2009

Fast Matrix Multiplication via Group Actions

Hendrik Orem
Harvey Mudd College

Fast Matrix Multiplication via Group Actions

Hendrik Orem

Michael Orrison, Advisor

Nicholas Pippenger, Reader

May, 2009

HARVEY MUDD
COLLEGE

Department of Mathematics

Abstract

Recent work has shown that fast matrix multiplication algorithms can be constructed by embedding the two input matrices into a group algebra, applying a generalized discrete Fourier transform, and performing the multiplication in the Fourier basis. Developing an embedding that yields a matrix multiplication algorithm with running time faster than naive matrix multiplication leads to interesting combinatorial problems in group theory. The crux of such an embedding, after a group G has been chosen, lies in finding a triple of subsets of G that satisfy a certain algebraic relation. I show how the process of finding such subsets can in some cases be greatly simplified by considering the action of the group G on an appropriate set X. In particular, I focus on groups acting on regularly branching trees.

Contents

Acknowledgments

First and foremost, I would like to thank my thesis advisor, Michael Orrison, for his support and guidance in the process of creating this thesis. My second reader, Nicholas Pippenger, has also provided substantial assistance. I am greatly indebted to Richard Bowen, Bo Chen, and Martijn van Schaardenburg for collaboration on some of the writing presented in the Background section, as well as for helpful conversations during my thesis project. Sam Just provided some programming assistance, and Claire Connelly has been extremely helpful with computing support.

Chapter 1

Introduction and Background

Group-theoretic matrix multiplication is a recent development in the history of fast matrix multiplication algorithms. The central challenge arising from this new approach consists of finding three subsets of a given group G which satisfy a certain algebraic relation and which are "large" relative to the character degrees of G. This chapter introduces a new method for finding such subsets, namely the use of group actions, and provides the background necessary for reading this thesis.

1.1 Introduction

This section will outline the history of fast matrix multiplication, introduce the group-theoretic approach, and give a survey of this document's results in the applications of group actions to fast matrix multiplication algorithms.

1.1.1 The Exponent of Matrix Multiplication

Computing the product of two $n \times n$ matrices is a problem of tremendous importance: computing determinants, solving systems of linear equations, and many other linear algebraic problems can be solved using an algorithm for matrix multiplication. Therefore determining the smallest number ω such that two $n \times n$ matrices over a field of characteristic zero can be multiplied in at most $O(n^{\omega})$ field operations has been an active area of research in theoretical computer science for the past forty years. It is known that ω is at least 2, which can be thought of intuitively as being due to the n^2 entries in the product matrix (Bürgisser et al., 1997).

The standard matrix multiplication algorithm computes the product in $2n^3$ field operations, and hence shows that ω is at most 3. Strassen took the first step in a long series of gradually improving upper bounds on ω in 1969 when he showed that the product of two 2×2 matrices can be computed with seven field multiplications, rather than the eight required by the standard algorithm (Strassen, 1969). Since his algorithm did not make use of the commutativity of the matrix entries, it can be applied recursively to any square matrix with dimension a power of 2 simply by treating the entries of a $2^n \times 2^n$ matrix as $2^{n-1} \times 2^{n-1}$ blocks. Padding any matrix with zeros yields one whose dimensions are a power of 2. In this way, Strassen's work showed that ω is at most $\log_2 7 \approx 2.81$.

Gradual improvements have been made to upper bounds on ω since Strassen's original paper, most recently with the work of Coppersmith and Winograd, who showed that $\omega < 2.38$ (Coppersmith and Winograd, 1990). It is widely believed that $\omega = 2$, but since 1990 no improved upper bound has been discovered.

1.1.2 Group-Theoretic Matrix Multiplication

In 2003, Cohn and Umans presented a new approach to proving upper bounds on ω (Cohn and Umans, 2003). They showed that, by fixing a finite group G and choosing three subsets of G satisfying the so-called *triple product property*, an $n \times n$ matrix multiplication can be transformed into the convolution of two elements in the group algebra $\mathbb{F}G$ (in this thesis, I always take $\mathbb{F}$ to be $\mathbb{C}$). This convolution can then be computed efficiently by applying a discrete Fourier transform, analogous to the linear-time convolution of signals in traditional signal processing. Intuitively, one can think of the two matrices as being encoded into the time domain as signal vectors and then multiplied efficiently in the frequency domain. Traditional signal processing deals with only the case $G = \mathbb{Z}_n$, but in order to develop useful algorithms we will need to consider the more complicated situation in which G is nonabelian.

The computational cost of computing the convolution of the two embedded matrices depends upon the dimensions of the irreducible representations of G, what we will call the *character degrees* of G. If the matrix entries can be "packed" into the group sufficiently well, that is, if the embedding allows for large enough matrices to be multiplied such that the group algebra convolution requires fewer field multiplications, then the embedding realizes a nontrivial bound (i.e., less than 3) on ω. Roughly speaking, if we embed a $k \times k$ matrix multiplication in the group algebra $\mathbb{C}G$ and G has character degrees $\{d_i\}$, then Cohn and Umans showed

$$k^{\omega} \leq \sum_i d_i^{\omega}.$$

The left-hand side represents the amount of work needed to compute the product of two $k \times k$ matrices, which is no more than the amount of work needed to multiply several smaller matrices whose dimensions are the character degrees of G.

1.1.3 Index Sets through Group Actions

Finding subsets which satisfy the triple product property, the relation which must be satisfied by the embedding of matrices into a group algebra for the Cohn-Umans algorithm to produce the correct output, presents an interesting group-theoretic challenge. One way to search for such embeddings is to consider the action of the group G on some set X, and then to use geometric or combinatorial intuitions gained from this approach to choose appropriate subsets.

As a first example of this method, I will use geometric methods to construct an embedding of matrices into the complex group algebra of $\mathrm{GL}_2(\mathbb{F}_q)$, the general linear group over a finite field of order q. By appealing to geometric intuition in Euclidean space, I will show that a certain triple of subsets satisfies the triple product property.

Since wreath products can be naturally viewed as acting on an associated tree, I will give two examples of embeddings resulting from such actions and then generalize to actions on an arbitrary set. In particular, I will show that, for any groups G, H, and K,

$$(|G|^n \cdot |H| \cdot |K|)^{\omega/3} \leq \sum_i d_i^{\omega},$$

where $\{d_i\}$ represent the character degrees of $K \times (G^n \rtimes H)$ (here, H permutes the n copies of G). Unfortunately, this is not sufficient to give a nontrivial bound on ω, as will be see in Chapter 3, but it does provide a framework for approaching matrix multiplication with group actions.

The character degrees of $K \times (G^n \rtimes H)$ can be difficult to compute explicitly, so I will present a bound for estimating the sum of the d_i^{ω}. In particular,

$$\sum_i d_i^{\omega} \leq l \cdot k_{\max}^{\omega} \cdot |H|^{\omega-1} \cdot |G|^{n-1},$$

where l is the number of irreducible representations of K and $k_{\max}$ is the maximum character degree of K. I will prove this inequality in Chapter 3.

1.1.4 Outline of the Thesis

Section 1.2 of this thesis will present the mathematical background in representation theory, group-theoretic fast matrix multiplication, and theoretical computer science necessary for the reader in remainder of the document. Chapter 2 describes three examples of index sets realized through group actions. The first of these relies on geometric intuition in the action of a matrix group on a vector space over a finite field; the second uses combinatorial intuition concerning the action of a semi-direct product group on a depth-two tree; the third is a similar but more complicated action of an iterated wreath product group on a regularly branching tree. Chapter 3 generalizes the latter two cases to that of the action of a more general semi-direct product group on an arbitrary set, and describes possible extensions of the theory developed in this thesis.

1.2 Mathematical Background

The representation theory of finite groups is central to group-theoretic matrix multiplication. This section summarizes the theoretical computer science and representation theory necessary for this document, and presents the existing theory of group-theoretic matrix multiplication.

1.2.1 Formalization of ω and Tensor Rank

Let $\mathrm{M}(n)$ be the number of field operations required to compute the product of two $n \times n$ matrices over a field of characteristic 0. To be precise, we now define ω with

$$\omega = \inf\{r \in \mathbb{R} | \mathrm{M}(n) = O(n^r)\}. \tag{1.1}$$

If ω were defined directly in terms of runtime rather than in terms of field operations, the fact that the output matrix has n^2 entries would imply that $\omega \geq 2$; however, with the above definition, a more complicated argument leads to an identical lower bound (Bürgisser et al., 1997).

