ON THE DIAMETER OF PERMUTATION GROUPS

James R. Driscoll *Merrick L. Furst*

Department of Computer Science
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

Abstract. We show that any group represented by generators that are cycles of bounded degree has $O(n^2)$ diameter, *i.e.*, that the longest product of generators required to reach any permutation in the group is $O(n^2)$. We also show how such "short" products can be found in polynomial time. The techniques presented are applicable to generalizations of many permutation–group puzzles such as Alexander's Star and the Hungarian Rings.

1. *Introduction*

One of the important abstractions central to the study of algorithm design and computational complexity is that of an operation on a finite structure. If a structure can be represented in some succinct way as a permutation of the objects of a finite set, and operations are all permutations, then the algorithmic problem of reaching a configuration is the problem of representing a given permutation as a product of a given set of others. Furst, Hopcroft, and Luks presented a polynomial–time algorithm for calculating whether it is possible to express one permutation as a product of others[FHL]. This problem was originally explored by Sims[S], and most recently by Jerrum[J]. Although it is possible to determine in polynomial time if a permutation can be expressed as a product of generators, it can be the case that the length of the shortest product expressing a permutation is exponential in the number of letters being permuted. In this paper we ask: when does a permutation group on n letters, generated by permutations $g_1, \ldots, g_i$ have polynomial diameter, *i.e.*, when is it the case that every element of G can be expressed as a polynomial–length product of the g_i.

It is fairly easy to concoct examples of permutation groups with exponential diameter. Consider the single permutation

$$g = (1\ \ 2)(3\ \ 4\ \ 5)(6\ \ 7\ \ 8\ \ 9\ \ 10)\cdots(\cdots n)$$

in which the length of the i^{th} cycle is the i^{th} prime number. The group $G = \langle g \rangle$ has order

This research is supported by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 3597, monitored by the Air Force Avionics Laboratory Under Contract F33615-81-K-1539.

$p_1 p_2 \cdots p_k$, where p_i is the i^{th} prime. Since the only words in G are of the form g^j, G has elements that are exponentially far from the identity. Is it possible for a group to have exponential diameter without there being a generator of exponential order? Yes. Consider the group D of symmetries of a collection of regular polygons of size $2, 3, \ldots, p_k$. D can be generated by two permutations a and b, where a is the product of k reflections, one for each polygon, and b is also the product of k reflections of the polygons, but along a different axis. Since $a^2 = b^2 = 1$, every word in the group is of the form $(a+b+\epsilon)(ab)^i(a+b+\epsilon)$. Since the order of D is at least $2 \times 3 \times \cdots \times p_k$, there are some permutations that are exponentially far away from 1, even though no generator has order greater than 2.

Actually, in these examples, the groups do not have strictly exponential diameter since the diameter cannot be expressed in the form c^n for some c. It is not known whether a particular group can have a truly exponential diameter. Answering this question may help resolve whether or not combinatorial objects with exponential diameter, like Towers of Hanoi, have small representations as permutation groups.

2. *Preliminaries*

DEFINITION 2.1: The *symmetric group*, S_n, is the group of all permutations on n letters. A permutation in S_n is *even* if it can be represented as a product of an even number of transpositions, *odd* otherwise. The set of even permutations form a group, the *alternating group* A_n. The *degree* of a group is the number of letters it acts on. The *order* of a group is the number of permutations it contains.

DEFINITION 2.2: G_Δ is the set of permutations fixing the letters of Δ. G^Δ is the group restricted to the letters of Δ.

DEFINITION 2.3: If H is a subgroup of G, the quotient G/H is the collection of cosets of H in G (equivalence classes of the elements of G with x equivalent to y if and only if $xy^{-1} \in H$). The *index* of H in G, $[G : H]$, is the number of distinct cosets of G/H.

DEFINITION 2.4: If $g_1, \ldots, g_i$ are elements in S_n, $G = \langle g_1, \ldots, g_i \rangle$ is the group of all permutations representable as a product of the generators $g_1, \ldots, g_i$.

DEFINITION 2.5: A permutation group is *transitive* if for every pair of letters α, β there is some group element p such that $\alpha^p = \beta$; it is *k-transitive* if for every pair of k-tuples $(\alpha_1, \alpha_2, \ldots, \alpha_k)$ and $(\beta_1, \beta_2, \ldots, \beta_k)$ there exists a permutation p such that $\alpha_i^p = \beta_i$. If a group is *intransitive*, it is composed of *transitive constituents*.

DEFINITION 2.6: A *block* B is a set of letters such that $B^p \cap B = B$, or $B^p \cap B = \emptyset$, for every group element p. Sets of single letters and the set of all letters are *trivial blocks*. A group that contains only trivial blocks is *primitive*. Every doubly transitive group is primitive (if it contained a non-trivial block, there would be a permutation that fixed one

block element and sent another out of the block). If B is a block, then $\mathcal{B} = \{B^\pi \mid \pi \in G\}$ is a *block system*. Let $G_\mathcal{B}$ be the group stabilizing a block system (the blocks are left fixed, but letters may move within a block). A block system is *minimal* if $G/G_\mathcal{B}$ is a primitive group. The blocks of a group generated by a cycle correspond to the divisors of its degree. For example, $\{1,4\}$ and $\{1,3,5\}$ are blocks of the group generated by (1 2 3 4 5 6).

DEFINITION 2.7: A member of a group G, defined by generators $g_1, \ldots, g_i$, is *k–expressible* if it can be written (expressed) as a product of $\leq k$ generators. Note that k–expressibility is highly dependent upon the generating set.

