute these large exponents.

The first polynomial-time, $O(n^8 + sn^2)$, algorithm for finding a composition series was given by Luks [20]. Although Luks's goal was only to demonstrate that the problem is solvable in polynomial time, he also established the requisite structure, group-theoretic and algorithmic, upon which to build improvements. Thus, the overall architecture of our algorithm still follows Luks's original procedure, while the subroutines have been replaced by much faster ones, taking advantage of new, mostly combinatorial estimates.

The case of primitive groups is a crucial instance of the composition-series problem for general groups. The algorithm described herein will play a key role in a forthcoming paper where the following general result will be proved.

THEOREM 1.2. [6] *Given s permutations generating $G \leq S_n$, a composition series for G, together with permutation representations for the composition factors, can be found in $O^{\sim}(sn^3)$ time.*

Following [20], we remark that it is trivial to reduce the composition-series problem to the primitive case in polynomial time. However, the naive reduction will not yield the strong timing estimates stated in Theorem 1.2; indeed, it is highly nontrivial to perform even the rudimentary task of finding strong generators within the time bound stated. The latter accomplishment is the main result of [6] and the results of the present paper are used as a subroutine in the SGS construction given there. By the time the SGS has been constructed, a composition series is a byproduct.

En route to Theorem 1.1, the present paper also resolves the following subcase of Theorem 1.2, which involves "small-base" groups (see Section 2.2).

THEOREM 1.3. *Let c be fixed. Given s permutations generating $G \leq S_n$, where $\log|G| = O(\log^c n)$, a composition series for G, together with permutation representations for the composition factors, can be found in $O^{\sim}(n^3 + sn)$ time.*

We also point out that our algorithm for this problem requires only $O^{\sim}(sn)$ space.

While no deterministic algorithm is currently known to beat the asymptotic timing of Theorem 1.3, Beals and Seress [8] have recently constructed a randomized (Monte Carlo) algorithm for finding a composition series in small-base groups, running in *nearly linear*, $O^{\sim}(ns)$ time. That algorithm follows the outline of the present paper and uses novel randomization techniques to overcome the n^2 and n^3 bottlenecks (cf. Section 4).

To handle the primitive groups of large order, our algorithm is strongly guided by their structure theory (cf. Theorems 2.1, 2.7, and 2.8 below).

The timing analysis of our algorithm depends on the classification of finite simple groups via estimates of the orders of primitive groups (Theorems 2.1, 2.2) and Schreier's Conjecture (Theorem 2.3). We shall also make use of the Odd Order Theorem [11].

2. Notation, Definitions and Background

2.1. Group Theory. For basic definitions and facts about groups and permutation groups, we refer the reader to [12] and [26].

For $H \leq G$, H^G denotes the *normal closure* of H and $\mathrm{N}_G(H)$ is the *normalizer* of H. The *automorphism group* and *outer automorphism group* of G are denoted by $\mathrm{Aut}(G)$ and $\mathrm{Out}(G)$, respectively. The *socle* of G, $\mathrm{Soc}(G)$, is the subgroup of G generated by all minimal normal subgroups. If $G = T_1 \times \cdots \times T_r$ is the direct product of groups each isomorphic to a fixed group T, the *diagonal* $\mathrm{Diag}(T_1 \times \cdots \times T_r)$ of G is the subgroup comprised of $\{f_1(t)f_2(t)\cdots f_r(t) \mid t \in T\}$ where, for each i, $f_i : T \to T_i$ is a fixed isomorphism. For $H \leq G$, G/H denotes the set of right cosets of H in G.

The full symmetric group acting on the domain Ω is denoted by $\mathrm{Sym}(\Omega)$, while A_n and S_n denote the alternating and symmetric groups of degree n, respectively. We say that G *acts* on a set Ω if a homomorphism $G \to \mathrm{Sym}(\Omega)$ is given; the action is *faithful* if its kernel is the identity. If G acts on Ω and $\omega \in \Omega$, we denote by ω^G the *orbit* of ω under G, namely $\{\omega^g \mid g \in G\}$; we say G acts *transitively* on Ω if Ω consists of one orbit. If G acts transitively on Ω, a nonempty subset $\Delta \subseteq \Omega$ is called a *block of imprimitivity* for G if, for all $g \in G$, $\Delta^g = \Delta$ or $\Delta^g \cap \Delta = \emptyset$; the singletons and Ω itself are the *trivial* blocks; G is said to act *primitively*, and is said to be *primitive* in the case $G \leq \mathrm{Sym}(\Omega)$, if it has no nontrivial blocks of imprimitivity. A G-invariant partition $\mathcal{B}$ of Ω is *minimal* if G acts primitively on $\mathcal{B}$. For $B \subseteq \Omega$, G_B denotes the *pointwise stabilizer* of B in G. A transitive group $G \leq \mathrm{Sym}(\Omega)$ is called *regular* if $G_\omega = 1$ for any (all) $\omega \in \Omega$.

When speaking of subgroups of S_n, we call S_n and A_n the *giants*. They are the two largest subgroups and, by far, the two largest primitive subgroups of S_n.

For the analysis of our algorithm, we need several consequences of the classification of finite simple groups. The first one originated with Cameron [9] who, based on results of Kantor [13], classified the primitive groups of degree n and order $> n^{c \log\log n}$. We state a simplified version due to Liebeck [17], suitable for algorithmic purposes. (Unless otherwise indicated, logarithms in this paper are to base 2.)

THEOREM 2.1. [17] *Let $G \leq \mathrm{Sym}(\Omega)$ be primitive and let $n = |\Omega|$. Then either $|G| \leq n^{9 \log n}$ or the socle of G is the product of isomorphic alternating groups.*

We also need an estimate for the order of non-giant primitive groups $G \leq S_n$. The following can be deduced from results of Cooperstein [10] and Liebeck, Saxl [19].