Because the standard matrix-multiplication algorithm computes the product in $2n^3$ field operations, we know that $\omega \leq 3$. Strassen's algorithm, which consisted of a noncommutative algorithm for multiplying 2×2 matrices with 7 multiplications rather 8, in effect showed that $\omega \leq \log_2 7$ by giving a recursive algorithm for computing the product of $2^k \times 2^k$ matrices (Strassen, 1969). Viewed differently, Strassen showed that the *rank* of the bilinear map

$$\mathbb{F}^{2\times 2} \times \mathbb{F}^{2\times 2} \to \mathbb{F}^{2\times 2}$$

defining 2×2 matrix multiplication is at most 7, that is, the minimum number of products of linear maps from the two copies of $\mathbb{F}^{2\times 2}$ which need to be summed in order to compute this bilinear map is 7 (the details of this are not important for this thesis, but the interested reader can consult (Bürgisser et al., 1997)). In general, computing the rank of $k \times k$ matrix multiplication to be at most l proves that

$$\omega \leq \log_k l. \tag{1.2}$$

Such bounds are difficult to come by: computing the rank of this collection of bilinear forms, that is, those encoding $k \times k$ matrix multiplication, is NP-hard (Håstad, 1990). We will see a way to develop such bounds without directly computing the rank of these forms.

1.2.2 Representation Theory

The Cohn-Umans algorithm, presented in the following section, relies upon a fundamental result from representation theory, namely that any complex group algebra is isomorphic to an algebra of complex block diagonal matrices. This is known as Wedderburn's theorem.

Theorem 1.1 (Wedderburn)**.** *The group algebra $\mathbb{C}G$ of a finite group G is isomorphic to a $\mathbb{C}$-algebra of block diagonal matrices:*

$$\mathbb{C}G \cong \bigoplus_{i=1}^{h} \mathbb{C}^{d_i \times d_i}, \tag{1.3}$$

where the d_i represent the character degrees of G, i.e., the dimensions of the irreducible representations of G, and h is the number of conjugacy classes of G. Additionally, we have that

$$|G| = d_1^2 + \cdots + d_h^2. \tag{1.4}$$

Proof. See Theorem 2.11 in Clausen and Baum (1993). □

Equation (1.4) will be of great use in estimating bounds on ω later in the thesis.

Definition 1.2. Any isomorphism of $\mathbb{C}$-algebras from $\mathbb{C}G$ to $\oplus_i \mathbb{C}^{d_i \times d_i}$ is called a *discrete Fourier transform*, or DFT, for $\mathbb{C}G$, or simply for G. The algebra of block diagonal matrices is sometimes referred to as the *frequency domain* of the group G, and $\mathbb{C}G$ is called the *time domain* of G. This analogy to signal processing will now be elucidated with an example.

Example 1.3. As an example, and in order to draw a parallel to traditional signal processing, consider $G = \mathbb{Z}_n$. Since G is abelian, the commutativity of the group algebra implies that the corresponding algebra of block diagonal matrices must also be commutative, and thus each irreducible representation of G has dimension 1 and we get diagonal matrices in the frequency domain. Taking $n = 3$, we have that

$$\mathbb{C}\mathbb{Z}_3 \cong \begin{pmatrix} z_1 & 0 & 0 \\ 0 & z_2 & 0 \\ 0 & 0 & z_3 \end{pmatrix}, \tag{1.5}$$

and the DFT has the form

$$DFT(\mathbb{Z}_3) = \begin{pmatrix} 1 & 1 & 1 \\ 1 & \alpha & \alpha^2 \\ 1 & \alpha^2 & \alpha \end{pmatrix}, \tag{1.6}$$

where α is a primitive 3^{rd} root of unity. In classical signal processing, the convolution of two elements of $\mathbb{C}G$ (vectors in the time domain) is turned into linear-time multiplication of diagonal matrices (which is the same as pointwise multiplication of vectors).

The abelian case cannot improve upon the naive algorithm because the commutative frequency domain requires as many field operations as the original matrix multiplication; using an abelian group is essentially the same as simply performing each multiplication called for in the naive algorithm. The following example illustrates how the frequency domain becomes more complicated in the noncommutative case. Ideally, the Cohn-Umans algorithm turns an $n \times n$ matrix multiplication into several smaller matrix multiplications in the frequency domain of some nonabelian group.

Example 1.4. Consider the complex group algebra $\mathbb{C}S_3$ of the symmetric group S_3 on three letters. This group has two irreducible representations of degree 1 and one of degree 2, so we get

$$\mathbb{C}S_3 \cong \begin{pmatrix} z_1 & 0 & 0 & 0 \\ 0 & z_2 & 0 & 0 \\ 0 & 0 & z_3 & z_4 \\ 0 & 0 & z_5 & z_6 \end{pmatrix}. \tag{1.7}$$

In order for this group to improve upon the naive matrix multiplication algorithm, i.e., use fewer than $O(n^3)$ field operations, we would need, roughly speaking, to embed a square matrix in $\mathbb{C}S_3$ such that the cube of the dimension of the square matrix exceeds the sum of the cubes of the character degrees. More precisely, an embedding of $n \times n$ matrices into $\mathbb{C}S_3$ demonstrates an upper bound on ω of less than 3 if and only if

$$n^3 > \sum_i d_i^3 = 10. \tag{1.8}$$

The left-hand side intuitively represents the amount of work we would have to do to get an answer using standard matrix multiplication on $n \times n$ matrices, and the right-hand side represents the amount of work needed for a multiplication in the frequency domain of $\mathbb{C}S_3$ (Cohn et al., 2005). An exhaustive search of the possible index sets in S_3, the smallest nonabelian group, showed that this particular group cannot achieve $\omega < 3$. All groups known to show that $\omega < 3$, presented in (Cohn et al., 2005), are much larger than S_3 (the smallest has order 250).

1.2.3 The Cohn-Umans Algorithm

This section follows closely one coauthored by Richard Bowen, Bo Chen, Martijn van Schaardenburg, and myself (Bowen et al., 2009). The definitions and theorems presented below were originally published in Cohn et al. (2005).

Definition 1.5. If S, T, U are ordered subsets of a group G, then the Cohn-Umans algorithm (Cohn and Umans, 2003) for matrix multiplication computes the product of matrices M and N of dimensions $|S| \times |T|$ and $|T| \times |U|$, respectively, as follows.

Index the rows of M by S^{-1}, the columns of M by T, the rows of N by T^{-1}, and the columns of N by U. Then let $f_M = \sum_{i,j} M_{i,j} s_i^{-1} t_j$ and $f_N = \sum_{j,k} N_{j,k} t_j^{-1} u_k$. Compute $f_P = f_M f_N$, and assign to $P_{i,k}$ the coefficient of $s_i^{-1} u_k$ in f_P.

Theorem 1.6. *The Cohn-Umans algorithm computes, in position i, k of the product matrix, the sum of all terms $M_{i',j} N_{j',k'}$, where*

$$s_{i'}^{-1} t_j t_{j'}^{-1} u_{k'} = s_i^{-1} u_k.$$

Proof. Every term in f_P is a product of a term in f_M with a term in f_N. The $s_i^{-1} u_k$ term is exactly the sum of all terms $(zm)(z'n)$, where $z, z' \in \mathbb{C}^{n \times n}$,

$m \in S^{-1}T$ and $n \in T^{-1}U$, and $mn = s_i^{-1}u_k$. But this is exactly the sum in the statement of the theorem. □

Corollary 1.7. *The Cohn-Umans algorithm is correct if and only if for all $s, s' \in S, t, t' \in T, u, u' \in U$, we have that $ss'^{-1}tt'^{-1}uu'^{-1} = e$ implies $s = s', t = t', u = u'$.*

Proof. This result follows from the previous theorem since any expression of the form

$$s_{i'}^{-1}t_j t_{j'}^{-1}u_{k'} = s_i^{-1}u_k$$

must satisfy the constraints $i = i', j = j', u = u'$, meaning that entry (i, k) of the product, indexed by s_i^{-1} and u_k, only contains terms formed by multiplying entry (i, j) by (j, k) in the left and right factor matrices, respectively. Furthermore, all such pairs will appear in that entry because multiplication in the group algebra obeys the distributive law. This is precisely the definition of matrix multiplication, so that the Cohn-Umans algorithm computes the correct product when the stated identity is satisfied. □

Definition 1.8. The property in Corollary 1.7 is called the *triple product property* (Cohn and Umans, 2003).