DEFINITION 2.8: The *diameter*, $\mathrm{diam}(G)$, of a group $G = \langle g_1, \ldots, g_i \rangle$ is the least integer k such that every $\pi \in G$ is k–expressible. That is, $\mathrm{diam}(G)$ is the diameter, in a graph theoretic sense, of the graph with group elements as vertices, and an edge between u and v if and only if $u\pi = v$ for some generator π.

DEFINITION 2.9: Let H be a subgroup of $G = \langle g_1, \ldots, g_i \rangle$. The diameter of G/H is the least integer k such that every coset of G/H has a k–expressible representative.

LEMMA 2.1: *If H is a subgroup of $G = \langle g_1, \ldots, g_i \rangle$, then*

$$\mathrm{diam}(G) \leq \mathrm{diam}(G/H) + \mathrm{diam}(H).$$

PROOF: Let π be an element of G. There is a permutation $p \in G$ such that $\pi \in Hp$ and p is *diam*(G/H)–expressible. Thus $\pi = hp$ for some $h \in H$, and hence π is $[\mathrm{diam}(G/H) + \mathrm{diam}(H)]$–expressible. □

DEFINITION 2.10: The *intersection graph*, $I(g_1, \ldots, g_i)$, of a set of generators is the graph with generators as vertices and with an edge joining two generators if and only if they act on a common letter.

3. *General Techniques and Observations*

In puzzles like Rubik's Cube, it is easy to send one face to another that it can reach because the face need not pass through any position more than once. By the same argument, every permutation group can be shown to be transitive by $O(n)$ length products. Lemma 3.1 makes the relationship between multiple transitivity and diameter explicit. Lemma 3.2 shows a similar relationship between diameter and primitivity.

LEMMA 3.1: *Suppose $G = \langle g_1, \ldots, g_i \rangle$ is a group of degree n. Let A and B be two k–tuples of letters of G. If S, the set of permutations mapping A to B, is nonempty, then there is a permutation $\pi \in S$ that is $O(n^k)$–expressible.*

PROOF: Let $p_1 \cdots p_j$ be the shortest product of generators that expresses a permuta-

tion in S. No partial products, $p_1 \cdots p_x$ and $p_1 \cdots p_x \cdots p_y$, can have the same effect on the letters of A, as that would imply that $p_1 \cdots p_x p_{y+1} \cdots p_j$ is a shorter product expressing a permutation in S. Since there are only $n!/(n-k)! = O(n^k)$ ways of permuting A over the n letters, j cannot be any larger than $O(n^k)$. □

LEMMA 3.2: *Suppose A is not a block of imprimitivity for $G = \langle g_1, g_1^{-1}, \ldots, g_i, g_i^{-1} \rangle$, a group of degree n. Then there exists a $2n/|A|$–expressible permutation $\pi \in G$ that acts as a witness to the fact that A is not a block, that is, $A^\pi \cap A \neq A$ and $A^\pi \cap A \neq \emptyset$.*

PROOF: Consider the collection of sets $C = \bigcup A^p$, where p runs over every $n/|A|$–expressible permutation. Suppose that for every pair of sets S and T in C, either $S \cap T = S$ or $S \cap T = \emptyset$. If there are $\leq n/|A|$ distinct sets, then every generator of G takes sets of C to sets of C, so the sets are blocks of imprimitivity, a contradiction. Therefore, there are at least $n/|A| + 1$ distinct sets. This is impossible, as it would require there to be $n + |A|$ letters. Thus, there must be two sets A^p and A^q such that $A^p \cap A^q$ is not A^p and is not $\emptyset$, with p and q both $n/|A|$–expressible. Hence, $A \cap A^{qp^{-1}}$ is not A and is not $\emptyset$, and $\pi = qp^{-1}$ is $2n/|A|$–expressible. □

We now argue that, in some sense, the hardest permutations to reach are the 3–cycles. This is a consequence of the following. Any even permutation is the product of $O(n)$ 3–cycles (Lemma 3.3) and every 3–cycle is within $O(n^3)$ of every other (Lemma 3.1).

LEMMA 3.3: *If π is an even permutation of degree n, π can be expressed as a product of no more than $\frac{n}{2}$ 3–cycles.*

PROOF: By induction on n. Obvious for $n = 3$. Let p be a permutation on $n > 3$ letters. If the order of p is not 2 then let α be some letter such that $\alpha^p = \beta$, $\beta^p = \gamma$, and $\alpha \neq \gamma$. Then $p' = (\gamma\ \beta\ \alpha)p$ is an even permutation with degree at most $n-2$ since it leaves β and γ fixed and $(\gamma\ \beta\ \alpha) = (\beta\ \gamma)(\gamma\ \alpha)$ is even. The inductive hypothesis applies to p', so it can be expressed as a product of no more than $(n-2)/2 = \frac{n}{2} - 1$ 3–cycles. Since $p = (\alpha\ \beta\ \gamma)p'$, p can be expressed as a product of no more than $\frac{n}{2}$ 3–cycles. If, on the other hand, the order of p is 2, $p = (\alpha\ \beta)(\gamma\ \delta)p' = (\alpha\ \beta\ \delta)(\gamma\ \delta\ \alpha)p'$, with the degree of $p' = n - 4$. By the inductive hypothesis, p can be expressed as a product of no more than $\frac{n}{2}$ 3–cycles. □

The following two lemmas, which we state without proof, can be found in Wielandt[W].

LEMMA 3.4(Jordan): *A primitive group that contains a 3–cycle is either alternating or symmetric.*

LEMMA 3.5(Marggraf): *Let G be a primitive group of degree n. If there is a transitive subgroup G_Δ with degree $< \frac{n}{2}$, then G is alternating or symmetric.*

In the following theorem we make the key observation that in order to show many groups to have polynomial diameter, it suffices to show that some 3–cycle can be expressed as polynomial–length product.