THEOREM 2.2. *Let $G \leq \mathrm{Sym}(\Omega)$ be a non-giant primitive group, and let $n = |\Omega|$. If $n \notin \{8, 11, 12, 24\}$ then $|G| \leq 5^{n/2}$. In the exceptional cases, only $AGL(3,2)$ (for $n = 8$) and the Mathieu groups (for $n = 11, 12, 24$) violate this bound. In any case, $|G| \leq 7(5^{n/2})$.*

Remark. Elementary estimates $|G| < 4^n$ and $|G| \leq n^{4\sqrt{n} \log n}$ have been obtained by Praeger, Saxl [22] and Babai [2, 3], respectively. While Babai's bound is asymptotically sharper than $7(5^{n/2})$, the latter is better, and more convenient, for small n.

We cite the following additional consequence of the classification, traditionally termed Schreier's "Conjecture".

THEOREM 2.3. *If G is simple then $\mathrm{Out}(G)$ is solvable. Moreover, if G has a permutation representation of degree n, then $|\mathrm{Out}(G)|$ has no prime factor $> \sqrt{n}$.*

2.2. Algorithmic Concepts and Results. It is assumed that permutation groups are input or output via a list of generators.

In his pioneering work [24],[25], Sims introduced the notions of a base and strong generating set as the fundamental data structures for computing with permutation groups. A *base* for a permutation group $G \leq \mathrm{Sym}(\Omega)$ of degree n is a subset $B = \{b_1, b_2, .., b_M\}$ of Ω such that $G_B = 1$. Viewing B as a linearly-ordered set, the *point-stabilizer chain* of G relative to B is the chain of subgroups

$$G = G^{(1)} \geq G^{(2)} \geq \ldots \geq G^{(M+1)} = 1,$$

where $G^{(i)} = G_{\{b_1,..,b_{i-1}\}}$. The base B is called *nonredundant* if there is strict inclusion $G^i > G^{i+1}$ for all $1 \leq i \leq M$. A *strong generating set* (SGS) for G relative to B is a set S of generators of G with the property that

$$\langle S \cap G^{(i)} \rangle = G^{(i)}, \text{ for } 1 \leq i \leq M+1.$$

With reference to some fixed $c > 0$, an (infinite) family $\mathcal{G}$ of groups is called a family of *small-base groups* if all members G of degree n admit bases of size $O(\log^c n)$. Equivalently, $\log|G| = O(\log^{c'} n)$ for a fixed constant c' and each $G \in \mathcal{G}$ of degree n.

Let $B = \{b_1, ..., b_M\}$ be a base of the group G and let $G = G^{(1)} \geq G^{(2)} \geq ... \geq G^{(M+1)} = 1$ be the corresponding point-stabilizer chain. Moreover, let R_i denote a *transversal* (complete set of right coset representatives) for $G^{(i+1)}$ in $G^{(i)}$, $1 \leq i \leq M$. Each $g \in G$ can be written uniquely in the form $g = r_M r_{M-1} \cdots r_2 r_1$, $r_i \in R_i$. The process of factoring g in this form is called *sifting* or *stripping*.

Knuth [16] has offered a version of the Schreier–Sims SGS construction that is particularly suitable for small-base groups. Although he did not explicitly analyze the timing in this fashion, it is straightforward to check the following.

THEOREM 2.4 ([16]). *Given $G = \langle S \rangle \leq S_n$, $|S| = s$, the following can be computed in $O(n^2 \log^3|G| + sn \log|G|)$ time.*

(i) *A base of size $\leq \log|G|$ and an SGS of size $\leq \log^2|G|$ relative to this base. The SGS allows for membership-testing in the group in time $O(n \log|G|)$.*

(ii) *Strong generators for the kernel of an action $G \to \mathrm{Sym}(\Delta)$ on a set Δ of size $O(n)$.*

Remark. The space requirement of Knuth's procedure can be expressed as $O(n^2 \log|G|)$. In [4, Lemma 2.2], Babai, Cooperman, Finkelstein, and Seress gave a version of Sims's Schreier vector data structure [24] for the storage of the (partial) transversals. This allows a time-space trade-off, adjusting the space demand to $O(n \log^2|G|)$ while increasing the time for SGS construction to $O(n^2 \log^4|G| + sn \log^2|G|)$ and membership-testing to $O(n \log^2|G|)$. Note, however, that for small-base groups this trade-off improves the space requirement by essentially an order of magnitude, from $O^\sim(n^2)$ to $O^\sim(n)$, while maintaining a time of the form $O^\sim(n^2 + sn)$ for SGS construction and $O^\sim(n)$ for membership-testing.

If a base and an SGS for a small-base group are already in hand, certain further elementary operations can be performed in nearly linear time. Standard methods, augmented by [4, Lemma 2.2], yield the following theorem. (Note: although [4] deals with probabilistic methods, the cited lemma is deterministic.)

THEOREM 2.5. *If a base B of size $O(\log|G|)$ and an SGS of size $O(\log^2|G|)$ are known for $G \leq \mathrm{Sym}(\Omega)$ then, given s' generators for $H \leq G$, the following can be computed in time $O(n \log^5|G| + s' log^3|G|)$.*

(i) *An SGS of size* $O(\log^2 |G|)$ *for* H *(relative to the base* B*).*
(ii) *An SGS of size* $O(\log^2 |G|)$ *for the normal closure of* H*.*

Atkinson's algorithm [1] shows that blocks of imprimitivity can be found in $O(sn^2)$ time. While this bound suffices for Theorem 1.1, we observe that, for small-base groups, Beals [7] has the following improvement.

THEOREM 2.6 ([7]). *Given* $G = \langle S \rangle \leq S_n$, $|S| = s$, *it can be decided in time* $O(n \log^3 |G| + sn \log |G|)$ *whether or not* G *is primitive. If not, the procedure finds a nontrivial block of imprimitivity.*

Finally, we quote two results from [5] which enable us to handle the large primitive groups.