Cohn et al. use this method to convert particular triples of subsets of a group into an algorithm for matrix multiplication and hence into an upper bound on ω. To be precise, Cohn et al. (2005) shows that

Theorem 1.9. *If $S, T, U \subset G$ satisfy the triple product property and G has character degrees $\{d_i\}$, then*

$$(|S||T||U|)^{\omega/3} \leq \sum_i d_i^{\omega}. \tag{1.9}$$

Intuitively, the left-hand side represents the multiplication we are able to do with our index sets, and the right-hand side represents the amount of work that must be done in the frequency domain to carry out the Cohn-Umans algorithm.

Because the character degrees themselves are often difficult to compute, the identity

$$|G| = \sum d_i^2$$

gives us the following result.

Corollary 1.10. *If $S, T, U \subset G$ satisfy the triple product property and G has character degrees $\{d_i\}$, then*

$$(|S||T||U|)^{\omega/3} \leq |G| d_{max}^{\omega-2}, \tag{1.10}$$

where d_{max} is the highest character degree of G.

Example 1.11. The analysis of the bounds on ω obtained from (1.9) and (1.10) are particularly easy to analyze in the case of the finite Heisenberg group, defined as follows.

Definition 1.12. The finite Heisenberg group H over the finite field $\mathbb{F}_q$ of size q consists of the set

$$\left\{ \begin{pmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix} \middle| \, a, b, c \in \mathbb{F}_q \right\}$$

under the usual operation of matrix multiplication.

As shown in (Terras, 1999), the Heisenberg group has q^2 1-dimensional representations and $(q-1)$ representations of dimension q. Therefore the inequality in Theorem (1.9) becomes

$$(|S||T||U|)^{\omega/3} \leq q^{\omega+1} + q^2 - q^{\omega},$$

and Corollary (1.10) yields

$$(|S||T||U|)^{\omega/3} \leq q^{\omega+1}. \tag{1.11}$$

Taking the latter (weaker) inequality, we can take logs and rearrange to obtain

$$\frac{\omega}{\omega+1} \leq \frac{\log q}{\log\left((|S||T||U|)^{1/3}\right)}. \tag{1.12}$$

If we view the left hand side as a function of ω, we see that it asymptotes to 1 for large ω. For a given S, T, U triple satisfying the triple product property, the right hand side of Inequality (1.12) is a constant function, and the intersection point of $\frac{\omega}{\omega+1}$ with this constant is the resulting bound on ω. This realizes no bound at all if the right hand side is one or greater, which occurs when $(|S||T||U|)^{1/3} \leq q$; in other words, for this inequality to realize any bound at all, the geometric mean of the sizes of the index sets must be larger than the size of the finite field, which is also the dimension of the largest irreducible representation.

This makes intuitive sense, because we do not gain anything by embedding a multiplication problem smaller than $q \times q$ in the frequency domain of this group as there is an irreducible representation of dimension q. Note that $(|S||T||U|)^{1/3} = q$ is obtained by the trivial construction where each of S, T, and U is made up of subsets with zeros in two of the three free matrix entries. To match the naive algorithm, substituting $\omega = 3$ into Inequality (1.11) implies that a construction would have to satisfy $(|S||T||U|)^{1/4} = q$, and to prove that $\omega = 2$ we would need $(|S||T||U|)^{2/9} = q$. No construction proving $\omega < 3$ is known for this group.

Example 1.13. The following sets in $D_{12} = \langle x, y | x^6 = y^2 = 1, xy = yx^{-1} \rangle$ have the triple product property:

$$\begin{aligned} S &= \{1, y\} \\ T &= \{1, yx^2, x^3, xy\} \\ U &= \{1, yx\}. \end{aligned}$$

Thus, S, T, and U can be used to index the product of a 2 × 4 matrix by a 4 × 2 matrix. Explicitly, if we use these sets to index a matrix multiplication of the form

$$\begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,3} & a_{1,4} \\ a_{2,1} & a_{2,2} & a_{2,3} & a_{2,4} \end{bmatrix} \times \begin{bmatrix} b_{1,1} & b_{1,2} \\ b_{2,1} & b_{2,2} \\ b_{3,1} & b_{3,2} \\ b_{4,1} & b_{4,2} \end{bmatrix} = \begin{bmatrix} c_{1,1} & c_{1,2} \\ c_{2,1} & c_{2,2} \end{bmatrix}$$

we have a guarantee that convolving the group algebra elements

$$a_{1,1} \cdot 1 + a_{1,2} \cdot x^3 + a_{1,3} \cdot yx^2 + a_{1,4} \cdot xy + a_{2,1} \cdot y + a_{2,2} \cdot yx^3 + a_{2,3} \cdot x^2 + a_{2,4} \cdot x^{-1}$$

and

$$b_{1,1} \cdot 1 + b_{1,2} \cdot yx + b_{2,1} \cdot x^3 + b_{2,2} \cdot x^2 y + b_{3,1} \cdot yx^2 + b_{3,2} \cdot x^{-1} + b_{4,1} \cdot xy + b_{4,2} \cdot x^2$$

will yield the correct coefficients on the elements of $S^{-1}U$.

This particular triple of index sets does not, however, show that $\omega < 3$. The maximum character degree of D_{12} is 2, and the order of the group is 12, so the right-hand of Corollary 1.10 becomes $12 \cdot 2^{\omega-2}$. Thus we have the inequality

$$16^{\omega/3} \leq 12 \cdot 2^{\omega-2}.$$

Solving numerically for the best upper bound on ω implied by this gives us that $\omega < 4.76$, much weaker than the naive algorithm.

Chapter 2

Geometric and Tree Constructions

Viewing the triple product property in terms of a group action can transform the search for index sets into a variety of different problems. This chapter will show how to find index sets using geometric intuition through an action of a matrix group on a vector space, and combinatorial intuition with an action of a semi-direct product on a tree.

2.1 Index Sets from Group Actions

I will now describe the main ideas underlying this thesis project. In this section, I will formulate the problem of finding index sets in terms of the action of a group on a set. While none of the constructions I develop achieve a bound of $\omega < 3$, they do show how to apply new skill sets to the problem of creating fast matrix multiplication algorithms.

2.1.1 Group Actions and the Triple Product Property

Before studying the effectiveness of group actions in the construction of index sets, we will recharacterize the triple product property in a way better suited to this setting.

Theorem 2.1. *Subsets S, T, U of G satisfy the triple product property if and only if*

$$|S^{-1}||U| = |S^{-1}U|, \text{ and} \tag{2.1}$$

$$(S^{-1}(T \cdot T^{-1} \setminus e)U) \cap S^{-1}U = \emptyset, \tag{2.2}$$

where S^{-1} is the set of inverses of elements in S.

Proof.
($\Rightarrow$): If S, T and U satisfy the triple product property, then for no choices of elements $s, s' \in S$ and $u, u' \in U$ can it be the case that

$$s^{-1}u = (s')^{-1}u',$$

unless $s = s'$ and $u = u'$. Therefore $|S^{-1}||U| = |S^{-1}U|$. Suppose now that the intersection in Equation (2.2) is not empty; then there is an expression

$$s^{-1}t_i t_j^{-1}u = (s')^{-1}u',$$

where $s \neq s'$ or $u \neq u'$ (otherwise the expression would be trivial), but this is a violation of the triple product property. Therefore the intersection must be empty and the two properties are satisfied.
($\Leftarrow$): Suppose that there existed an expression

$$s^{-1}t_i t_j^{-1}u = (s')^{-1}u'.$$

Then there are two cases. First, if $i = j$, then the resulting expression clearly violates Equation (2.1), which means that if there were an expression violating the triple product property, then $i \neq j$. If, however, $i \neq j$ then the intersection in Equation (2.2) is nonempty. Thus such an expression cannot exist and S, T, and U satisfy the triple product property. □

Note that this formulation of the triple product property deals with the set S^{-1} instead of S. This does not make any difference to the algebraic properties of the index sets, and so we omit the inverse with this recharacterization and for the remainder of this document. The first index set will still be referred to simply as S, but we will check the triple product property in terms of the two properties in Equations (2.1) and (2.2) and hence not invert the elements of S when embedding matrices into the group algebra.

We now present a special case of the triple product property where subsets satisfying the triple product property arise through a group action. The idea is to formulate the properties necessary to correctly multiply matrices in the frequency domain without appealing directly to the structure of the group; in this sense, Definition 2.2 is unsatisfactory, as it does not rewrite $|SU| = |S||U|$ in terms of the group action. In the following example (whose development motivated Definition 2.2), however, this shortcoming of the definition was not difficult to overcome; this is discussed further in the example.