THEOREM 3.2: *If G is a primitive group containing a polynomially expressible 3–cycle, then the diameter of G is polynomially bounded.*

PROOF: G must be alternating or symmetric by Lemma 3.4, so in particular, G is triply transitive if it is of degree at least 5. (If the degree of G is < 5, the diameter of G is a constant.) Let t be a $q(n)$–expressible 3–cycle, for some fixed polynomial $q(n)$. Any 3–cycle can be expressed as a conjugate $p^{-1}tp$. Since p is defined by its action on three letters, p is $O(n^3)$–expressible by Lemma 3.1, and consequently an arbitrary 3–cycle is $O(q(n)+n^3)$–expressible. Any even permutation can be expressed as a product of at most $\frac{n}{2}$ 3–cycles by Lemma 3.3, so any even permutation is $O(n \cdot q(n)+n^4)$–expressible. If there is any odd generator of G, the odd permutations can be expressed as the product of this generator and an even permutation. Therefore, the diameter of G is $O(n \cdot q(n)+n^4)$. □

EXAMPLE: Consider any group generated by a set of cycles that pairwise intersect in at most one letter. If p and q are two cycles intersecting in exactly one letter, then $pqp^{-1}q^{-1}$ can be seen to be a 3–cycle, so we may conclude by the previous theorem that the diameter of such groups is $O(n^4)$. Alexander's Star (if the orientation of each "block" is disregarded) is such a group. Thus, generalizations of Alexander's Star can be solved in $O(n^4)$ moves.

The following lemma shows that to get to a permutation in a polynomial number of moves, it is good enough to "almost" reach the permutation in a polynomial number of moves. For instance, if a generalized Rubik's Cube could always be solved with the exception of 4 faces in no more than $f(n)$ twists, then the remaining faces can be solved in $O(f(n))$ twists.

LEMMA 3.6: *Let $G = \langle g_1, \ldots, g_i \rangle$, be a group of degree n, with a subgroup H of order $\leq k$, for some constant k. If the diameter of G/H is $f(n)$, then the diameter of G is $O(f(n))$.*

PROOF: Take $p = p_1 \cdots p_j$ to be a product of generators that expresses a permutation of H. We will transform the p_k so that each partial product expresses a permutation of H. Suppose that $p_1 \cdots p_{k-1}$ is in the coset Hc_{k-1} and $p_1 \cdots p_k$ is in the coset Hc_k, where c_{k-1} and c_k are $f(n)$–expressible coset representatives. For each k, replace p_k by $p'_k = c_{k-1}p_kc_k^{-1}$ to get a new product p'. We now have $p = p'$, and each partial product of p' expresses an element of H. Suppose some pair of partial products $p'_1 \cdots p'_x$ and $p'_1 \cdots p'_x \cdots p'_y$ express the same permutation of H, then $p'_1 \cdots p'_x p'_{y+1} \cdots p'_j$ also represents the same permutation but is a shorter product. By an induction on the number of such

equal partial products there is an equivalent product with length at most order$(H) \leq k$. Since each permutation in the sequence is $O(f(n))$–expressible, the diameter of H is $O(f(n))$. Every permutation of G is in some coset Hx, where x is a coset representative that is $f(n)$–expressible, therefore the diameter of G is $O(f(n))$. □

It is sometimes convenient to think of a generating set in terms of its intersection graph. Lemma 3.7 shows that the connected components of the intersection graph are just the transitive constituents of the group.

LEMMA 3.7: *Let $c_1, \ldots, c_i$ be cycles. The intersection graph $I = I(c_1, \ldots, c_i)$ is connected if and only if the group $G = \langle c_1, \ldots, c_i \rangle$ is transitive.*

PROOF: Suppose I is connected. Let α and β be any pair of letters acted on by the generators x_0 and x_j, respectively, and let $x_0 x_1 \cdots x_j$ be a path from x_0 to x_j in I. By an induction on the length of the path, there is a product of the form $x = x_0^{k_0} x_1^{k_1} \cdots x_j^{k_j}$ with $\alpha^x = \beta$. Conversely, suppose G is transitive. Let a and b be any pair of generators that act on α and β, respectively, and let p be a permutation such that $\alpha^p = \beta$. If $x_{k_0} x_{k_1} \cdots x_{k_j}$ is the shortest product expressing p, then it is a path in I. □

4. *Groups Generated by Cycles of Bounded Degree*

It is easy to construct examples of groups that have super–polynomial diameter. What restrictions can be placed upon generators to ensure that diameter remains small? We conjecture that if generators are constrained to operate on no more than a constant number of letters, then the group they generate has polynomial diameter. At present, the best we can show is if generators act on a constant number of letters and are cycles, then the group they generate has diameter $O(n^2)$. (This is the best one can hope for, as there exist groups generated by constant–sized cycles that have diameter $O(n^2)$, Theorem 4.10.) The techniques used to prove this result may be help in proving the conjecture.

DEFINITION 4.1: A permutation π is *d–cyclic* if π is a cycle of order (degree) $\leq d$. Throughout this section, d is assumed to be a constant.

If a group G is generated by d–cyclic permutations, then the transitive constituents of G are generated by subsets of the generators. Therefore, the diameter of G is bounded by the sum of the diameters of the constituents. From now on we assume that G is transitive. The primitive and imprimitive cases are handled separately.