THEOREM 2.7 ([5]). *Given a primitive group* $G = \langle S \rangle \leq S_n$, $|S| = s$, *it can be recognized in* $O^{\sim}(n^3 + sn^2)$ *time whether or not* G *contains* A_n*. If it does, then, within the same time bound, a strong generating set of size* n *can also be constructed from* S*.*

Remark. The phrase "constructed from S" in the above signals the construction of a *straight-line program*, that is, a sequence of group elements starting with elements of S, in which each successive element is determined as a product or power of some predecessor(s). Thus, for example, having identified a group as A_n, we do not allow the luxury of writing down arbitrary even permutations as members. The reason for such restriction is that we must deal with groups that are alternating only on some orbit or some induced representation, but permutations must be specified globally.

THEOREM 2.8 (NATURAL ACTION [5]). *If the order of the primitive group* $G = \langle S \rangle \leq S_n$ *exceeds* $n^{9 \log n \log \log n}$ *and* $\mathrm{Soc}(G) \simeq (A_k)^r$ *then these properties can be recognized and the natural imprimitive action of* G *on a set of size* kr*, with the factors of the socle acting on blocks of size* k *as* A_k*, can be constructed in* $O^{\sim}(sn^2)$ *time.*

Remark. The discrepancy between the bounds $n^{9 \log n}$ in Theorem 2.1 and $n^{9 \log n \log \log n}$ in Theorem 2.8 is due to the fact that, for small values of n, our algorithm may not recognize the imprimitive action if $|G| \leq n^{9 \log n \log \log n}$. Asymptotically, the cutoff point can be pushed down to $n^{5 \log n \log \log n}$ [5].

3. Primitive Permutation Groups

In this section, we state the results about primitive groups that are necessary to prove the correctness and timing of the composition-series algorithm.

A central ingredient is the O'Nan-Scott theorem, which classifies primitive groups according to their socles [23],[9],[18]. The statement below represents a regrouping of the subcases to suit our purposes. Note that a primitive group may belong to more than one subcase.

THEOREM 3.1. *Let G be a primitive permutation group of degree n. Then at least one of the following holds:*

(i) *G is cyclic of prime order.*

(ii) *G has a proper normal subgroup of index $\leq 7\sqrt{n}$.*

(iii) *G has one or two regular normal subgroups.*

(iv) *$N = \mathrm{Soc}(G)$ is the unique minimal normal subgroup of G, and it admits a decomposition $N = N_1 \times \ldots \times N_m$, with $m!/2 > 7\sqrt{n}$ (and, consequently, $m \geq 1 + \log n/(2 \log\log n)$), where the N_i are conjugate under the action of G. Moreover, G acts on $\{N_i\}_{1 \leq i \leq m}$ as the full alternating group A_m. In addition, one of the following holds.*

(a) $N_\omega = (N_1)_\omega \times \ldots \times (N_m)_\omega$, $1 \neq (N_i)_\omega \neq N_i$.

(b) $N = T_1 \times \cdots \times T_r$ for some $m|r$, where the T_i are isomorphic nonabelian simple groups and $N_\omega = \mathrm{Diag}(T_1 \times \ldots \times T_r)$.

(v) *G is nonabelian simple.*

PROOF. Let G be a primitive group. If G has a regular normal subgroup then G falls in case (i) or (iii). Otherwise, G has a unique minimal normal subgroup, $\mathrm{Soc}(G) = T_1 \times \ldots \times T_r$ where the T_i are isomorphic nonabelian simple groups, and $r \leq \log_5 n$. If $r = 1$ then $T_1 \leq G \leq \mathrm{Aut}(T_1)$ and, by Theorem 2.3, G falls in case (ii) or (v).

If $r > 1$ then G acts, by conjugation, transitively on $\{T_1, \ldots, T_r\}$. Let M be the pointwise stabilizer of a minimal block system in this transitive action. Then $M \lhd G$ and $G/M \simeq L$ for some primitive group L, $L \leq S_m$ for some $m|r$. If L does not contain A_m then, by Theorem 2.2, $|L| \leq 7(5^{r/2}) \leq 7\sqrt{n}$ and G falls in case (ii). If $L = S_m$ then G has a normal subgroup of index 2. Finally, if $L = A_m$ then either $m!/2 \leq 7\sqrt{n}$ and G falls in case (ii) or $m!/2 > 7\sqrt{n}$ and G is in case (iv). Each N_i is the direct product of the T_j in a block of the minimal block system. The two subcases of case (iv) involve an elementary rearrangement (as in [20, pp. 96-97]) of cases in the O'Nan-Scott Theorem. □

Remark. To establish Theorem 1.1, it would suffice to work with the coarser bound n in case (ii) of Theorem 3.1, eliminating the need for Theorem 2.2. However, the $7\sqrt{n}$ bound eliminates one n^3 bottleneck in the algorithm.

DEFINITION. We say that primitive groups of the types described in Theorem 3.1 (iv)(a) and (iv)(b) are of *type (iv)a* and *type (iv)b* respectively.

In order to speed up the handling of primitive groups of types (iv)a and (iv)b, we must strengthen some of the results of [20]. We start by examining the connection between the conjugation action of G on $\{N_1, \ldots, N_m\}$ and the orbits of G_ω, for $\omega \in \Omega$.

We have $N = \mathrm{Soc}(G) = T_1 \times \ldots \times T_r$ where the T_i are isomorphic nonabelian simple groups and G acts (via conjugation) as A_m on a minimal block system in $\{T_j\}_j$ for some $m|r$; the blocks in $\{T_j\}_j$ correspond to the N_i in case (iv) of Theorem 3.1. Furthermore, since $G/N \simeq G_\omega/N_\omega$, G_ω also acts as A_m on

$\{N_1, \dots, N_m\}$. Let $K \lhd G_\omega$ be the kernel of this action. Then $G_\omega/K \simeq A_m$, and conjugation by K fixes each N_i. Note that $N_\omega \leq K$.

