Separation of Variables and the Computation of Fourier Transforms on Finite Groups, II

David Maslen[1], Daniel N. Rockmore[2†], and Sarah Wolff[2‡]

[1]*HBK Capital Management, New York, NY 10036,*
[2]*Departments of Mathematics and Computer Science, Dartmouth College, Hanover, NH 03755*
[3]*Department of Mathematics and Computer Science, Denison University, Granville, OH 43023*

Abstract. We present a general diagrammatic approach to the construction of efficient algorithms for computing the Fourier transform of a function on a finite group. By extending work which connects Bratteli diagrams to the construction of Fast Fourier Transform algorithms we make explicit use of the path algebra connection and work in the setting of quivers. In this setting the complexity of an algorithm for computing a Fourier transform reduces to path counting in the Bratelli diagram, and we generalize Stanley's work on differential posets to provide such counts. Our methods give improved upper bounds for computing the Fourier transform for the general linear groups over finite fields, the classical Weyl groups, and homogeneous spaces of finite groups.

Résumé. Nous présentons une approche schématique générale à la construction d'algorithmes efficaces pour le calcul de la transformée de Fourier d'une fonction sur un groupe fini. En étendant le travail qui relie diagrammes Bratelli à la construction d'algorithmes efficaces nous faisons usage explicite de la connexion de l'algèbre des chemins et le travail dans le cadre de carquois. Dans ce cadre la complexité d'un algorithme de calcul de transformée de Fourier réduit à des comptages de certains sous-carquois du diagramme Bratelli. Nos méthodes donnent améliorées limites supérieures pour le calcul de la transformée de Fourier pour les groupes linéaires généraux sur les corps finis, les groupes de Weyl classiques, et des espaces homogènes de groupes finis.

Keywords. Fast Fourier Transform, Bratteli diagram, path algebra, quiver

1   Introduction

The *Fast Fourier Transform* (FFT) remains among the most important family of algorithms in information processing [Roc00]. It efficiently computes the *discrete Fourier transform* (DFT) which is equivalent to the matrix-vector multiplication

$$\left(e^{2\pi ijk/n}\right)_{j,k}\vec{f} \tag{1}$$

for $i = \sqrt{-1}$, $j, k = 0, \ldots n-1$, and $\vec{f}$ a complex-valued vector of length n [Roc00]. This calculation can be framed in a number of ways. We take a representation theoretic point of view and cast the DFT as a

---

change of basis in $\mathbb{C}[C_N]$, the complex group algebra of the cyclic group of order N, from a natural basis of group element indicator functions to a basis of irreducible matrix elements. This perspective suggests a generalization of the DFT to finite nonabelian groups G as the computation of a change of basis in $\mathbb{C}[G]$ from a basis of indicator functions to a basis of irreducible matrix elements and raises attendant questions of computational complexity addressed herein.

Let $T_G(R)$ denote the computational complexity of the Fourier transform on a group G at a set of inequivalent irreducible representations R. Then $C(G)$ denotes the *complexity of the group G*, defined as

$$C(G) := \min_R \{T_G(R)\}.$$

For N a "highly composite" number, Cooley and Tukey in 1965 famously presented an algorithm to show $C(\mathbb{Z}/N\mathbb{Z}) \leq O(N \log_2 N)$ [CT65]. Yavne [Yav68] later showed that for $N = 2^m$, $C(\mathbb{Z}/N\mathbb{Z}) \leq \frac{8}{3}N \log_2 N - \frac{16}{9}N - \frac{2}{9}(-1)^{\log_2(N)} + 2$. More recently, Johnson and Frigo [JF07] and Lundy and Van Buskirk [LVB07] further reduced the total number of complex multiplications required, but without affecting the overall group complexity $C(\mathbb{Z}/N\mathbb{Z})$. More generally, for A an abelian group of size N, various efficiencies can be combined to prove the complexity of the DFT on A is bounded above by $O(N \log_2 N)$ [Dia80]. The deep and ongoing study of this problem has been motivated by a wide range of applications in digital signal processing and beyond (see e.g. [AT79, DMT13, TAL97]).