4.1 *Primitive Groups with d–Cyclic Generators*

Suppose G is primitive and has d–cyclic generators. Any constant–size subset of the

generators can be used to generate a transitive subgroup of G by building $O(n)$–expressible conjugates of the cycles that act on a common letter. The diameter of such a subgroup is $O(n)$, since any permutation can be written as a constant–length product of the $O(n)$–expressible conjugates. If the subgroup is also primitive and of degree $> 2d+1$, then it contains an $O(n)$–expressible 3–cycle since it is alternating or symmetric (Lemma 3.5). Theorem 4.1 shows that because only a constant number of the generators are required to ensure the primitivity of the group, such a subgroup can always be found. The other generators are only required to maintain transitivity. This result and Theorem 3.2 imply an $O(n^4)$ bound on the diameter of G. This bound many then be improved to $O(n^2)$ by using the same special subgroup to construct arbitrary $O(n)$–expressible 3–cycles.

THEOREM 4.1: *Let the generators of $G = \langle c_1, \ldots, c_i \rangle$ be d–cyclic. If G is primitive of degree at least $2d+1$ then G contains a subgroup H such that:*

(i) $2d+1 \leq \text{degree}(H) \leq d^2$,
(ii) *H is alternating or symmetric, and*
(iii) *H has $O(n)$ diameter.*

PROOF: Consider the intersection graph $I = I(c_1 \ldots, c_i)$. Since G is transitive this graph must be connected by Lemma 3.7. Without loss of generality, assume that the cycle c_1 is not properly contained in any other cycle. Let T be a spanning tree of I, rooted at c_1. The procedure we are about to describe successively removes cycles from the generating set T and adds conjugates of some cycles to S such that $H = \langle T \cup S \rangle$ is always primitive and of degree $> 2d$. When the procedure terminates the set S will consist of cycles that are "special" in the sense that they play a central role in what makes G primitive.

```
set S = ∅;
while degree(⟨T ∪ S⟩) > 3d and T ≠ {} do
    let c_j be a leaf of T
    set T ← T − {c_j}
    if ⟨T ∪ S⟩ is imprimitive then
        let a be a letter on the root cycle c_1
        let b be a letter on the cycle c_j
        let π be O(n)–expressible from the
              members of T such that b^π = a
        set S ← S ∪ {π^{-1} c_j π}
```

We claim that the above procedure terminates, and that when it does the group $H = \langle T \cup S \rangle$ is primitive of degree between $2d+1$ and d^2. Termination follows from the fact that in each iteration the cardinality of T decreases by one. The group $\langle T \cup S \rangle$ is transitive at every stage because every cycle of S intersects with the root of T, so the removal of

a leaf c_j of T leaves $\langle T \cup S \rangle$ transitive. If removing the leaf leaves $\langle T \cup S \rangle$ imprimitive, then adding the conjugate $\pi^{-1} c_j \pi$ to S restores primitivity since it is becomes the same group as before the leaf was removed. Since the removal of a cycle decreases the degree of $\langle T \cup S \rangle$ by at most $d-1$, the degree of $\langle T \cup S \rangle$ is $> 2d$. Thus, $\langle T \cup S \rangle$ is always primitive and of degree $> 2d$.

We now argue that the degree of $\langle T \cup S \rangle$ is $\leq d^2$ because the number of cycles that can be added to the set S is at most d. To see that $|S| \leq d$, we must first understand the nature of the blocks that are created when a cycle is removed from T. Let B be a block of size r, and let a be some d–cyclic generator not contained in any other. Suppose that a acts on some, but not all, of the letters of B. If a also acts on some letters not in B, then a^k, for some k, is a witness to the fact that B is not a block. If, on the other hand, a acts exclusively on letters of B, then there is some other cycle b that acts on letters of B and letters not of B, or else B contains all the letters. The cycle b does not act on all of the letters of B, or a would be contained in b, so b^k, for some k, is a witness that shows b is not a block. Thus, the block must be wholly contained in a. Moreover, the block is also a block of a because a must stabilize the block and the block cannot be the entire cycle since any disjoint intersecting cycle is a witness that shows it is not a block (or a acts on all of the letters). As B is a block of a, the size of the block, r, must divide the degree of a. (Recall that the blocks of cycles correspond to the divisors of the degree.)

Now, suppose the removal of some cycle creates blocks of size r. It must be the case that r divides the degree of c_1. After the conjugate is added, the removal of any other cycle cannot again cause the blocks to be of size r, since the remainder stabilize all the blocks of that size. As there are certainly no more than d divisors of an integer $\leq d$ (actually, it is $O(d^\epsilon)$, for any $\epsilon > 0$ [C]), the cardinality of S is at most d and therefore the degree of $\langle S \rangle$ is at most d^2.

The conclusion is that the group G contains a primitive subgroup H ($= \langle T \cup S \rangle$ at termination) with degree between $2d+1$ and d^2. From Lemma 3.5 we know that H is either alternating or symmetric since it has degree at least $2d+1$ and there is a cycle acting on no more than d letters. Furthermore, the degree of H is a constant, so any permutation in H can be expressed as a constant length product of cycles from $T \cup S$. As the cycles in S can each be written as $O(n)$–length products using the cycles $c_1, \ldots, c_j$, the diameter of H is $O(n)$. □

The above theorem can be strengthened to allow a more careful specification of the size of H.

COROLLARY 4.1: *If G is primitive and generated by d–cyclic permutations, then, for any constant k with $d^2 \leq k \leq n-d$, a subgroup H satisfying* (ii) *and* (iii) *above can be found with degree between k and $k+d$.*

The algorithm of Theorem 4.1 requires an efficient primitivity test, which can be devised with the same idea used in the proof of Lemma 3.2. Efficient ($O(n^2)$) primitivity

tests are known (see, for example, [J]). However, for the special case of $\mathcal{G}$ consisting of d-cyclic permutations, the following algorithm runs in $O(|\mathcal{G}|) + O(n)$ time.

ALGORITHM 4.1: Given a set $\mathcal{G}$ of cycles, determine whether $G = \langle \mathcal{G} \rangle$ is primitive or imprimitive.