Since N is transitive, Ω can be identified with N/N_ω via $\omega^x \leftrightarrow N_\omega x$; considering the induced conjugacy-action of G_ω on N/N_ω, this identification is a G_ω-map. In particular, the orbits of G_ω correspond to the G_ω conjugacy classes in N/N_ω. Also, if G is of type (iv)a, we have

$$\Omega \stackrel{\cdot}{\simeq} N/N_\omega \simeq N_1/(N_1)_\omega \times \cdots \times N_m/(N_m)_\omega. \tag{1}$$

In this situation, let $\psi \in \Omega$ and take $x \in N$ such that $\psi = \omega^x$. Then $x = x_1 \cdots x_m$ with $x_i \in N_i$. We call $(N_i)_\omega x_i$ the i^{th} *coordinate* of ψ and call this coordinate *nontrivial* if $x_i \notin (N_i)_\omega$.

We also need the following general observation.

FACT 1. *Let $\bar{G} \leq \mathrm{Sym}(\bar{\Omega})$ be primitive and $\bar{N}$ be a nontrivial, nonregular normal subgroup of $\bar{G}$. For $\bar{\omega} \in \bar{\Omega}$, the only point fixed by the stabilizer $\bar{N}_{\bar{\omega}}$ is $\bar{\omega}$.*

Indeed, the equivalence relation on $\bar{\Omega}$ defined by

$$\bar{\omega} \sim \bar{\psi} \text{ iff } \bar{N}_{\bar{\omega}} = \bar{N}_{\bar{\psi}}$$

is clearly $\bar{G}$-invariant. □

LEMMA 3.2. *Let $G \leq \mathrm{Sym}(\Omega)$ be a primitive group of type (iv)(a). Then, for any $\omega \in \Omega$, the size of the smallest orbit of G_ω on $\Omega \setminus \{\omega\}$ is $< \log^3 n$. Furthermore, the elements of such a smallest orbit have precisely one nontrivial coordinate.*

PROOF. By (1), $n = |N_1 : (N_1)_\omega|^m$ and therefore $|N_1 : (N_1)_\omega| \leq \log^2 n$, since $m \geq 1 + \log n/(2 \log\log n)$.

As above, let $K \lhd G_\omega$ be the kernel of the conjugation action of G_ω on $\{N_1, \dots, N_m\}$. Let q be the size of the smallest K-orbit in $N_1/(N_1)_\omega$. We claim that if $\psi \in \Omega$ has k nontrivial coordinates, then the G_ω-orbit of ψ has at least $\binom{m}{k} q 2^{k-1}$ conjugates under G_ω. To show this, we first note that the $\binom{m}{k}$ term accounts for the action of $A_m \cong G_\omega/K$. We may now assume that the k nontrivial coordinates of ψ occur in $N_i/(N_i)_\omega$, $i = 1, \dots, k$. The term q derives from the action of K on $N_1/(N_1)_\omega$. Finally, we see that there is an additional independent factor of 2 contributed by the conjugacy action of $(N_i)_\omega$ on $N_i/(N_i)_\omega$, for each i, $2 \leq i \leq k$; for this, observe, if a coset $(N_i)_\omega v$ were fixed by $(N_i)_\omega$ then $N_\omega v$ would be fixed by N_ω, contrary to Fact 1.

On the other hand, there is an element $\psi = \omega^u \in \Omega$ with one nontrivial coordinate, such that $N_\omega u$ has q conjugates under K, so that the G_ω-orbit of ψ has cardinality mq.

Hence coset representatives in the smallest G_ω-orbit have only one nontrivial coordinate and the size of the smallest G_ω-orbit is $mq \leq (\log_5 n)|N_1 : (N_1)_\omega| < \log^3 n$. □

LEMMA 3.3. *Let $G \leq \mathrm{Sym}(\Omega)$ be a primitive group of type (iv)(a) and let $\omega \in \Omega$. Then there exist points $\psi, \chi, \phi \in \Omega$ such that*

(i) *ψ, χ belong to the smallest orbit of G_ω in $\Omega \setminus \{\omega\}$.*

(ii) *ϕ belongs to the smallest orbit of G_χ in $\Omega \setminus \{\chi\}$.*

(iii) *The subgroup $H = \langle G_{\omega\psi}, G_{\chi\phi} \rangle \neq G$, and the action of G on the cosets of any maximal subgroup containing H is either nonfaithful or of degree $\leq n/2$.*

PROOF. By Lemma 3.2, for any i, $1 \leq i \leq m$, there is an element in the smallest G_ω-orbit in $\Omega \setminus \{\omega\}$ that is nontrivial only in the i^{th} coordinate. Let ψ, χ belong to that smallest orbit, with ψ nontrivial in the second coordinate and with χ nontrivial in the first coordinate. There exist $t \in N_1$, $u \in N_2$, such that $\omega^t = \chi$, $\omega^u = \psi$. Let $\phi = \chi^u = \psi^t$ so that ϕ belongs to the smallest orbit of G_χ $(= G_\omega^t)$ in $\Omega \setminus \{\chi\}$. Thus ψ, χ, ϕ satisfy (i) and (ii).

To show that ω, ψ, χ, ϕ satisfy (iii), we first show that $H = \langle G_{\omega\psi}, G_{\chi\phi} \rangle \leq \mathrm{N}_G(N_2) < G$. Suppose $x \in G_{\omega\psi}$. If $N_2^x = N_j$ then $u^x u^{-1} \in N_\omega \cap N_j N_2 = (N_j)_\omega (N_2)_\omega$. Since $u \notin (N_2)_\omega$, we must have $j = 2$. Thus, $G_{\omega\psi}$ normalizes N_2. Similarly for $G_{\chi\phi}$.