Definition 2.2. Let G be a finite group with a left action on the set X. We say that $S, T, U \subseteq G$ and $X_S, X_U, X_D \subseteq X$ (D stands for "destination") *realize* $(|S|, |T|, |U|)$ *as an action* if

$$U \cdot X_U \subseteq X_S, \tag{2.3}$$

$$S \cdot X_S = X_D, \tag{2.4}$$

$$((TT^{-1} \setminus e) \cdot X_S) \cap X_S = \emptyset, \tag{2.5}$$

$$|S||U| = |SU|. \tag{2.6}$$

Theorem 2.3. *If $S, T, U \subseteq G$ and $X_S, X_U, X_D \subseteq X$ realize $(|S|, |T|, |U|)$ as an action, then S, T, U satisfy the triple product property.*

Proof. There are two things to show in order to establish the triple product property: $|S||U| = |SU|$ and $S(TT^{-1} \setminus e)U \cap SU = \emptyset$. The former is assumed in the construction; we now show the latter.

Let $g \in S(TT^{-1} \setminus e)U \cap SU$. Then we know that, because $g \in SU$, $g \cdot X_U \subset X_D$. However, since $g \in S(TT^{-1} \setminus e)U$, g can be written in the form

$$g = st_1t_2^{-1}u. \tag{2.7}$$

Thus the action of g can also be viewed as

$$g : X_U \overset{t_1t_2^{-1}u}{\rightarrow} (TT^{-1} \setminus e) \cdot X_S \overset{s}{\rightarrow} X_D, \tag{2.8}$$

where we know that the last arrow holds by the above argument that $g \in SU$. By Equations (2.8) and (2.7), an appropriate choice of $s \in S$ yields a group element $s^{-1}g$ with the property that

$$s^{-1}g : X_U \rightarrow (TT^{-1} \setminus e) \cdot X_S;$$

however, by Equation (2.4) in the definition of the structure of the group action, we also have that

$$s^{-1}g : X_U \rightarrow X_S,$$

so that for any $x \in X_U$, $s^{-1}g$ satisfies

$$s^{-1}g \cdot x \in ((TT^{-1} \setminus e) \cdot X_S) \cap X_S = \emptyset,$$

meaning that $S(TT^{-1} \setminus e)U \cap SU = \emptyset$ by Equation (2.5) if X_U is nonempty. □

Example 2.4. I now present an application of this formulation of the triple product property that, while it does not show $\omega < 3$, does demonstrate a way in which this framework can apply geometric intuition to the problem of finding index sets. Let $G = GL_2(\mathbb{F}_q)$, and $X = \mathbb{F}_q^2$ (as column vectors under the obvious left action), with

$$X_U = X_D = \left\{ \left(\begin{array}{c} \lambda \\ 0 \end{array} \right) \middle| \lambda \in \mathbb{F}_q, \lambda \neq 0 \right\},$$
$$X_S = \left\{ \left(\begin{array}{c} 0 \\ \lambda \end{array} \right) \middle| \lambda \in \mathbb{F}_q, \lambda \neq 0 \right\}.$$

We then choose our subsets to be

$$S = \left\{ \left(\begin{array}{cc} \alpha & 1 \\ \beta & 0 \end{array} \right) \middle| \alpha, \beta \in \mathbb{F}_q, \beta \neq 0 \right\},$$
$$T = \left\{ \left(\begin{array}{cc} 1 & \lambda \\ 1 & 1 \end{array} \right) \middle| \lambda \in \mathbb{F}_q, \lambda \neq 0, 1 \right\},$$
$$U = \left\{ \left(\begin{array}{cc} 0 & 1 \\ \lambda & 0 \end{array} \right) \middle| \lambda \in \mathbb{F}_q, \lambda \neq 0 \right\}.$$

These choices are shown visually, with $\mathbb{F}_q^2$ depicted as two-dimensional Euclidean space to provide an intuitive picture of these sets, in Figure 2.1.

First, notice that an element of SU has the form

$$su = \left(\begin{array}{cc} \lambda & \alpha \\ 0 & \beta \end{array} \right),$$

which is clearly uniquely determined by s and u, thus showing Equation (2.6). Properties (2.3) and (2.4) are similarly clear: if $u \in U$ acts on a vector $x \in X_U$, then we have that

$$u \cdot x = \begin{pmatrix} 0 & 1 \\ \lambda_1 & 0 \end{pmatrix} \begin{pmatrix} \lambda_2 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ \lambda_1 \lambda_2 \end{pmatrix} \in X_S,$$

which shows the containment (2.3). If we now have $s \in S$ and $x \in X_S$, then we see that

$$s \cdot x = \begin{pmatrix} \alpha & 1 \\ \beta & 0 \end{pmatrix} \begin{pmatrix} 0 \\ \lambda \end{pmatrix} = \begin{pmatrix} \lambda \\ 0 \end{pmatrix} \in X_D.$$

This shows that $S \cdot X_S \subseteq X_D$. Elements of S^{-1} have the form

$$\left(\begin{array}{cc} 0 & \alpha \\ \beta & \gamma \end{array} \right),$$

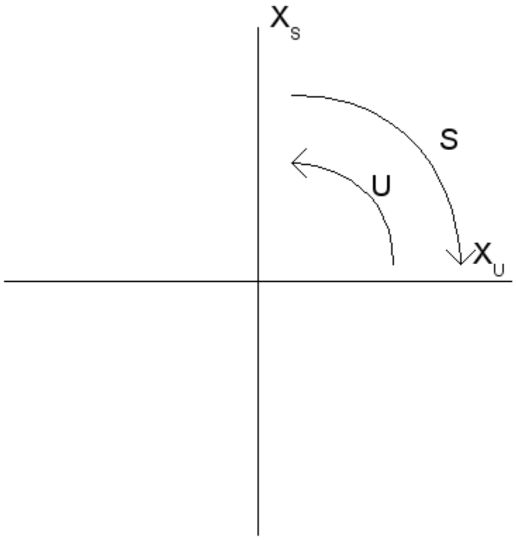

Figure 2.1: A visualization of the index sets described in Example 2.4.

and a similar argument shows that $S^{-1} \cdot X_D \subseteq X_S$ so that $S \cdot X_S = X_D$, as desired in Equation (2.4). The final property follows because each $t_i \in T$ takes

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

to a nontrivial linear combination

$$\begin{pmatrix} \lambda_i \\ 1 \end{pmatrix},$$

so that each element of T essentially "rotates" the coordinate system to a different extent and every element of $(TT^{-1} \setminus e) \cdot X_S$ will have a nonzero entry in the first coordinate and therefore the intersection with X_S is empty.

This construction demonstrates some of the advantages of realizing triple product sets through a group action: S and U send the vertical axis to the horizontal axis and vice versa, respectively. Each element of T, if we think of this vector space as having Euclidean geometry, rotates the coordinate system by a different angle so that $TT^{-1} \setminus e$ will shift X_S (the vertical axis). For this reason SU must act on the standard basis vector $(1,0)^{\mathrm{T}}$ in a different way than does $S(TT^{-1} \setminus e)U$. In this case, the group action and hence the index sets were informed by geometric intuition.

While the above characterization of the triple product property in terms of group actions is unnecessarily rigid for the examples and theory developed in the remainder of this thesis, it provides an example of geometric intuition, as opposed to the combinatorial intuition used hereafter in this thesis, applied to index set construction. The following section will demonstrate how to view finding index sets in terms of the action of a semi-direct product on the leaves of a tree. This hints at much broader possibilities for a group action-based approach, as diverse skill sets may be brought to bear.

2.2 Index Sets via Leaves and Subtrees

In this section I will present a construction of index sets in the group $G = \mathbb{Z}_{m^2} \times (\mathbb{Z}_{m^2} \wr S_n)$ based on the action of G on leaves and subtrees of a particular tree. This and the following construction, an iterated wreath product of cyclic groups, will lead us to a more general way of finding index sets in semi-direct product groups in the next chapter.