1. Create a graph having the letters $\{1, \ldots, n\}$ as vertices. Let there be an edge between two vertices i and j if there is a generator $g \in \mathcal{G}$ such that $i^g = j$. Label every edge with such a generator.
2. Find a cycle c such that the set of letters acted on by c is not properly contained in the set of letters acted on by another cycle, and pick some letter α in c. Guess a block S of $\langle c \rangle$.
3. Do a depth first search starting at α. When a new vertex ω is reached, compute $S_\omega = S^{p_\omega}$, where p_ω is the product of the generators along the direct path from α to ω. Label the vertex ω with S_ω.
4. If, for some vertex x, there is a vertex $y \in S_x$ such that $S_x \neq S_y$ (that is, S is not a block), then go to Step 6.
5. If there is a generator $g \in \mathcal{G}$ for which it is the case that $S_x^g \neq S_y$, for an x acted on by g, and a $y \in S_x^g$, then go to Step 6. Otherwise quit, as S is a block and G is imprimitive.
6. Guess another block S of $\langle c \rangle$, and go back to Step 3. If all the blocks of $\langle c \rangle$ have been checked, and none is a block, then quit and conclude that G is primitive.

THEOREM 4.2: *Algorithm 4.1 correctly determines whether a transitive group G of degree n generated by $\mathcal{G}$, a set of d-cyclic permutations, is primitive in $O(|\mathcal{G}|) + O(n)$ time.*

PROOF: Each set S_ω is a candidate for being a block, as $S_\omega = S^{p_\omega}$. Furthermore $\omega \in S_\omega$, since $\alpha^{p_\omega} = \omega$. Since every letter ω is in some set S_ω, S is not a block unless the sets partition the letters (checked in Step 4). If every generator acts on a partition of the elements as blocks, then S must be a block (Step 5). Suppose G is imprimitive. From the proof of Theorem 4.1, some block of G is also a block of $\langle c \rangle$. Therefore, the algorithm correctly determines if G is primitive, as it must find a block if G is imprimitive.

Steps 1 and 2 can be performed in time $O(|\mathcal{G}|)$. The depth first search (Step 3) takes time $O(|\mathcal{G}|) + O(n)$, since there are $O(|\mathcal{G}|)$ edges, and S^{p_ω} can be computed in constant time at each vertex. Step 4 can be done in $O(n)$ time since there are n constant size sets that have to be compared with a constant number of other sets. Step 5 can be done in $O(|\mathcal{G}|)$ time, as each generator acts on a constant number of letters. Since there are at most a constant number of blocks of $\langle c \rangle$ (recall that c is of bounded degree), Steps 3–6 are performed a constant number of times. Therefore, the algorithm terminates in $O(|\mathcal{G}|) + O(n)$ time. □

COROLLARY 4.2: *If α and β are two letters of a transitive group $G = \langle \mathcal{G} \rangle$ and the generators are d–cyclic, then a product of length $O(n)$ that sends α to β can be found in $O(|\mathcal{G}|) + O(n)$ time.*

PROOF: The permutation p_β of Algorithm 4.1 is the product of $O(n)$ generators, and $\alpha^{p_\beta} = \beta$. □

Since the number of generators can be quite large compared to n, the algorithm of Theorem 4.1 (or Corollary 4.1) could require as many as $O(n^d)$ iterations to terminate. Reducing the number of generators is simplified by the fact that G is either alternating or symmetric if $n \geq 2d + 1$ (Lemma 3.5). Any subset T of the generators that generates a primitive group over all the letters is alternating or symmetric, so the problem is reduced to finding a small set T, to which is added an odd generator of G (if it exists). The following algorithm reduces the size of the generating set to $O(n)$.

ALGORITHM 4.2: Given a set $\mathcal{G}$ of d–cyclic permutations that generates a primitive group G with degree $n \geq 2d + 1$, find a set $T \subset \mathcal{G}$ such that $\langle T \rangle = \langle \mathcal{G} \rangle$ and $|T| = O(n)$.

In the algorithm, C_ω is a letter indicating which transitive constituent ω is in. Initially, $C_\omega = \omega$, but as generators are added to T the constituents are consolidated. If $C_\alpha = C_\beta$, then there is a permutation $p \in \langle T \rangle$ with $\alpha^p = \beta$.

1. Set $C_\omega \leftarrow \omega$ for each letter ω, and set $T \leftarrow \emptyset$.
2. Pick a generator g from $\mathcal{G}$.
3. For each pair of letters α,β acted on by g, if $C_\alpha \neq C_\beta$ then set $T \leftarrow T + g$ and set $C_\omega = C_\alpha$ for all ω with $C_\omega = C_\beta$. (Combine the constituents containing α and β.)
4. Pick another generator from $\mathcal{G}$ and return to Step 3. If there are no more generators, then continue with Step 5.
5. Check if $\langle T \rangle$ is primitive. If it isn't, add to T a cycle from $\mathcal{G} - T$ that does not stabilize the blocks of $\langle T \rangle$. Repeat until $\langle T \rangle$ is primitive.
6. If there is an odd generator in $\mathcal{G}$, add it to T.

THEOREM 4.3: *Algorithm 4.2 terminates in $O(|\mathcal{G}|) + O(n^2)$ time. On termination $\langle T \rangle = G$ and $|T| = O(n)$.*

PROOF: Initially, each of the n letters is in its own constituent. A generator is added to T in Step 3 if it reduces the number of constituents, so $O(n)$ generators can be added in this way. Thus, steps 2–4 insure that $\langle T \rangle$ is transitive and that $|T| = O(n)$. Step 5 insures that $\langle T \rangle$ is primitive. From the proof of Theorem 4.1, we know that no more than a constant number of generators will be added here, so $|T| = O(n)$. Step 6 insures that $\langle T \rangle = G$, as G is alternating or symmetric.