Let M be a maximal subgroup of G containing H. Suppose that the action of G on the right cosets of M is faithful. We show that $|G : M| \leq n/2$.

Since $t \in N_1$, $(N_i)_\omega = (N_i)_\chi$ for $i > 1$, but, by Fact 1, $N_\omega \neq N_\chi$; thus, $(N_1)_\omega \neq (N_1)_\chi$. Since u centralizes N_1, $(N_1)_\omega \leq G_{\omega\psi}$; similarly, $(N_1)_\chi \leq G_{\chi\phi}$. We have

$$(N_1)_\omega < \langle (N_1)_\omega, (N_1)_\chi \rangle \leq (N_1)_H \leq (N_1)_M$$

(the latter two being "point" stabilizers in the actions of N_1 on G/H and G/M, respectively). In particular, $|(N_1)_M| \geq 2|(N_1)_\omega|$. But, since G_M acts on $\{N_i\}_{1 \leq i \leq m}$ as does G, we know that $(N_i)_M$ is conjugate under G_M to $(N_1)_M$. (Of course, G_M is just M but it is more useful to emphasize the point-stabilizer interpretation.) Thus $|N_M| \geq (2|(N_1)_\omega|)^m$. Since N is transitive in the (primitive) action of G on G/M,

$$|G : M| = |N : N_M| \leq (|N_1 : (N_1)_\omega|/2)^m = n2^{-m}. \qquad \square$$

The next two lemmas give us the machinery to deal with primitive groups of type (iv)b.

Let $\Omega^{(2)}$ denote the set of unordered pairs of elements of Ω. Then G acts naturally on $\Omega^{(2)}$.

LEMMA 3.4. *Let $G \leq \mathrm{Sym}(\Omega)$ be a primitive group of type (iv)(b) and let $\omega \in \Omega$. Then there is a point $\psi \neq \omega \in \Omega$ such that*

(i) *ψ belongs to a G_ω-orbit of size $< \log^3 n$.*

(ii) *If $\mathcal{B}$ is a minimal block system in the G-orbit of the pair $\{\omega, \psi\} \in \Omega^{(2)}$ then the action of G on $\mathcal{B}$ is either nonfaithful or $|\mathcal{B}| \leq n/2$.*

PROOF. We have $N = \mathrm{Soc}(G) = T_1 \times \ldots \times T_r$, $n = |T_1|^{r-1}$, and $|N_1| = |T_1^{r/m}|$. Therefore $n = |N_1|^{m(1-1/r)}$. On the other hand, $r \geq m \geq 1 + \log n/(2 \log\log n)$, hence $|T_1| \leq |N_1| < \log^2 n$. Also, since $r - 1 < \log_{|T_1|} n$, we have $r < \log n$.

Let $t \in T_1$ be an involution (guaranteed to exist by the Odd Order Theorem[11]) and set $\psi = \omega^t$. To see that (i) holds, note that any $t \in T_1$ has at most $|T_1| r < \log^3 n$ conjugates under G_ω so that the G_ω-orbit of $N_\omega t$ in N/N_ω is also bounded by $\log^3 n$. The result follows now by the G-set identification $N/N_\omega \simeq \Omega$.

Suppose that the action of G on $\mathcal{B}$, as in (ii), is faithful. We show that $|\mathcal{B}| \leq n/2$.

Let b denote the block in $\mathcal{B}$ containing $\{\omega, \psi\}$. Since t fixes $\{\omega, \psi\}$, it fixes b in the action on $\mathcal{B}$. It follows from $(T_1)_b \neq 1$ that $|\mathcal{B}| = |T_1 : (T_1)_b|^r$ (e.g., by consideration of the cases of the O'Nan-Scott Theorem [9]). Since $\sqrt{n} < r!/2$, we have $|T_1| = n^{1/(r-1)} < (r!/2)^{2/(r-1)}$ and the latter is less than 2^{r-1} for all $r \geq 2$. Therefore, $|\mathcal{B}| = |T_1 : (T_1)_b|^r \leq |T_1|^r/2^r = n|T_1|/2^r < n2^{r-1}/2^r$. □

In our application of Lemma 3.4, we need to locate such a ψ without having to search through many points. For this, we have

LEMMA 3.5. *Let $G \leq \mathrm{Sym}(\Omega)$ be a primitive group of type (iv)(b) and let $\omega \in \Omega$. The union of all G_ω-orbits that have $< \log^3 n$ points has cardinality $O^\sim(1)$.*

PROOF. We consider points in Ω as equivalence classes of N, each class of size $|N_\omega| = |\mathrm{Diag}(T_1 \times \ldots \times T_r| = |T_1| \leq \log^2 n$. It suffices then to show that there are only $O^\sim(1)$ elements of N with fewer than $\log^5 n$ conjugates under G_ω, for if an an element has at least $\log^5 n$ conjugates, then the number of equivalence classes among these conjugates is at least $\log^5 n/|N_\omega| \geq \log^3 n$.

Let k^* be the smallest integer such that $\binom{m}{k^*} > \log^5 n$. (For large n, we have $k^* = 6$). Let $u \in N$ be arbitrary and write u in the form $u = u_1 u_2 \cdots u_m$, $u_i \in N_i$. We note that the diagonal subgroup $N_\omega = \mathrm{Diag}(T_1 \times \cdots \times T_r)$ defines a G_ω-invariant identification (specific isomorphism) between the T_i; this in turn determines an identification between the N_i. So we can compare $u_1, u_2, \ldots$. If a particular element of N_1 occurs exactly ℓ times among the u_i where $k^* \leq \ell \leq m - k^*$ then u has at least $\binom{m}{\ell} > \log^5 n$ conjugates, due to the A_m action. If there are at least $k^* + 2$ different values among the u_i then u has at least $m^{k^*} > \log^5 n$ conjugates. Finally, we observe that there are $< \binom{m}{k^*}^{k^*+1} |N_1|^{k^*+1} = O^\sim(1)$ elements $u \in N$ which belong to neither of the above cases. □

We shall also need the following simple lemma.