2.2.1 An Example from Cohn et al.

The construction presented in this section was an attempt to generalize the following construction from (Cohn et al., 2005). Although it did not succeed in improving upper bounds on ω, in fact it does not show $\omega < 3$, it provides a beginning for the application of group actions to index set construction. Let G be the semi-direct product

$$G = (\mathbb{Z}_n^3 \times \mathbb{Z}_n^3) \rtimes \mathbb{Z}_2, \tag{2.9}$$

where $\mathbb{Z}_2$ acts by permuting the two copies of $\mathbb{Z}_n^3$ in G. If we denote the three factors of $\mathbb{Z}_n^3$ by H_i, $i = 1, 2, 3$, and write elements of G in the form (a, b, z), where $a, b \in \mathbb{Z}_n^3$ and $z \in \mathbb{Z}_2$, then Cohn et al. define index sets S_1, S_2, and S_3 by

$$S_i = \{(a, b, z) | a \in H_i \setminus \{0\},\ b \in H_{i+1},\ z \in \mathbb{Z}_2\}.^{1} \tag{2.10}$$

In this definition, we take $H_4 = H_1$ for notational convenience. This is their first example of a group and index sets realizing a nontrivial bound on ω; the optimal value of the parameter n is at $n = 17$, yielding

$$\omega < 2.9088.$$

[1] Cohn et al. use the convention presented in Section 1.2.3 of indexing with $S_1^{-1}S_2$ and $S_2^{-1}S_3$.

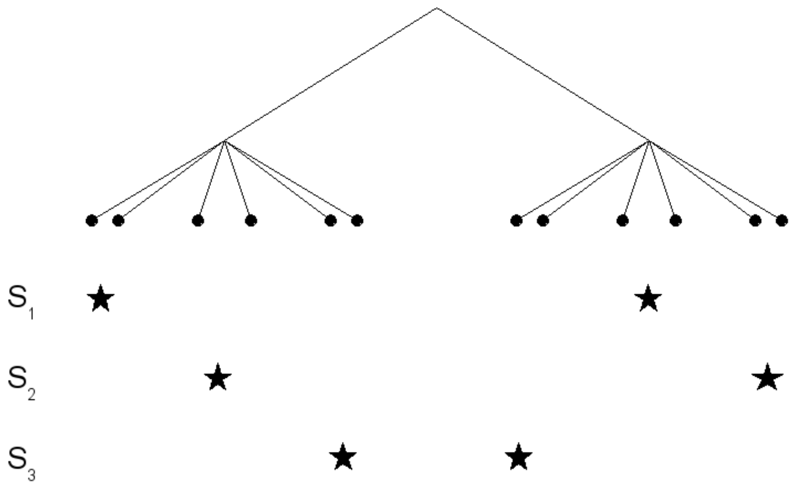

Figure 2.2: The tree naturally acted on by $(\mathbb{Z}_n^3 \times \mathbb{Z}_n^3) \rtimes \mathbb{Z}_2$ in the case of $n = 2$, together with the subtrees permuted by the index sets S_i, $i = 1, 2, 3$, indicated with stars.

Unlike in the cases of the index sets presented below, the proof that these sets satisfy the triple product property (which can be found in Cohn et al. (2005)) is purely algebraic. Semi-direct products of this form, however, naturally act on a tree of depth two where the subtrees correspond to the factor groups on the left (leaves corresponding to the factor group elements) and are permuted by the group on the right. From this perspective, the subsets defined above can be viewed in terms of their actions, as shown in Figure 2.2 for the case $n = 2$. The subtrees permuted by each S_i are indicated with stars.

The interaction of the last coordinates in the S_i (those which permute the two subtrees) make proving that these sets have the triple product property very complicated in the group action setting. The following subsection details and makes use of the natural action of semi-direct products on trees to easily show that a certain simpler triple of subsets satisfies the triple product property.

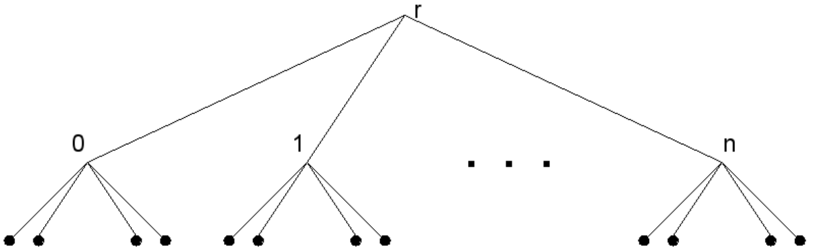

Figure 2.3: The tree acted on by $\mathbb{Z}_m^2 \times (\mathbb{Z}_m^2 \wr S_n)$, shown for the case of $m = 2$.

2.2.2 Partitioning the Group by Leaves and Subtrees

Consider the group $G^n \rtimes H$, where $H \leq S_n$ acts by permuting the n copies of G. This group has a natural action on the following tree (shown in a particular case in Figure 2.3): label the root r, its n children $\{1, \ldots, n\}$, and the $|G|$ children of vertex i with i_g for each $g \in G$. A group element of the form $(e, \ldots, e, h)$, where e is the identity of G, acts by

$$\begin{aligned} r &\mapsto r, \\ i &\mapsto h \cdot i, \\ i_g &\mapsto (h \cdot i)_g. \end{aligned}$$

The j-th coordinate g of an element of $G^n \rtimes H$ maps $j_{g'}$ to $j_{gg'}$ and fixes vertices outside of the j-th subtree (i.e., $i_{g'}$ for $i \neq j$ is fixed).

We now wish to make use of this natural action in order to find subsets with the triple product property. Recall from Theorem 2.1 that we would like to find three subsets S, T, U of a group G satisfying

$$|S||U| = |SU|, \tag{2.11}$$

$$(S(T \cdot T^{-1} \setminus e)U) \cap SU = \emptyset. \tag{2.12}$$

Let us consider the group $G = \mathbb{Z}_m^2 \times (\mathbb{Z}_m^2 \wr S_n)$; this group can also be viewed as $(\mathbb{Z}_m^2)^{n+1} \rtimes S_n$ in which the action of S_n fixes one copy of Z_m^2. Here, we view the leftmost factor group $\mathbb{Z}_m^2$ (those vertices labeled with a 0) as acting on an additional subtree of the tree described above, say with vertices labeled $\{0, 0_{g_1}, 0_{g_2}, \ldots\}$ for $g_i \in \mathbb{Z}_m^2$. Thus we may view G as acting on a tree consisting of vertices

$$\{r, 0, 1, \ldots, n, 0_{g_1}, \ldots, 0_{g_{2m}}, 1_{g_1}, \ldots, n_{g_{2m}}\},$$

shown in Figure 2.3, where r is the root, i is a child of r, and i_j is a child of vertex i labeled with $j \in [1, 2m]$. We take the leaves i_1 through i_m to be permuted by the left factor group of $\mathbb{Z}_m^2 = \mathbb{Z}_m \times \mathbb{Z}_m$, and i_{m+1} through i_{2m} to be permuted by the right factor group.

We can view this group as a subgroup of $\mathbb{Z}_m^2 \wr S_{n+1}$ (and, more fully, of $S_{2m} \wr S_{n+1}$), which presents an interesting opportunity in the search for effective index sets. Above, we chose a subgroup G of this total group of symmetries in which we can, hopefully, construct index sets satisfying the triple product property whose sizes compare favorably with the character degrees of G. In effect, there is a trade-off between minimizing the character degrees of the group chosen (i.e., trying to making G "small") and trying to build large index sets satisfying the triple product property. Understanding this trade-off would be a significant advance in group-theoretic matrix multiplication.

One way in which to choose subsets such that Equation (2.12) is satisfied would be to let T consist only of those elements with the identity in all but the first coordinate, that is,

$$T = \{(g, e, \ldots, e) | g \in \mathbb{Z}_m^2\}. \tag{2.13}$$

In terms of the group action, we can see this index set as the allowed automorphisms (that is, cyclic permutations) of the subtree labeled with a 0. Notice that T does not permute any of the vertices except the children of 0, meaning that if we choose S and U such that they fix this subtree, then any element of SU will stabilize 0 and its children, while any element of $S(T \cdot T^{-1} \setminus e)U$ will nontrivially permute the children of 0 (since the permutations of 0_{g_i} in T are all distinct). In other words, if each element of T is associated with a unique permutation of the children of 0 and each element of S and U is in the stabilizer of $\{0_j\}$, then Equation (2.12) is satisfied.