Step 1 requires $O(n)$ time. Step 3 is repeated $O(|\mathcal{G}|)$ times. Since $O(n)$ generators are

added to T, it takes a total of $O(n^2)$ time to combine the various constituents. Thus, steps 2–4 require $O(|\mathcal{G}|) + O(n^2)$ time. The number of times Step 5 is repeated is bounded by a constant, so by Theorem 4.2, Step 5 takes $O(|\mathcal{G}|) + O(n)$ time. Finally, Step 6 requires $O(|\mathcal{G}|)$ time. Therefore the algorithm terminates in $O(|\mathcal{G}|) + O(n^2)$ time. □

The preceding algorithms provide an asymptotically efficient algorithm for finding generators for the subgroup H of Theorem 4.1.

THEOREM 4.4: *If $G = \langle \mathcal{G} \rangle$ satisfies the conditions of Theorem 4.1, generators for the subgroup H of Theorem 4.1 (or Corollary 4.1) can be found in $O(|\mathcal{G}|) + O(n^2)$ time.*

PROOF: First use Algorithm 4.2 to reduce the number of generators to $O(n)$. Next apply the algorithm of Theorem 4.1 to find generators for H. The intersection graph can be found in $O(n^2)$ time, and the leaves can be removed by a depth first search in $O(n^2)$ time. Since there are $O(n)$ generators, the loop is repeated $O(n)$ times. By Theorem 4.2 and Corollary 4.2, every operation in the loop can be computed in $O(n)$ time. Therefore, the algorithm takes $O(|\mathcal{G}|) + O(n^2)$ time. □

The following lemma shows how to improve the bound of Lemma 3.1 when G is a primitive group generated by d–cyclic permutations.

LEMMA 4.1: *Let k be a constant. Suppose G is a primitive group of degree n, generated by d–cyclic permutations. If $S \subset G$ is the set of permutations mapping a k–tuple A onto a k–tuple B and S is nonempty, then there is a permutation $\pi \in S$ that is $O(n)$–expressible.*

PROOF: Assume G has order at least $2(k+d)$; otherwise the degree (and hence the diameter) of G is a constant. Let H be the subgroup of Corollary 4.1 with degree at least $2k+d$. First we will show that there is an $O(n)$–expressible permutation p such that A^p and B^p are letters of H. Let x be a letter of A or B. Find a minimal path in the intersection graph from a generator containing x to a generator containing a letter of H. Let X be the transitive group generated by the cycles along the path. By Lemma 3.1, there is an $O(n)$–expressible permutation $p_x \in X$ that sends x to a letter of H. X acts on no more than d letters of H (or the path was not minimal), and thus leaves at least $2k$ letters of H fixed. Suppose we wish to leave $2k$ distinguished letters of H fixed. There is some $O(n)$–expressible permutation $q \in H$ that sends the $2k$ distinguished letters to the $2k$ letters left fixed by X. Thus $p'_x = q p_x q^{-1}$ is $O(n)$–expressible, sends x to a letter of H, and leaves the $2k$ distinguished letters fixed. For each x in A or B let the $< 2k$ distinguished letters be those letters of A and B already sent to H. The product of these permutations $p = \prod p'_x$, with x running over the letters of A and B, sends the $2k$ letters of A and B to the letters of H and is $O(n)$–expressible. Let $h \in H$ map A^p to B^p. Then $\pi = php^{-1}$ maps A to B, and is $O(n)$–expressible. □

The following two theorems show that a primitive group G generated by d–cyclic

permutations has $O(n^2)$ diameter and that a product of $O(n^2)$ generators expressing any permutation of G can be found in $O(n^2)$ time.

THEOREM 4.5: *Let G be generated by d-cyclic permutations. If G is primitive with degree n, then the diameter of G is $O(n^2)$.*

PROOF: Assume the order of G is at least $2d+1$. Let H be the subgroup of Theorem 4.1. Since H is alternating, H contains a 3–cycle t that is $O(n)$–expressible. Any three letters can be sent to the 3–cycle by a permutation p expressible in $O(n)$ by Lemma 4.1. Thus an arbitrary 3–cycle can be expressed as a conjugate $p^{-1}tp$ that is $O(n)$–expressible. Since $O(n)$ 3–cycles and an odd generator can express any permutation of G, the diameter of g is $O(n^2)$. □

THEOREM 4.6: *Let $G = \langle \mathcal{G} \rangle$ be a primitive group of degree n, and let $\mathcal{G}$ be a set of d–cyclic permutations. If $p \in G$, then a product of $O(n^2)$ generators that expresses p can be found in $O(|\mathcal{G}|) + O(n^2)$ time.*

PROOF: Use the method of Theorem 4.5 and Lemma 4.1. Finding $O(n)$–length products for elements of H can be done in constant time, since the degree of H is a constant. This could result in products with lengths exponential in the degree of H, but a representation of the product can be found in time polynomial in the degree of H by using the algorithms of [FHL]. □

4.2. *Imprimitive Groups with d–Cyclic Generators*

We now consider the imprimitive case. Let G be an imprimitive group generated by d–cyclic permutations, and let $G_{\mathcal{B}}$ be the subgroup stabilizing a minimal block system. The problem of reaching some permutation of G can be decomposed into two parts: arranging the blocks and arranging the elements within blocks. In the case of Rubik's Cube, this corresponds to arranging the physical blocks of the cube, and then arranging the faces of the blocks. The problem of arranging the blocks is equivalent to reaching a particular coset of $G/G_{\mathcal{B}}$. Below, Lemma 4.2 shows that the diameter of $G/G_{\mathcal{B}}$ is $O(n^2)$, so we have

$$\text{diam}(G) \leq \text{diam}(G/G_{\mathcal{B}}) + \text{diam}(G_{\mathcal{B}}) \leq O(n^2) + \text{diam}(G_{\mathcal{B}}).$$