The Cooley-Tukey algorithm is undoubtedly the most famous of the FFTs. It is a divide-and-conquer algorithm whose key step is to rewrite the DFT on a cyclic group C_N as a linear combination of DFTs on $C_n < C_N$ (for $n \mid N$). Iterating this step for a chain of subgroups of C_N yields algorithms more efficient than a direct matrix-vector multiplication.

In this paper we continue a line of work that generalizes this approach to nonabelian groups [MR95, MR97b, MR00]. The computations are encoded via paths in a Bratteli diagram associated to the group of interest, which in turn means that matrix elements correspond to pairs of paths in the diagram, which for a given group element may only be nonzero when of a particular form. The "repeated units" of our divide-and-conquer approach are subgraphs of a Bratteli diagram and efficiencies are gained by recognizing their multiple appearances in the corresponding calculation. This is the guts of the "separation of variables" (SOV) approach first introduced in [MR97b] and extended in [Mas98].

We take on the problem of laying a proper axiomatic and logical foundation for this approach and in so doing also produce improved algorithms for the important families of classical Weyl groups B_n and D_n and the general linear groups over finite fields $GL_n(\mathbb{F}_q)$:

Theorem 1.1 $C(B_n) \leq n(2n-1)|B_n|$.

Theorem 1.2 $C(D_n) \leq \frac{n(13n-11)}{2}|D_n|$.

Theorem 1.3 $C(Gl_n(\mathbb{F}_q)) \leq \left(\frac{4^n q^{n+1} - q}{4q-1} + \frac{n(n+1)(2n+1)(q-1)}{6} \right) |Gl_n(q)|$.

Improvements for the complexity of Fourier transforms on related homogeneous spaces are also gained. For example, let B_n/B_{n-k} denote the homogenous space of the Weyl group B_n.

Theorem 1.4 $C(B_n/B_{n-k}) \leq k(4n-2k-1)\frac{|B_n|}{|B_{n-k}|}$.

The following is an extended abstract of [MRW]. Although we have made an effort to keep it self-contained, full details can be found in the full paper. In Sections 2 and 3 we outline the preliminaries

needed for our results, including a discussion of the main ideas behind the SOV approach. Our main theorem of Section 3, Theorem 3.5, reduces complexity counts to combinatorial path-counting. We conclude in Section 3.1 by generalizing Stanley's work on differential posets [Sta88, Sta90] to provide the main combinatorial counts needed for Theorems 1.1, 1.2, 1.3, and 1.4.

2 Background

2.1 Fourier transforms and the group algebra

Definition 2.1 *Let G be a finite group and f a complex-valued function on G.*

(i) *Let ρ be a matrix representation of G. Then the **Fourier transform of f at** ρ, denoted $\hat{f}(\rho)$, is the matrix sum*

$$\hat{f}(\rho) = \sum_{s \in G} f(s)\rho(s).$$

(ii) *Let R be a set of matrix representations of G. Then the **Fourier transform of f on R** is the direct sum of Fourier transforms of f at the representations in R:*

$$\mathcal{F}_R(f) = \bigoplus_{\rho \in R} \hat{f}(\rho) \in \bigoplus_{\rho \in R} \mathrm{Mat}_{\dim \rho}(\mathbb{C}).$$

When we compute the Fourier transform for a complete set of inequivalent irreducible representations R of G we refer to the calculation as the **computation of a Fourier transform on** G (with respect to R).

Definition 2.2 *Let G be a finite group, R a set of matrix representations of G.*

(i) *A **straight-line program** is a list of instructions for performing the operations $\times, \div, +, -$ on inputs and precomputed values.*

(ii) *The **arithmetic complexity** of a Fourier transform on R, denoted $T_G(R)$, is the minimum number of complex arithmetic operations needed to compute the Fourier transform of f on R via a straight-line program for an arbitrary complex-valued function f defined on G. As in [Mas98] we will always define the arithmetic complexity to be the maximum of the number of complex multiplications and the number of complex additions.*

(iii) *The **complexity of the group** G, denoted $C(G)$ is defined by*

$$C(G) := \min_R \{T_G(R)\},$$

where R varies over all complete sets of inequivalent irreducible representations of G.