It then only remains to choose S and U from the stabilizer of $\{0_j\}$ in such a way that Equation (2.11) is satisfied. Notice that the i-th factor group $\mathbb{Z}_m^2 = \mathbb{Z}_m \times \mathbb{Z}_m$ can be conveniently partitioned into elements permuting i_1 through i_m and those permuting i_{m+1} through i_{2m}. If we let S permute the first m leaves of subtrees 1 through n and U permute the latter m leaves, and furthermore let the last coordinate of S (that is, the coordinate coming from S_n) be any element of S_n, then we have that

$$S = \{(e, g_1, \ldots, g_n, \pi) | g_i \in \mathbb{Z}_m \times \{0\}, \pi \in S_n\}, \ |S| = m^n \cdot n!,$$
$$U = \{(e, g_1, \ldots, g_n, e) | g_i \in \{0\} \times \mathbb{Z}_m, e\}, \ |U| = m^n.$$

The map $\varphi : S \times U \to SU$ defined by $\varphi(s, u) = su$ is clearly injective so that $|S||U| = |SU|$. Notice that SU consists of all elements of G not permuting the 0-th subtree, so that this choice of S and U is optimal given our previous choices (i.e., SU is the entire stabilizer of $\{0_j\}$). These sets satisfy Equations (2.11) and (2.12), and hence S, T, and U satisfy the triple product property.

2.2.3 Analysis of the Corresponding Bound on ω

To summarize, we chose $T \subset \mathbb{Z}_m^2 \times (\mathbb{Z}_m^2 \wr S_n) = G$ to permute one of the subtrees of an appropriate tree and, because S and U were chosen from the stabilizer of that subtree, could immediately see that Equation (2.12) was satisfied. Because we wanted S and U to be subsets of the stabilizer of a particular subtree, it was clear that SU must be a subset of the stabilizer; a way to partition the stabilizer into S and U such that $|S||U| = |SU|$ was then demonstrated. The result was

$$\begin{aligned} G &= \mathbb{Z}_m^2 \times (\mathbb{Z}_m^2 \wr S_n), \\ |S| &= m^n \cdot n!, \\ |T| &= m^2, \\ |U| &= m^n. \end{aligned}$$

In order for index sets to realize a nontrivial bound on ω, it is necessary for them to compute more information than the naive algorithm would produce if applied to the resulting frequency domain multiplication; symbolically, this means

$$|S||T||U| > \sum_i d_i^3. \tag{2.14}$$

Notice that the above construction gives $|S||T||U| = |G|$. A well-known identity from representation theory states that

$$\sum_i d_i^2 = |G|.$$

Therefore this particular construction cannot, for any values of m and n, prove that $\omega < 3$. It is, however, helpful in demonstrating a combinatorial approach to the construction of index sets with the triple product property.

Arriving at this result another way, we apply Theorem 1.9, which yields

$$(m^{2n+2} \cdot n!)^{\omega/3} \leq \sum_i d_i^\omega.$$

As the d_i are difficult to compute, we apply Corollary 1.10 to get

$$(m^{2n+2} \cdot n!)^{\omega/3} \leq m^{2n+2} \cdot n! \cdot d_{\max},$$

and we know that $d_{\max} \leq n!$, the index of $(\mathbb{Z}_m^2)^n$ in S_n; this upper bound on character degrees is stated in (Cohn et al., 2005). Rearranging this inequality to solve for ω yields

$$\omega \leq 3\left(\frac{2n\log m - \log n!}{2n\log m - 2\log n!}\right). \tag{2.15}$$

It is easy to see that this bound will, unfortunately, never show $\omega < 3$ for any choices of m and n since the denominator is always less than the numerator. (Note that the inequality is only valid for m and n for which the denominator is positive, since otherwise the direction of the inequality would be reversed and we would be left with a useless, negative lower bound on ω.)

2.3 An Iterated Wreath Product Group

This section will describe a construction similar to that of the last section, but for a very different group. I will define this group recursively, with $G_{1,p}$ being simply $\mathbb{Z}_p$ for p a prime, and then iteratively wreath the previous group with $\mathbb{Z}_p$. It will again act on an appropriate tree, and it will be partitioned based on a subtree and its stabilizer.

2.3.1 Defining $G_{n,p}$ and its Tree

Let $G_{1,p}$ be $\mathbb{Z}_p$ for $p > 2$ a prime (this constraint on the order of the underlying cyclic group is only necessary for the later discussion of the group's character degrees, but I put it here because I will confine my analysis to this case).

Definition 2.5. Given $G_{n-1,p}$, we define

$$G_{n,p} = G_{n-1,p} \wr \mathbb{Z}_p.$$

It is easy to see that

$$|G_{n,p}| = |G_{n-1,p}|^p \cdot p, \tag{2.16}$$

since we have p copies of the previous group, with a final coordinate from $\mathbb{Z}_p$ permuting them.

We define the action in a similarly recursive way: $G_{1,p}$ acts on the tree $T_{1,p}$ with vertices $\{r, 1, 2, \ldots, p\}$, where r is the root and has the p children

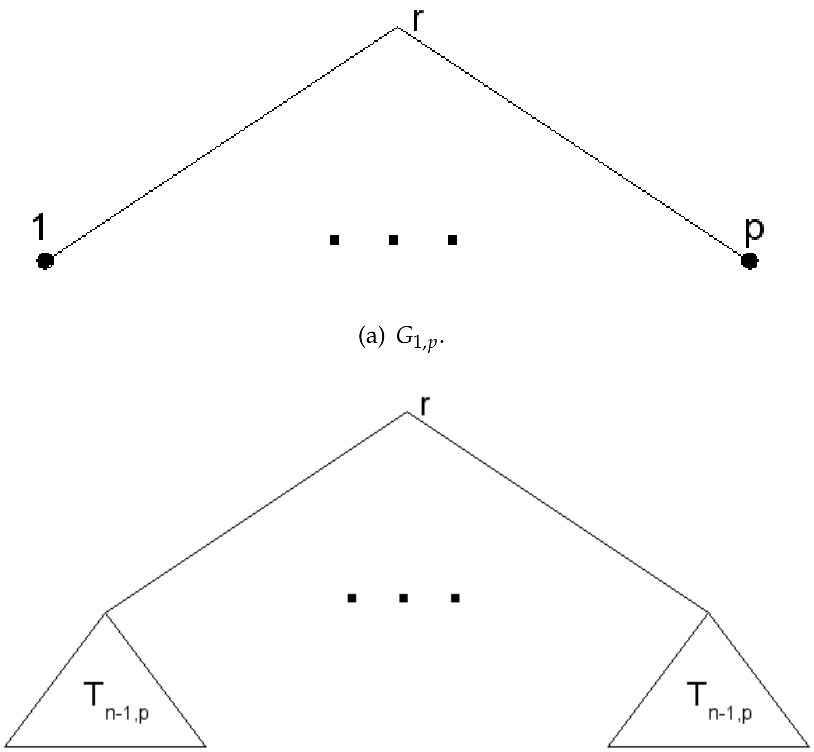

(a) $G_{1,p}$.

(b) $G_{n,p}$.

Figure 2.4: The trees acted on by $G_{1,p}$ and $G_{n,p}$.

labeled 1 through p, by cyclically permuting the vertices (the group element k cyclically permutes the vertices by the number k); this case of the action is shown in Figure 2.4(a). Now define $T_{n,p}$ to have p subtrees identical to T_{n-1}, with their roots r_i adjacent to a common root r, as shown in Figure 2.4(b). The i-th factor of $G_{n-1,p}$ permutes the subtree i-th subtree in the previously understood way, and the last coordinate of $G_{n,p}$ permutes the subtrees.

2.3.2 Index Sets in $G_{n,p}$

As in the previous section, let T be the subset of $G_{n,p}$ permuting the leftmost subtree and fixing the others:

$$T = \{(g, e, \ldots, e) | g \in G_{n-1,p}\}.$$

If we choose S and U from the stabilizer of this subtree then, as before, the condition in Equation (2.12) becomes trivial, since all elements of SU stabilize the leftmost subtree, while any element of $S(T \cdot T^{-1} \setminus e)U$ will nontrivially permute the leftmost subtree. This only leaves us with the problem of choosing S and U, subsets of the stabilizer of the leftmost subtree, such that $|S||U| = |SU|$. As an example of such a choice, divide the subtrees other than the first one evenly into sets A and B (this is possible for any odd prime), and let S be those permutations fixing all subtrees not in A, with U permuting those in A. Notice that, unlike in the previous construction, it is not possible to include the top-level permutations (the last coordinate of $G_{n,p}$) in SU, as any element of $\mathbb{Z}_p$ would move the leftmost subtree.