Arranging the elements within the blocks after the position of the blocks has been settled is a bit tricky. First we observe in Lemma 4.3 that any single block can be arranged by an $O(n)$–expressible permutation that leaves fixed all but k distinguished blocks, for some constant k determined by d. In the proof of Theorem 4.8 we show that all of the $O(n)$ blocks, with the exception of the k distinguished blocks, can be arranged by an $O(n^2)$–expressible permutation. Thus, if K is the group fixing all but the k blocks, then the diameter of $G_{\mathcal{B}}/K$ is $O(n^2)$. Since the order of K is a constant, the diameter of K is

$O(n^2)$ by Lemma 3.6. Therefore:

$$\text{diam}(G) \leq \text{diam}(G/G_{\mathcal{B}}) + \text{diam}(G_{\mathcal{B}}/K) + \text{diam}(K),$$

which we show to be $O(n^2)$.

LEMMA 4.2: *Let G be an imprimitive group generated by d-cyclic permutations. If $G_{\mathcal{B}}$ is the group stabilizing a minimal block system, then the diameter of $G/G_{\mathcal{B}}$ is $O(n^2)$.*

PROOF: $G_{\mathcal{B}}$ is normal and stabilizes a minimal block system, so $G/G_{\mathcal{B}}$ is a primitive group. The natural homomorphism from G to $G/G_{\mathcal{B}}$ maps cycles of G to cycles of $G/G_{\mathcal{B}}$. Since the generating cycles of G have degree $\leq d$, the diameter of $G/G_{\mathcal{B}}$ is $O(n^2)$ by Theorem 4.5. □

LEMMA 4.3: *Let G be an imprimitive group generated by d-cyclic permutations. Take B to be a block of a minimal block system of G, G^B to be the group restricted to the letters of B, and S to be a set of k distinguished blocks, with k a constant depending on d. For every permutation of G^B, there is a corresponding permutation of G that is $O(n)$-expressible and that acts only on B and the blocks of S.*

PROOF: The blocks contain a constant number of letters, b, since the size of a block divides the size of a containing cycle, which is bounded by a constant d. Let $p = p_1 \cdots p_j$ be the shortest product of generators expressing a permutation of G^B. As the partial products of p send B to one of the $O(n)$ blocks in one of at most $b!$ arrangements, p and every permutation of G^B can be expressed as a product no longer than $b! \cdot n$ generators.

While p is $O(n)$-expressible, it may act on an arbitrary number of blocks. However, it is possible to transform p to another $O(n)$-expressible permutation that is the same permutation of G^B, but only acts on the blocks of S. Let $q = p_1 \cdots p_k$ be the largest partial product that leaves fixed at least d/b blocks (one cycle of blocks), but any smaller partial product leaves fixed fewer than d/b blocks. The length of q is at least $(n/b)/(d/b) = n/d$, the total number of blocks divided by the largest number of blocks any generator can act on. Let r be a permutation in a coset of $G/G_{\mathcal{B}}$ that sends the block B^q and some of the blocks left fixed by q to the letters of a cycle c that is a generator of G. The conjugate $c_r = rcr^{-1}$ is a cycle acting on B^q and blocks left fixed by q. The permutation r, and hence c_r, is $O(n)$-expressible by Lemma 4.1.

Now consider the permutation $q' = qc_rq^{-1}c_r^{-1}$. It is the same permutation with respect to B, since qc_r takes B to a permutation left fixed by q^{-1}. Since qc_rq^{-1} is the cycle c_r with one block replaced by B, the product of this cycle and c_r^{-1} acts on B and two other blocks, leaving everything else fixed. This is because q and c_r intersect in exactly one block, so the commutator is a 3-cycle with respect to the blocks. Since all of the components of q' are $O(n)$-expressible, q' is $O(n)$-expressible.

In this way, products of length at least n/d are converted to an equivalent product with respect to B acting only on B and two other blocks and having $O(n)$ terms. Since p

has at most $(n \cdot b!)/(n/d) = d \cdot b!$ products of length n/d, p can be converted to a product p' of no more than $d \cdot b!$ such products that have $O(n)$ length and act on only two additional blocks. Thus p' is the same permutation of G^B as p, and it acts on no more than $2d \cdot b!$ blocks. As there are at most $b!$ permutations of G^B, the total number of blocks acted on by any such permutation of G^B is $(2d \cdot b!) \cdot b!$, which is large, but a constant.

Let S be a set of at least $k = 2d \cdot (b!)^2$ blocks. By Lemma 4.1 there is an $O(n)$–expressible permutation π that fixes the block B and takes S to the blocks disturbed by any p'. The permutation π may permute the letters of B, but this permutation π_B is in G^B, so there is another permutation $p'' = \pi_B^{-1} p' \pi_B$ in G^B such that $\pi p'' \pi^{-1}$ is the desired permutation of G^B, acts only on the blocks of S, and is $O(n)$–expressible. □

Actually, for Lemma 4.3 to apply, the generators need not be restricted to single cycles. What is important is that the quotient group has $O(n^2)$ diameter and that the blocks are of constant size. While Alexander's Star has generators that are products of cycles, the quotient group is generated by constant degree cycles and the blocks are of size 2. Thus, a block of a generalized Alexander's Star can be arranged with an $O(n)$ length product.

ALGORITHM 4.3: Given a set $\mathcal{G}$ of d–cyclic permutations generating an imprimitive group G, a block B of a minimal block system, and a set S as in Lemma 4.3, find a product of length $O(n)$ for each permutation of G^B that acts only on B and S.