(iv) *The **reduced complexity**, denoted $t_G(R)$, is defined by*

$$t_G(R) = \frac{1}{|G|}T_G(R).$$

Let $\rho_1, \ldots, \rho_m$ be a complete set of inequivalent irreducible matrix representations of a group G of dimensions $d_1, \ldots, d_m$, respectively. A direct computation of a Fourier transform would require at most $|G| \sum d_i^2 = |G|^2$ arithmetic operations. Rewriting, for a direct computation we have

$$C(G) \leq T_G(R) \leq |G|^2.$$

Fast Fourier transforms (FFTs) are algorithms for computing Fourier transforms that improve on this naive upper bound.

The group algebra $\mathbb{C}[G]$ is the space of all formal complex linear combinations of group elements under the product

$$\left(\sum_{s \in G} f(s)s \right) \left(\sum_{t \in G} h(t)t \right) = \sum_{s,t \in G} f(s)h(t)st.$$

Elements of $\mathbb{C}[G]$ are in one-to-one correspondence with complex-valued functions on G, and the group algebra product corresponds to convolution of functions.

A complete set R of inequivalent irreducible matrix representations of a group G determines a basis for $\mathbb{C}[G]$ and in this case the Fourier transform of a function f on G with respect to a complete set of inequivalent irreducible representations R of G is an algebra isomorphism

$$\mathbb{C}[G] \xrightarrow{\ \mathcal{F}_R\ } \bigoplus_{\rho \in R} M_{\dim(\rho)}(\mathbb{C}),$$

and so as in [Cla89, Mas98]:

Lemma 2.3 *The computation of the Fourier transform of a function f on G with respect to a complete set of irreducible representations R is equivalent to computation of*

$$\sum_{s \in G} f(s)s$$

in the group algebra, relative to a fixed basis for R.

2.2 Bratteli diagrams and quivers

The fundamental idea of the SOV approach is a recasting of the Cooley-Tukey algorithm in terms of graded quivers, which is an elaboration of path algebras derived from Bratteli diagrams.

Definition 2.4 *A **quiver** Q is a directed multigraph with vertex set $V(Q)$ and edge set $E(Q)$. For an arrow $e \in E(Q)$ from vertex β to vertex α, we call α the **target** of e and β the **source** of e.*

Let Q be a quiver. For each $e \in E(Q)$, let $t(e)$ denote the target of e and $s(e)$ the source of e.

Definition 2.5 *A quiver Q is **graded** if there is a function $gr : V(Q) \to \mathbb{N}$ such that $gr(t(e)) > gr(s(e))$, for each $e \in E(Q)$.*

Example 2.6 *Figure 1 is an example of a graded quiver. Each vertex v is labeled by its grading, $gr(v)$.*
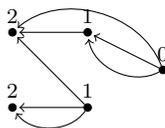
Fig. 1: A graded quiver

Definition 2.7 *A **Bratteli diagram** is a finite graded quiver such that:*

*(i) there is a unique vertex with grading 0, called the **root**,*

(ii) if $v \in V(Q)$ is not the root then v is the target of at least one arrow,

(iii) if $v \in V(Q)$ does not have grading of maximum value then v is the source of at least one arrow,

(iv) for each $e \in E(Q)$, $gr(t(e)) = 1 + gr(s(e))$.

Example 2.8 *Note that the quiver of Figure 1 is not a Bratteli diagram. However, by removing the top arrow and adding an arrow from vertex 0 to the bottom vertex of grading 1, we produce a Bratteli diagram.*

Consider a group algebra chain $\mathbb{C}[G_n] > \mathbb{C}[G_{n-1}] > \cdots > \mathbb{C}[G_1] > \mathbb{C}[G_0] = \mathbb{C}$. To associate a Bratelli diagram to this chain we follow the language of [Ram97]. Let ρ be an irreducible representation of G_i, i.e., an irreducible $\mathbb{C}[G_i]$-module. Upon restriction to G_{i-1}, $\rho \downarrow_{G_{i-1}}$ decomposes as a direct sum of irreducible $\mathbb{C}[G_{i-1}]$-modules. For γ an irreducible representation of G_{i-1}, let $M(\rho, \gamma)$ denote the multiplicity of γ in $\rho \downarrow_{G_{i-1}}$.