LEMMA 3.6. *Suppose that the primitive group $G \leq \mathrm{Sym}(\Omega)$ has a regular normal subgroup N, $\omega, \psi \in \Omega$, and let M be a maximal subgroup containing $G_{\omega\psi}$. Then $|G : M| \leq |\Omega|$.*

PROOF. Let $|\Omega| = n$. We know that $|G : G_{\omega\psi}| \leq n(n-1)$ and $G_{\omega\psi} \cap N = 1$. We distinguish two cases:

(i) M contains N. In this case, $|M| \geq |G_{\omega\psi}||N|$, so $|G : M| \leq n - 1$.

(ii) M does not contain N. In this case, since M is maximal, $MN = G$. Hence $|G| = |MN| \leq |M||N|$ and $|G : M| \leq n$. □

4. The Algorithm

We reduce Theorem 1.1 to Theorem 1.3 as follows. Given a primitive G, then by Theorem 2.8, in $O^{\sim}(sn^2)$ time it is possible to detect whether or not $|G| \geq n^{9 \log n \log\log n}$ and by constructing the imprimitive natural action of G, to reduce the composition-series problem to the giant and the small-base case. However, for giant G, we can, by Theorem 2.7, identify G/A_n and construct an SGS for A_n.

Thus, it remains to prove Theorem 1.3.

The core of the algorithm used for the problem in Theorem 1.3 is a procedure that solves the following problem and is especially efficient for small-base groups.

PROBLEM. GENERALIZED PROPER NORMAL SUBGROUP (GPNS)
INPUT: $G = \langle S \rangle \leq \mathrm{Sym}(\Omega)$, $|\Omega| = n$, $|S| = s$.
OUTPUT: *One of the following.*
(1) *The report "G is simple".*
(2) *Generators for a proper normal subgroup of G.*
(3) *A faithful action of G on a domain of size at most n/2.*

Our procedure for GPNS follows the outline of the algorithm given in [20, Section 5]; the novel parts are the handling of primitive groups with a regular normal subgroup (see Step 4) and the improvement in the handling of primitive groups with a unique nonabelian normal subgroup (see Steps 5 and 6).

In most cases, outputs of types (2) and (3) are discovered in induced actions of G. Each constructed action gives rise to a primitive action in which we test the kernel and the size of the domain. The procedure TEST_ACTION formalizes these steps. The input, Δ, is a set of size > 1, on which G acts transitively.

```
procedure TEST_ACTION(Δ)
  begin
    B:= a minimal G-block system on Δ;
    N:= the kernel of the G-action on B;
    if N ≠ 1 then (output N; halt);
    if |B| ≤ |Δ|/2 then (output {G → Sym(B)}; halt)
  end.
```

TEST_ACTION is called at several points in procedure for GPNS. If it does not detect a proper normal subgroup (output of type (2)) or a smaller domain (output of type (3)) then control is passed back to the calling procedure.

The algorithm for the GPNS problem:

procedure GPNS

Step 0. Construct a nonredundant base and an SGS for G

Step 1. $\Psi :=$ any nontrivial orbit of G;
TEST_ACTION(Ψ).

Step 2. **if** $|G| = |\Psi|$ **then** (**output** "G is simple of prime order"; **halt**).

Step 3. $A :=$ any subset of G of size $7\sqrt{n}+1$;
for each pair $\{a, b\} \in A$ **do**
$N := \langle ab^{-1} \rangle^G$;
if $N \neq G$ **then** (**output** N; **halt**).

Step 4. $\omega, \psi :=$ any two points of Ψ;
$H := \mathrm{N}_G(G_{\omega\psi})$;
$M :=$ a maximal subgroup of G containing H;
if $|G : M| \leq n$ **then** TEST_ACTION(G/M).

Step 5. $\omega :=$ any point in Ψ;
for all ψ, χ in smallest G_ω-orbit **do**
for all ϕ in smallest G_χ-orbit **do**
$H := \langle G_{\omega\psi}, G_{\chi\phi} \rangle$;
if $H \neq G$ **then** TEST_ACTION(G/H).

Step 6. $\omega :=$ any point in Ψ;
for all ψ in G_ω-orbits of size $< \log^3 n$ **do**
$\Gamma_\psi :=$ the G-orbit of $\{\omega, \psi\}$ in set of unordered pairs;
TEST_ACTION(Γ_ψ).

Step 7. **output** "G is nonabelian simple".

LEMMA 4.1. *The output of the procedure GPNS is correct.*

PROOF. It is clear that the output is correct if the algorithm encounters a proper normal subgroup or a smaller domain. We have to prove that G is simple if the output says so.

If Step 1 is passed then G acts primitively on Ψ. A primitive regular group is cyclic of prime order, so the output of Step 2 is correct. Suppose that the algorithm passed Step 2; it is sufficient to show that, if it reaches Step 7, then G is simple.

If G falls in case (ii) of Theorem 3.1 then the algorithm halts at Step 3.

We claim, if G falls in case (iii), that the algorithm halts at Step 4. To see this, let N be a regular normal subgroup of G. By Lemma 3.6, $|G : M| \leq n$, so TEST_ACTION will be called. The unique $x \in N$ carrying ω to ψ centralizes $G_{\omega\psi}$, so $x \in M$. This means that N cannot act regularly on the cosets of M; so, if the action on the cosets is faithful then the number of cosets is $\leq |N|/2 = |\Psi|/2$

(Since N is normal in G, it acts transitively in any primitive action of G, whence the size of the domain is the index of a point-stabilizer in N.)