This construction gives us

$$\begin{aligned} |SU| &= |G_{n-1,p}| \cdot (p-1), \\ |T| &= |G_{n-1,p}|. \end{aligned}$$

Therefore we have that, for $S, T, U \subset G_{n,p}$,

$$|S||T||U| = |G_{n-1,p}|^2 \cdot (p-1).$$

As before, this construction cannot show that $\omega < 3$, since the sizes of the sets grow more slowly than that of the group, and hence does not exceed $\sum_i d_i^2$, much less $\sum_i d_i^3$. Nonetheless, an analysis of the character degrees of this group is worthwhile, in the event that a way to expand or modify these index sets is found.

2.3.3 Estimates and Computation of Character Degrees

A method for determining the character degrees of $G_{n,r}$ (for general r) is given in (Orellana et al., 2004); I will describe a simple algorithm for determining the character degrees in the special case where r is a prime p. Pseudocode for this algorithm is shown in Appendix A.

Proceeding inductively, it is clear that $G_{1,p}$ has p distinct 1-dimensional irreducible representations. Given $\{\rho_i\}$, a complete set of irreducible representations of $G_{n-1,p}$, we begin by defining the inertia group of a p-tuple of the ρ_i tensored together, that is, an irreducible representation of $G_{n-1,p}^p$. The *inertia group* of any such tuple $\rho_1 \otimes \cdots \otimes \rho_p$ is $\{e\}$ if any pair of irreducible representations in the tensor product differs, and $\mathbb{Z}_p$ otherwise. More generally, the inertia group of such a tuple t for $G \wr H$ is the subgroup of H under whose action t is invariant; when $H = \mathbb{Z}_p$, this is always either $\{e\}$ or $\mathbb{Z}_p$.

To produce the irreducible representations of $G_{n,p} = (G_{n-1,p})^p \rtimes \mathbb{Z}_p$, first notice that the irreducible representations of $(G_{n-1,p})^p$ are of the form $\rho = \rho_1 \otimes \cdots \otimes \rho_p$, where the ρ_i are, as above, irreducible representations of $G_{n-1,p}$. Let H be the inertia group of ρ, and $\{\sigma_i\}$ its irreducible representations. The irreducible representations of $G_{n,p}$ are then of the form $\rho \otimes \sigma_i$, induced from $(G_{n-1,p})^p \rtimes H$ to the full group $(G_{n-1,p})^p \rtimes \mathbb{Z}_p$. The details of this process are not important for this thesis, as we are only interested in the dimensions of these irreducible representations. To determine these dimensions, we have to consider two cases: if $H = \{e\}$, then the induction increases the dimension of $\rho \otimes \sigma_i$ (notice that σ_i is always one-dimensional) by a factor of p, and in this case one gets p equivalent copies of each irreducible representation; if $H = \mathbb{Z}_p$, then the dimension of the resulting irreducible representation is the same as that of $\rho \otimes \sigma_i$ and each such representation is inequivalent.

Chapter 3

General Theory and Possible Extensions

The method used in Section 2.2 to construct index sets can be generalized to a class of semi-direct products and associated group actions on essentially arbitrary sets. This suggests a multitude of possible extensions of group actions applied to fast matrix multiplication, including acting on structured sets.

3.1 The Stabilizer Construction: A Generalization of Tree Constructions

This section introduces the general theory of the *stabilizer construction*, a method for constructing index sets with the triple product property based on an abstraction of the constructions in Chapter 2.

3.1.1 Action on a General Set

Suppose we have a set X_0 and n copies of another set X, denoted with $X_1, \ldots, X_n$. Then we let $K \leq \mathrm{Aut}(X_0)$. Additionally, let G be some group with an action on $X = X_i, i = 1, \ldots, n$: $G \leq \mathrm{Aut}(X)$. Finally, let H be some subgroup of S_n. The constructions of the previous section will now be generalized to the action of the group

$$K \times (G^n \rtimes H), \tag{3.1}$$

where H permutes the n copies of G, on the set

$$X_0 \times X_1 \times \cdots \times X_n.$$

This action is defined in the natural way, with the first coordinate of the group (corresponding to K) permuting the coordinate corresponding to X_0, the i-th copy of G permuting X_i, and H permuting the indices of the X_i, $i = 1, \dots, n$. With this action, I will demonstrate index sets which satisfy the triple product property and achieve $|S||T||U| = |K||G|^n|H|$, the size of the group in question.

Theorem 3.1. *If $\{d_i\}$ are the character degrees of the group defined in Equation (3.1), then*

$$(|K||G|^n|H|)^{\omega/3} \leq \sum_i d_i^\omega.$$

Proof. Let T be those elements of $K \times (G^n \rtimes H)$ with identity elements in all coordinates except possibly the first. Let S be those group elements with an identity in all coordinates except possibly the last (the coordinate corresponding to H), and let U be those group elements with identity elements in the first and last coordinates, but with the other coordinates ranging freely over G. Then S and U stabilize X_0, but each element of T permutes X_0 differently so that every element of $S(T \cdot T^{-1} \setminus e)U$ permutes X_0 nontrivially. Therefore these index sets satisfy Equation (2.12). Given any element of SU, we can determine its preimage under the product map $\varphi : (s, u) \mapsto su$ by reading s off of the last coordinate, and u off of the entries other than the first and the last. Thus $|S||U| = |SU|$, and S, T, and U satisfy the triple product property. The result then follows from Theorem 1.9. □

Unfortunately, this construction cannot immediately show that $\omega < 3$ for any choice of K, G, and H. We can see this because

$$|S||T||U| = |K||G|^n|H| = \sum_i d_i^2 \leq \sum_i d_i^3.$$

However, the above argument presents a combinatorial way to analyze the class of groups defined in Equation (3.1). The obvious inequality

$$\sum_i d_i^3 \leq d_{\max} \sum_i d_i^2 \tag{3.2}$$

shows that we need to gain no more than a factor of the maximum character degree of the group in order to prove a nontrivial bound on ω. It seems possible that an extension of this technique, possibly combined with the results in group-theoretic partial matrix multiplication developed by the Harvey Mudd College Applied Representation Theory Group (Bowen et al., 2009), could improve the index sets sufficiently to prove substantial bounds.

3.1.2 Estimating $\sum d_i^\omega$

In order for index sets in $K \times (G^n \rtimes H)$ to be easily analyzed, I developed a simple upper bound on $\sum_i d_i^\omega$, the right-hand side in Theorem 1.9, which requires much less understanding of the character degrees of this fairly complicated group. Although I believe it to be true in greater generality, I have only found a proof for the case where G is abelian.

We will need the following lemma, which is Lemma 1.2 from Cohn et al. (2005):

Lemma 3.1. *Let $\{d_i\}$ be the character degrees of a finite group G and $\{c_j\}$ the character degrees of $G^n \rtimes S_n$, where S_n acts by permuting the n copies of G. Then*

$$\sum_j c_j^\omega \leq (n!)^{\omega-1} \left(\sum_i d_i^\omega \right)^n . \tag{3.3}$$

If G is abelian, then the proof given in Cohn et al. (2005) shows that we may replace S_n with any group H and obtain a special case of Inequality (3.3):

$$\sum_j c_j^\omega \leq |H|^{\omega-1} |G|^n. \tag{3.4}$$

Theorem 3.2. *If $\{d_i\}$ are the character degrees of $K \times (G^n \rtimes H)$, G abelian, then*

$$\sum_i d_i^\omega \leq l k_{max}^\omega |H|^{\omega-1} |G|^n, \tag{3.5}$$

where l is the number of irreducible representations of K, of which the highest dimension is k_{max}.

Proof. Let $\{c_i\}$ be the character degrees of $G^n \rtimes H$. Since the irreducible representations of the direct product of K and $G^n \rtimes H$ are the tensor products of all pairs of irreducible representations of K and $G^n \rtimes H$, we can write

$$\sum_i d_i^\omega = \sum_{i=1}^h \sum_{j=1}^l (c_i k_j)^\omega,$$

where $\{k_j\}$ are the character degrees of K. Since each term is made no smaller by replacing k_j with k_{max}, we obtain

$$\begin{aligned} \sum_{i=1}^h \sum_{j=1}^l (c_i k_j)^\omega &\leq \sum_{i=1}^h \sum_{j=1}^l k_{\text{max}}^\omega c_i^\omega \\ &= l k_{\text{max}}^\omega \sum_i c_i^\omega . \end{aligned}$$

Applying Inequality (3.4) to the above result gives

$$\sum_i d_i^\omega \leq l k_{\max}^\omega |H|^{\omega-1}|G|^n,$$

as desired. □

Example 3.3. Applying this inequality to the construction given in Section 2.2 for the group $\mathbb{Z}_m^2 \times (Z_m^2 \wr S_n)$ yields

$$\omega \leq 3\left(\frac{2n\log(m) - \log(n!)}{2n\log(m) - 2\log(n!)}\right), \tag{3.6}$$

which is clearly always greater than or equal to 3, consistent with the observation at the end of Section 3.1.1 that such index sets do not prove a nontrivial bound on ω. This is also precisely the same bound as was found for this construction by a different method in Section 2.2.3 (since $K = \mathbb{Z}_m^2$ was also abelian).