Definition 2.9 *For a chain of group algebras $\mathbb{C}[G_n] > \mathbb{C}[G_{n-1}] > \cdots > \mathbb{C}[G_0]$, the **associated Bratteli diagram** consists of vertices of grading i labeled by the (equivalence classes of) irreducible representations of G_i, and each vertex labeled by an irreducible representation γ of G_{i-1} is connected to a vertex labeled by an irreducible representation ρ of G_i by $M(\rho, \gamma)$ arrows.*

Example 2.10 *Figure 2 shows two examples of Bratteli diagrams, with the gradings listed at the top. On the left we see the Bratteli diagram for a chain of group algebras for C_6 while on the right we see the Bratteli diagram for a chain of group algebras for the symmetric group S_3, viewing S_i as the subgroup of S_n that fixes the elements $\{i+1, \ldots, n\}$. Note that we distinguish $\mathbb{C}[S_1]$ from $\mathbb{C}$ only so that vertices at level i correspond to representations of $\mathbb{C}[S_i]$.*

For the group algebra $\mathbb{C}[C_N]$, irreducible representations are naturally indexed at each level by the integers $0, \ldots, N-1$, while for $\mathbb{C}[S_n]$, the irreducible representations are indexed by partitions of n.
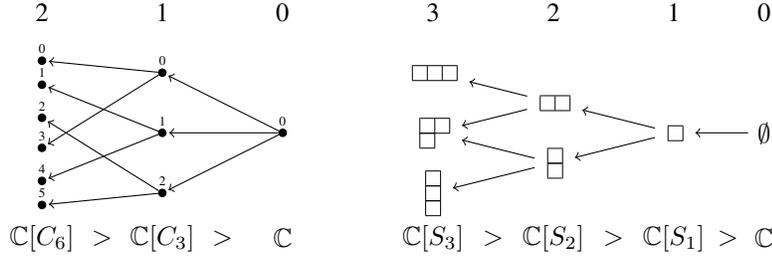
Fig. 2

Given a Bratteli diagram $\mathcal{B}$, there is a canonical chain of algebras associated to $\mathcal{B}$ called the chain of path algebras.

Definition 2.11 *Let $\mathcal{B}$ be a Bratteli diagram. The **path algebra (at level i)**, denoted $\mathbb{C}[\mathcal{B}_i]$, is the $\mathbb{C}$-vector space with basis given by ordered pairs of paths of length i in $\mathcal{B}$ which start at the root and end at the same vertex at level i.*

Example 2.12 *For the Bratteli diagram $\mathcal{B}$ of Figure 2 associated to the chain $\mathbb{C}[S_3] > \mathbb{C}[S_2] > \mathbb{C}[S_1] > \mathbb{C}$, let P_1, P_2, P_3, P_4 be the paths from the root to level 3 in $\mathcal{B}$, labeled from top to bottom. Then the path algebra $\mathbb{C}[\mathcal{B}_3]$ has basis $\{(P_1, P_1), (P_2, P_2), (P_2, P_3), (P_3, P_2), (P_3, P_3), (P_4, P_4)\}$.*

Note that for a vertex v, labeled by a representation ρ, the dimension of ρ is given by the number of paths from the root to v.