1. Use steps 1–4 of Algorithm 4.3 to find a set T such that $|T| = O(n)$ and $\langle T \rangle$ is transitive on the letters of G.
2. Construct a table of the possible permutations of the letters of B (that is, a representative from each of the cosets of G/G_B). This can be done by multiplying each permutation p in the table by every element of T that acts on a letter of B^p. If this is repeated $n \cdot b!$ times, then every permutation will be in the table (since $\text{diam}(G/G_B) \leq [G : G_B] \leq n \cdot b!$). If, when a new permutation is added to the table, both the element of T and the permutation already in the table that formed the new product are added along with it, then the actual product can be recovered.
3. If the product of a generator $g \in \mathcal{G}$ and a permutation in the table is not also in the table, then add g to T and return to Step 2.
4. For each of the fewer than $b!$ elements of G^B, convert the product in the table to one that acts only on B and S using the method of Lemma 4.3.

THEOREM 4.7: *Algorithm 4.3 is correct and terminates in* $O(|\mathcal{G}|) + O(n^2)$ *time.*

PROOF: Step 1 insures that $\langle T \rangle$ is transitive and $|T| = O(n)$ in $O(|\mathcal{G}|) + O(n^2)$ time (Theorem 4.3). Step 2 computes all of the permutations of $\langle T \rangle^B$, and takes $O(n^2)$ time, since there are $O(n)$ generators in T. At most $b!$ generators are added in Step 3, since each new generator must increase the order of $\langle T \rangle^B$, which is at most $b!$. Thus, Step 2

is repeated a constant number of times. Since $\langle T \rangle^B = G^B$ after Step 3, Step 4 finds the desired products for every permutation of G^B. □

THEOREM 4.8: *Let G be an imprimitive group generated by d-cyclic permutations. The diameter of G is $O(n^2)$.*

PROOF: Let $\mathcal{B}$ be a minimal block system. By Lemma 4.2 the diameter of $G/G_{\mathcal{B}}$ is $O(n^2)$, so the diameter of G is $O(n^2) + \text{diam}(G_{\mathcal{B}})$. Let $G_{\overline{S}}$ be the group fixing all but the set S of Lemma 4.3. The diameter of $G_{\mathcal{B}}/G_{\overline{S}}$ is $O(n^2)$, since each of the $O(n)$ blocks not in S can be arranged independently with products $O(n)$ long. Since $G_{\overline{S}}$ has constant order (the size of S is a constant), the diameter of $G_{\mathcal{B}}$ is $O(n^2)$ by Lemma 3.6. Therefore, the diameter of G is $O(n^2)$. □

THEOREM 4.9: *Let $G = \langle \mathcal{G} \rangle$ be an imprimitive group of degree n, and let $\mathcal{G}$ be a set of d-cyclic permutations. If $p \in G$, then a product of $O(n^2)$ generators that expresses p can be found in $O(|\mathcal{G}| \cdot n^2) + O(n^6)$.*

PROOF: The method of Lemma 3.6 cannot be used to achieve a polynomial-time solution, since there are an exponential number of representatives. However, this can be done in polynomial time using the sift and close method [FHL] as a basis. Each of $O(n^2)$ coset representatives can be found in $O(n^2)$ time using the method of Theorem 4.8, except for a constant number of cosets corresponding to each of the permutations of $G_{\overline{S}}$. Each of the generators of $\mathcal{G}$ are sifted into the table, and the table is closed to compute a set of strong generators. (Note that sifting will not increase the length of the product to $O(n^3)$, because the permutations reached by sifting are also $O(n^2)$-expressible.) Since a constant number of representatives are added to the table, the length of the representatives that are the permutations of $G_{\overline{S}}$ is $O(n^2)$. Then, using the method of Theorem 4.8, a product of $O(n^2)$ generators expressing any permutation in G can be found in $O(|\mathcal{G}| \cdot n^2) + O(n^6)$, which is essentially the time to sift $|\mathcal{G}|$ generators and close the table. □

Both the primitive and imprimitive case give an $O(n^2)$ diameter bound, which is also the best bound possible.

THEOREM 4.10: *If G is a group of order n generated by d-cyclic permutations, the diameter of G is $O(n^2)$. For some groups, this is the best possible bound.*

PROOF: The diameter of G is the sum of the diameters of the transitive constituents. If a constituent has order m, then its diameter is $O(m^2)$ by Theorem 4.5 or Theorem 4.8. Therefore the diameter of G is $O(n^2)$. This bound is the best possible, because the group $\langle (1\ \ 2), (2\ \ 3), \ldots, (n-1\ \ n) \rangle$ has $O(n^2)$ diameter. The permuation $(1\ \ n)(2\ \ n-1)\cdots(\lfloor n/2 \rfloor\ \ \lceil n/2 \rceil)$ requires a product of $O(n^2)$ generators, since it has $O(n^2)$ inversions and each generator increases or decreases the number of inversions by one. □

References.

[C] K. Chandrasekharan, "Introduction to Analytic Number Theory." Springer–Verlag, Berlin, 1968.

[FHL] M. Furst, J. Hopcroft, E. Luks, "Polynomial–Time Algorithms for Permutation Groups." 21st *Symposium on the Foundations of Computer Science*, 1980, pp. 36-41.

[J] M. Jerrum, "A compact representation for permutation groups." 23rd *Symposium on the Foundations of Computer Science*, 1982, pp. 126-133.

[S] C. Sims, "Computational methods in the study of permutation groups." In Computational Problems in Abstract Algebra, J Leech (ed.). Pergamon Press, 1970.

[W] H. Wielandt, "Finite Permutation Groups." Academic Press, New York, 1964.