If G falls in case (iv)(a) of Theorem 3.1 then Lemma 3.3 implies that the algorithm halts at Step 5.

If G falls in case (iv)(b) then, by Lemma 3.4, the algorithm halts at Step 6.

So, if Step 7 is reached then G is simple. □

As in Theorems 2.4 and 2.5, we express the timing of GPNS as a function of n and $|G|$. At this point, however, it is convenient to introduce an extension of our "soft O" notation. In considering numerical functions on the pairs (G, n), where $G \leq S_n$, we write $f(G,n) = O^{\approx}(g(G,n))$ if, for some constants $c, c', C > 0$, $f(G,n) \leq Cg(G,n)\log^c n \log^{c'} |G|$. If we were dealing exclusively with small-base groups $G \leq S_n$, the $\log^{c'} |G|$ term would be superfluous. However, we ultimately need to use the GPNS procedure in induced representations of a small-base group. With respect to these, possibly smaller, domains the image of G need not be a small-base group.

LEMMA 4.2. *The procedure GPNS runs in $O^{\approx}(n^3 + sn)$ time.*

PROOF. First, we point out that a call of TEST_ACTION takes $O^{\approx}(n^2)$ time on a set of size $O^{\approx}(n)$ since, by Theorem 2.6, a minimal block system $\mathcal{B}$ can be computed in $O^{\approx}(n)$ time and, by Theorem 2.4, the kernel of the action on $\mathcal{B}$ can be obtained in $O^{\approx}(n^2)$.

By Theorem 2.4, Step 0 runs in $O^{\approx}(n^2 + sn)$ and, as observed above, Step 1 runs in $O^{\approx}(n^2)$. Theorem 2.5 shows that Step 3 takes $O^{\approx}(n^2)$ time.

Step 4 contains the n^3 bottleneck of the algorithm. One can choose ω, ψ to be the first two points of the base returned in Step 0. Let $\{u_i\}$ and $\{v_j\}$ denote the set of coset representatives for G_ω in G and for $G_{\omega\psi}$ in G_ω, respectively and let $S_{\omega\psi}$ be the set of strong generators for $G_{\omega\psi}$. Then $|S_{\omega\psi}| = O^{\approx}(1)$. The subgroup $H = \mathrm{N}_G(G_{\omega\psi})$ can be computed in $O^{\approx}(n^2)$ time since for each of the $< n^2$ coset representatives $v_j u_i$ for $G_{\omega\psi}$ in G, it is enough to trace the image of ω and ψ at the conjugates $y^{v_j u_i}$, for $y \in S_{\omega\psi}$. Computing M can be done by setting $M := H$ and then, for all coset representatives of $G_{\omega\psi}$ in G, computing $M^* = \langle M, v_j u_i \rangle$. If $M^* = G$ then discard $v_j u_i$; otherwise, set $M := M^*$. By Theorem 2.5, one computation of M^* takes $O^{\approx}(n)$ time. Eventually, a maximal subgroup M is obtained in $O^{\approx}(n^3)$ time. The action of each of the $O^{\approx}(1)$ (strong) generators y of G on the cosets of M can be constructed in $O^{\approx}(n^2)$ time since for each coset representative z of M in G, the coset of zy can be found in $O^{\approx}(n)$ time, by testing which of the ratios $zy\bar{z}^{-1}$, $\bar{z} \in G/M$, are in M. Testing one $\bar{z}$ amounts to a sifting through an SGS of M. This can be done in $O^{\approx}(1)$ time, since it is enough to follow the images of the (common) base of G and M.

Step 5 runs in $O^{\approx}(n^2)$ time since, by Lemma 3.2, we have to test only $O^{\approx}(1)$ triples ψ, χ, ϕ. By Theorem 2.5, H can be computed in $O^{\approx}(n)$ time. Since $H = \langle G_{\omega\psi}, G_{\chi\phi} \rangle$ contains $G_{\omega\psi}$, the action of G on the cosets of H is the action

of G on a block system of cosets of $G_{\omega\psi}$. The action of G on the cosets of $G_{\omega\psi}$ is the action of G on $(\omega,\psi)^G$, and can be found in $O^{\approx}(n)$ time since the index of $G_{\omega\psi}$ in G is $O^{\approx}(n)$. The block system corresponding to H can be found in $O^{\approx}(n)$ time since we can compute generators for H. Finally, calling TEST_ACTION on G/H costs $O^{\approx}(n^2)$.

Step 6 also runs in $O^{\approx}(n^2)$ time since there are only $O^{\approx}(1)$ choices for ψ (by Lemma 3.5), for each ψ, the images of the unordered pair $\{\omega,\psi\}$ can be found as a block system of cosets of $G_{\omega\psi}$, and, as above, this block system is obtained in $O^{\approx}(n)$. Calling TEST_ACTION on Γ_ψ costs $O^{\approx}(n^2)$. □

As indicated, we apply the above algorithm to $\phi(G)$ where $\phi : G \to S_m$ is some induced representation of G and always $m = O^{\sim}(n)$. Although not essential for the asymptotic result, it seems most reasonable to carry out such procedure calls "locally", e.g. elementary operations, such as permutation multiplications, are performed only in the m element domain. To facilitate the "lifting" of the answer back to $G \leq S_n$, in an initial construction of the base and SGS for $\phi(G)$, preimages in G of the SGS are kept; for small-base G, doing this via the method of Theorem 2.4 costs $O^{\sim}(nm+sn)$ time. The lifting of elements (e.g., generators of targeted subgroup) is done at a per-element cost of $O^{\sim}(n)$ by sifting through the base of $\phi(G)$, copying the computation in the lifting of the SGS. The complete lifting of a subgroup also includes the kernel of the action (Theorem 2.4(ii)).