Further, $\mathbb{C}[\mathcal{B}_i]$ embeds into $\mathbb{C}[\mathcal{B}_{i+1}]$ as a subalgebra by mapping any pair of paths $(P, Q) \in \mathbb{C}[\mathcal{B}_i]$ to the sum

$$\sum_e (e \circ P, e \circ Q),$$

over all arrows e such that the source of e is the target of P (equivalently, of Q), and $\circ$ denotes concatenation of paths. Thus, elements in these subalgebras are effectively determined by the initial "legs" of their paths. This is also equivalent to a choice of basis in the corresponding Wedderburn decomposition of the group algebra as a direct sum of matrix algebras, recognizing that for a given element, a number (equal to the total number of distinct paths that have the common middle "source" of tail of P) of irreducible matrix elements will take on the same value. Identification of this kind of common "unit" (formalized by the injection of one quiver into another) is the fundamental observation and technique of the quiver-based SOV approach.

Multiplication in the path algebra $\mathbb{C}[\mathcal{B}_i]$ linearly extends $(P, Q) * (P', Q') = \delta_{QP'}(P, Q')$:

$$\sum_{(P,Q)} a_{PQ}(P, Q) * \sum_{(P',Q')} b_{P'Q'}(P', Q') = \sum \left(\sum_Q a_{PQ} b_{QQ'} \right) (P, Q').$$

Lemma 2.13 *Let $\mathbb{C}[G] = \mathbb{C}[G_n] > \mathbb{C}[G_{n-1}] > \cdots > \mathbb{C}[G_1] > \mathbb{C}[G_0] = \mathbb{C}$ be a chain of group algebras with Bratteli diagram $\mathcal{B}$. Then the chain of path algebras associated to $\mathcal{B}$ is isomorphic to the group algebra chain.*

Lemma 2.13 allows for Lemma 2.3 to be recast in terms of elements of the path algebra.

3   The Separation of Variables Approach

The first steps of the SOV approach involve expressing a path algebra element as a factorization over subsets of the Bratteli diagram in such a way as to disentangle the dependencies in the sum. To do so we first factor the Fourier transform through the subalgebras $\mathbb{C}[G_i]$. If we do this for a simple two-step chain, $\mathbb{C}[G] > \mathbb{C}[H] > \mathbb{C}$, we get a corresponding factorization

$$\mathcal{F} := \sum_{s \in G} f(s)\tilde{s} = \sum_{y \in Y} \sum_{h \in H} f(yh)\tilde{y}\tilde{h} = \sum_{y \in Y} \tilde{y} \sum_{h \in H} f(yh)\tilde{h} = \sum_{y \in Y} \tilde{y}F_y, \tag{2}$$

for Y a set of coset representatives for G/H such that for each $y \in Y$,

$$F_y = \sum_{h \in H} f_y(h)\tilde{h} \in \mathbb{C}[\mathcal{B}_H]$$

with $f_y(h) := f(yh)$.

   This factorization allows us to obtain a simple, but key complexity estimate: given a set of coset representatives Y for G/H with F_y (for each $y \in Y$) an arbitrary element of $\mathbb{C}[\mathcal{B}_H]$, define

$$m_G(R, Y, H) = \frac{1}{|G|} \times \{\text{minimum number of operations required to compute } \sum_{y \in Y} \tilde{y}F_y\}.$$

Then a restatement of Lemma 2.10 of [Mas98] and Proposition 1 of [DR90] gives

$$t_G(R) \leq t_H(R_H) + m_G(R, Y, H).$$

This bound shows that to compute the Fourier transform of a complex function defined on G we compute

$$\mathcal{F}_Y := \sum_{y \in Y} \tilde{y}F_y,$$

for Y a set of coset representatives for G/H. In doing so, the complexity estimate "reduces" to a close study of the computation of F_Y.

   The heart of the SOV approach is the efficient computation of $\mathcal{F}_Y$. It is comprised of three main steps:

1 Factor each term $\tilde{y}F_y$ and use these factorizations to rearrange $\mathcal{F}_Y$ into a recursive summation.

2 Each factor x_i of each term corresponds to an element of the path algebra, and thus a particular subgraph of the Bratteli diagram. These subgraphs are given a vector space structure through an identification with a space of quiver morphisms.

3 By virtue of the vector space identification, the element multiplication $x_i x_{i+1}$ becomes a bilinear map whose complexity can be calculated directly in terms of the number of certain subgraphs in the Bratteli diagram.