3.1.3 Conclusion

We have seen how two different sorts of mathematical intuitions, namely geometric and combinatorial, can be brought to bear on the problem of finding index sets via group actions. A more general way to view the latter was presented, with a construction that comes within a simple factor of realizing a nontrivial bound on ω. The following section will discuss some possible next steps to more fully utilize this new approach of index sets via group actions.

3.2 Future Work

Group-theoretic matrix multiplication is a relatively recent development in the study of ω, and as a result there is a plethora of possible extensions. I have catalogued some possible further work, mostly related to group actions, below for interested researchers.

- It may provide additional insights to consider actions of a group G on a set X with some structure, for example a partially ordered set. If X is partially ordered, then we could perhaps describe the index sets in terms of this ordering. For example, elements of S could have the property that their action makes a certain pair of incomparable elements $x_0, x_1 \in X$ satisfy $s \cdot x_0 < s \cdot x_1$.

The added structure of X might allow for a less coarse choice of index sets than was presented in this thesis (that is, more of the automorphisms of the subsets X_i presented in Chapter 3 could be included in S, T, and U). In this vein, it may also be interesting to construct index sets for G in terms of the action of G on X, together with the way this action interacts with some function

$$f : X \to \mathbb{R}.$$

For example, one might take S and U to be subsets of G which decrease the value of f on a certain point $x \in X$, with T disrupting this property in some way such that SU and $S(T \cdot T^{-1} \setminus e)U$ have empty intersection.

- Figure 2.2 shows how a particular triple of index sets developed in Cohn et al. (2005) which prove $\omega < 3$ can be understood in terms of the action of a group G on a tree. Notice that each set S_i permutes one third of the subtrees, whereas the construction presented in Chapter 3 assigns one subtree to T and divides the others evenly between S and U.

  It would be interesting to see whether there are other ways of generalizing the construction of Cohn et al., such that the sizes of the three index sets remain more "balanced", that is, they permute roughly equal numbers of subtrees. Qualitatively, it seems reasonable that keeping the three index sets roughly the same size results in $|S||T||U|$ being larger than when S and U are disproportionately large relative to T. For this reason, I am hopeful that a more balanced generalization of the construction of Cohn et al. might yield useful upper bounds on ω.

- Group-theoretic matrix multiplication computes products of matrix entries other than those desired for the output. The triple product property guarantees that these terms do not appear as coefficients of those group elements which index the output (namely SU). These additional terms can be understood as encoding the products of certain other rows and columns of matrices not included in the original problem. For example, if we compute the product

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix}$$

group-theoretically, then we could make use of the "unwanted" terms to multiply the larger, structured matrices,

$$\begin{bmatrix} a & b \\ c & d \\ b & a \end{bmatrix} \begin{bmatrix} \alpha & \beta \\ \gamma & \delta \end{bmatrix},$$

at no extra computational cost. This is because the nature of the group algebra multiplication leaves us with the dot product of the row vector $\begin{bmatrix} b & a \end{bmatrix}$ with each column of the other matrix. Studying the sorts of structured matrices that arise in this way (that is, which patterned matrices can be multiplied at little or no additional computational cost) could lead to further interesting interactions of group theory with matrix multiplication algorithms.

- The Cohn-Umans algorithm for group-theoretic matrix multiplication does not require that the elements used to index matrices be simply group elements: the matrix entries could be indexed with any elements of the group algebra $\mathbb{C}G$ satisfying the triple product property. Using linear combinations of group elements allows for increased flexibility in our choice of index sets and thus could yield useful upper bounds on ω.

- Cohn et al. showed that it is possible to construct an algorithm faster than $O(n^3)$ for matrix multiplication by transforming the problem of multiplying two $n \times n$ matrices into that of multiplying two block diagonal matrices. If, additionally, we can guarantee that the block diagonal matrices are sparse, then it may be possible to realize interesting upper bounds on ω by applying fast algorithms for sparse matrix multiplication. Doing so requires an appropriate choice of index sets and this is, to my knowledge, a completely unexplored problem.

Appendix A

An Algorithm for Computing the Character Degrees of the Iterated Wreath Product of $\mathbb{Z}_p$

The following page contains pseudocode for the algorithm given in Section 2.3.3. In the pseudocode, I abbreviate irreducible representation with "irrep".

```
Input: characterDegreeList $(N, p)$
Result: List of pairs: $[irrepindex, dimension]$
if $N$ == $1$ then
    L:= empty list
    for $x \leftarrow 1$ to $p + 1$ do
        Append to L an irrep with index $x$ and dimension 1
        return L
else
    previousCharacterDegrees:= characterDegreeList $(N - 1, p)$
    L:= all $N$-tuples of length $p$ with entries in
    previousCharacterDegrees
    WPDegrees:= empty list
    irrepIndex:= 1
    WPDegreeOrbits:= empty dictionary
    foreach $l$ in L do
        if all first coordinates of tuples in $l$ are equal then
            for $x \leftarrow 1$ to $p + 1$ do
                Add to WPDegrees an irrep of index irrepIndex and
                dimension $(l[0][1])^p$
                irrepIndex = irrepIndex + 1
        else
            tensorProductDegree:= $p$ times the product of the degrees
            of irreps in $l$
            if tensorProductDegree is in WPDegreeOrbits then
                Increment the value of WPDegreeOrbits at key
                tensorProductDegree
            else
                Create new key tensorProductDegree in
                WPDegreeOrbits and set it to 1
    foreach irrepDegree in the list of keys of WPDegreeOrbits do
        for $x \leftarrow 0$ to $\frac{(\text{WPDegreeOrbitsatkeyirrepDegree})}{p}$ do
            Add to WPDegrees an irrep of index irrepIndex and
            dimension irrepDegree
            irrepIndex = irrepIndex + 1
    return WPDegrees
```

Bibliography

Bowen, R.S., B. Chen, H. Orem, and M. van Schaardenburg. 2009. Group-Theoretic Partial Matrix Multiplication. URL `http://www.citebase.org/abstract?id=oai:arXiv.org:0902.2407`. 1.2.3, 3.1.1

Bürgisser, P., M. Clausen, and M.A. Shokrollahi. 1997. *Algebraic Complexity Theory*. Springer. 1.1.1, 1.2.1

Clausen, M., and U. Baum. 1993. *Fast Fourier Transforms*. Wissenschaftsverlag. 1.2.2

Cohn, H., R. Kleinberg, B. Szegedy, and C. Umans. 2005. Group-Theoretic Algorithms for Matrix Multiplication. *Foundations of Computer Science 46th Annual IEEE Symposium on* 23–25. 1.4, 1.2.3, 1.2.3, 2.2.1, 2.2.1, 2.2.3, 3.1.2, 3.1.2, 3.2

Cohn, H., and C. Umans. 2003. A Group-Theoretic Approach to Fast Matrix Multiplication. *Foundations of Computer Science, 2003 Proceedings 44th Annual IEEE Symposium on* 438–449. 1.1.2, 1.5, 1.8

Coppersmith, D., and S. Winograd. 1990. Matrix Multiplication via Arithmetic Progressions. *Journal of Symbolic Computation* 9(3):251–280. 1.1.1

Håstad, J. 1990. Tensor Rank is NP-Complete. *J Algorithms* 11(4):644–654. 1.2.1

Orellana, R.C., M.E. Orrison, and D.N. Rockmore. 2004. Rooted Trees and Iterated Wreath Products of Cyclic Groups. *Advances in Applied Mathematics* 33(3):531–547. 2.3.3

Strassen, V. 1969. Gaussian Elimination is not Optimal. *Numerische Mathematik* 13(4):354–356. 1.1.1, 1.2.1

Terras, A. 1999. *Fourier Analysis on Finite Groups and Applications*. Cambridge University Press. 1.11