We complete, finally, the

PROOF OF THEOREM 1.3. It suffices to indicate how one obtains a maximal normal subgroup M together with a permutation representation for G/M, for this may be repeated $O^{\sim}(1)$ times to obtain a composition series of G. We outline the construction of M, which follows that in [20].

To find a maximal normal subgroup M, we apply the procedure for GPNS to G. If the output of GPNS is of type (3), we rerun GPNS with the induced action. Thus, for nonsimple groups, at most $\log n$ applications of GPNS will produce a proper normal subgroup N (output of type (2)). In such case, we can construct a nontrivial action $\phi : G \to \mathrm{Sym}(\Psi)$ with $N \leq \ker(\phi)$ and $|\Psi| \leq |\Omega|$, namely, take $\Psi = G/G^{(i)}N$, where i is minimal such that $G > G^{(i)}N$ (by Theorems 2.4 and 2.5), this action is constructible in $O^{\sim}(n^2)$ time). Again, we may assume the output is of type (1) or (2). If $\phi(G)$ is simple then $M = \ker(\phi)$ is maximal normal in G and ϕ induces a faithful representation of G/M. Otherwise, if $\bar{N}$ is a proper normal subgroup of $\phi(G)$, we have $N < \phi^{-1}(\bar{N}) < G$. Replace N by $\phi^{-1}(\bar{N})$ and repeat, i.e., construct an action whose kernel includes N, etc. □

Acknowledgment. The authors wish to thank the referees for their effort, their interest, and their suggestions.

REFERENCES

1. M.D. Atkinson, *An algorithm for finding the blocks of a permutation group*, Math. Comp. **29** (1975), pp. 911–913.
2. L. Babai, *On the order of uniprimitive permutation groups*, Ann. of Math. **113** (1981), pp. 553–568.
3. L. Babai, *On the order of doubly transitive permutation groups*, Inventiones Math. **65** (1982), pp. 473-484.
4. L. Babai, G. Cooperman, L. Finkelstein, and Á. Seress, *Nearly linear time algorithms for permutation groups with a small base*, Proc. ISSAC'91 (Internat. Symp. on Symbolic and Algebraic Computation), Bonn 1991, pp. 200–209.
5. L. Babai, E.M. Luks, and Á. Seress, *Fast management of permutation groups*, Proc. 29th IEEE FOCS (1988), pp. 272–282.
6. L. Babai, E.M. Luks, and Á. Seress, *Fast deterministic management of permutation groups*, in preparation.
7. R. Beals, *Constructing blocks of imprimitivity in nearly linear time for small-base groups*, in these Proceedings.
8. R. Beals and Á. Seress, *Structure forest and composition factors for small base groups in nearly linear time*, Proc. 24th ACM STOC (1992), pp. 116–125.
9. P. J. Cameron, *Finite permutation groups and finite simple groups*, Bull. London Math. Soc. **13** (1981), pp. 1–22.
10. B.N. Cooperstein, *Minimal degree for a permutation representation of a classical group*, Israel J. Math. **30** (1978), pp. 213–235.
11. W. Feit and J.G. Thompson, *Solvability of groups of odd order*, Pacific J. Math. **13** (1963), pp. 775–1029.
12. M. Hall, Jr., The Theory of Groups, Macmillan, New York 1959.
13. W.M. Kantor, *Permutation representations of the finite classical groups of small degree or rank*, J. Algebra **60** (1979), pp. 158–168.
14. W.M. Kantor, *Sylow's theorem in polynomial time*, J. Comp. and Syst. Sci. **30** (1985), pp. 359–394.
15. W.M. Kantor and E.M. Luks, *Computing in quotient groups*, Proc. 22nd ACM STOC, 1990, pp. 524–563.
16. D.E. Knuth, *Notes on efficient representation of perm groups*, Combinatorica **11** (1991), pp. 57–68 (preliminary version circulated since 1981).
17. M.W. Liebeck, *On minimal degrees and base sizes of primitive groups*, Arch. Math. **43** (1984), pp. 11–15.
18. M.W. Liebeck, C.E. Praeger, and J. Saxl, *On the O'Nan-Scott theorem for finite primitive permutation groups*, J. Australian Math. Soc. **44** (1988), pp. 389-396.
19. M.W. Liebeck and J. Saxl, *The orders of maximal subgroups of the finite exceptional groups of Lie type*, Proc. London Math. Soc. (3) **55** (1987), pp. 299–330.
20. E. M. Luks, *Computing the composition factors of a permutation group in polynomial time*, Combinatorica **7** (1987), pp. 87–99.
21. E. M. Luks, *Permutation groups and polynomial-time computation*, in these Proceedings.
22. C. E. Praeger and J. Saxl, *On the orders of primitive permutation groups*, Bull. London Math Soc. **12** (1980), pp. 303–307.
23. L. L. Scott, *Representations in characteristic p*, Proc Santa Cruz Conf. on Finite Groups, AMS (1980), pp. 319–322.
24. C.C. Sims, *Computational methods in the study of permutation groups*, in: Computational Problems in Abstract Algebra, J. Leech, ed., Pergamon Press 1970, pp. 169–183.
25. C.C. Sims, *Computation with Permutation Groups*, in Proc. Second Symposium on Symbolic and Algebraic Manipulation, S.R. Petrick, ed., ACM, New York, 1971, pp. 23–28.
26. H. Wielandt, Finite Permutation Groups , Acad. Press, New York 1964.

Department of Computer Science, University of Chicago, Chicago, IL 60637-1504,
Department of Algebra, Eötvös University, Budapest, Hungary H-1088
E-mail address: laci@cs.uchicago.edu

Department of Computer and Information Science, University of Oregon, Eugene, OR 97403
E-mail address: luks@cs.uoregon.edu

Department of Mathematics, The Ohio State University, Columbus, OH 43210
E-mail address: akos@function.mps.ohio-state.edu