   To give the general idea, the "gluing" and summing operations that are multiplication in the path algebra mean that only certain kinds of "middle paths" contribute when two path algebra elements are multiplied. A complexity estimate thus becomes counting the number of subgraphs (subquivers) wherein this compatibility is respected – counting the occurrences of subquivers $\mathcal{Q}$ in the corresponding Bratteli diagram $\mathcal{B}$. Ultimately, this is the number of *morphisms* from $\mathcal{Q}$ into $\mathcal{B}$. We give a general example below.

Q_i
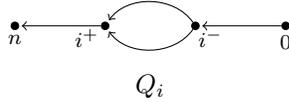
Fig. 4

Example 3.1 *Suppose $y \in \mathbb{C}[G]$ factors as $y = x_1 x_2$ with $x_i \in \mathbb{C}[G_{i+2}] \cap \text{Centralizer } \mathbb{C}[G_i]$. When we express x_i in in the path algebra as $\tilde{x}_i = \sum_{(P,Q)} [x_i]_{PQ}(P,Q)$, an application of Schur's Lemma shows that $[x_i]_{P,Q}$ is 0 unless P and Q are paths in $\mathcal{B}$ that agree from level n to level $i + 1$, and from level $i - 1$ to level 0, as in the quivers Q_i of the lefthand side of Figure 3 (see also [MR97a]). The product $\tilde{x}_1 \tilde{x}_2$ is indexed by a triple of paths resulting from gluing Q_2 to Q_1 by identifying the bottom path of Q_1 with the top path of Q_2, but maintaining the structures of Q_1 and Q_2 (the quiver on the righthand side of Figure 3). The complexity count is thus the careful counting of these compatible structures, which can be recast as the computation of the dimension of a space of quiver morphisms.*
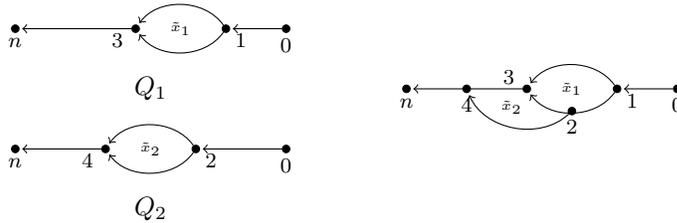
Fig. 3

For products with more factors we iterate this gluing process. Example 3.3 below gives further details. The SOV approach consists of factoring $\tilde{y}F_y$, forming the graph (akin to Fig. 3) and determining the subgraphs corresponding to each individual product.

Definition 3.2 *For a path algebra product $x_1 \cdots x_m$, let i^+ denote the smallest integer such that $x_i \in \mathbb{C}[\mathcal{B}_{i^+}]$ and let i^- denote the largest integer less than or equal to i^+ such that $x_i \in \text{Centralizer}(\mathbb{C}[\mathcal{B}_{i^-}])$. Then for $1 \le i \le m$ define*

$$X_i := \mathbb{C}[\mathcal{B}_{i^+}] \cap \text{Centralizer}(\mathbb{C}[\mathcal{B}_{i^-}]).$$

To each space X_i, associate the quiver Q_i of Figure 4. Note that Q_i is the quiver associated to every element of X_i. We show in the full paper [MRW] that X_i has dimension equal to the number of occurrences of Q_i in the Bratteli diagram $\mathcal{B}$. Denote this number by $\# \text{Hom}(Q_i; \mathcal{B})$. An "occurrence" of Q_i is the same as an injective map from Q_i into $\mathcal{B}$. Thus, $\# \text{Hom}(Q_i; \mathcal{B})$ is also the dimension of this space of morphisms of Q_i into $\mathcal{B}$.

In this setting (bilinear) group algebra multiplication is transformed into a bilinear map on products of associated spaces of quiver morphisms. Call this map $*$. As the notation and details are more technical than illuminating, we defer the explicit definition of $*$ and discussion of its properties to the full paper. However, even with deferring this we can present an example and the final algorithm.

$$Q_2 \cup Q_3 \qquad\qquad (Q_2 \triangle Q_3) \cup Q_1$$

Fig. 6

Example 3.3 *Suppose $\tilde{y} = x_1 x_2 x_3$, with*

$$
\begin{aligned}
1^+ &= 7, & 1^- &= 4, \\
2^+ &= 3, & 2^- &= 1, \\
3^+ &= 5, & 3^- &= 2.
\end{aligned}
$$

Figure 5 shows the quivers Q_i and the quiver $\mathcal{Q}$ formed by gluing Q_1 to Q_2 to Q_3.
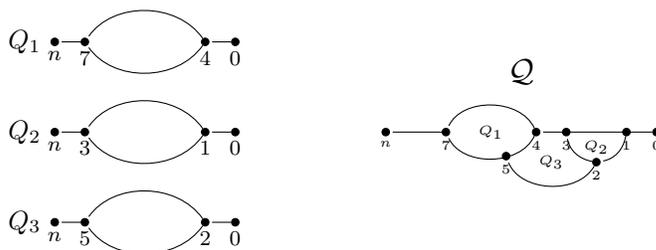
Fig. 5

    Translated to the space of quiver morphisms, the complexity of computing $x_1 x_2 x_3$ in the group algebra becomes the complexity of computing $w_1 * w_2 * w_3$, where w_i is in the space of morphisms from Q_i to $\mathcal{B}$. In this space, the complexity of computing $w_1 * w_2 * w_3$ is equivalent to the complexity of computing $w_{\sigma(1)} * w_{\sigma(2)} * w_{\sigma(3)}$, for $\sigma \in S_3$.

    For example, for $\sigma = (123)$, we determine the complexity of $w_2 * w_3 * w_1$. The complexity of $w_2 * w_3$ is $\# \operatorname{Hom}(Q_2 \cup Q_3; \mathcal{B})$, where $Q_2 \cup Q_3$ is as in Figure 6, the subquiver of $\mathcal{Q}$ corresponding to Q_2 and Q_3 (in Figure 6 we need only consider the subquiver formed by the segments of $Q_2 \cup Q_3$ where not all three of the paths agree). The complexity of $(w_2 * w_3) * w_1$ is $\# \operatorname{Hom}((Q_2 \triangle Q_3) \cup Q_1; \mathcal{B})$, where $Q_2 \triangle Q_3$ is the quiver of Figure 6 associated to the space containing $w_2 * w_3$. Note that as per the notation $Q_2 \triangle Q_3$ is in fact the symmetric difference *of Q_2 and Q_3, i.e., the edges of $Q_2 \cup Q_3$ not in $Q_2 \cap Q_3$.*

Lemma 3.4 *For Q_i (respectively, Q_j) the quiver associated to X_i (respectively, X_j), $x_i * x_j$ requires at most $\# \operatorname{Hom}(Q_i \cup Q_j; \mathcal{B})$ scalar multiplications and fewer additions.*

With Lemma 3.4 we now have our main general result. Choose $m \in \mathbb{N}$ and a subset $X \subseteq (\mathbb{C}[\mathcal{B}_n])^m = \mathbb{C}[\mathcal{B}_n] \times \cdots \mathbb{C}[\mathcal{B}_n]$ such that $|X| = |Y|$ and for each $y \in Y$ there exists $(x_1, \ldots, x_m) \in X$ with $\tilde{y} F_y = x_1 \cdots x_m$. Thus, X can be thought of as a choice of factorization into at most m elements of each term $\tilde{y} F_y$.

Theorem 3.5 *For x_i as above and $\sigma \in S_m$, let Q_i^σ denote the quiver associated to $w_i = x_{\sigma(i)}$. Let $W_i = \{(w_{i+1}, \ldots w_m) | (x_1 \ldots x_m) \in X\}$. Then we may compute $\sum_{y \in Y} \tilde{y} F_y$ in at most*

$$\sum_{i=1}^{m-1} |W_{i-1}| \# \operatorname{Hom}((Q_1^\sigma \triangle \cdots \triangle Q_i^\sigma) \cup Q_{i+1}^\sigma; \mathcal{B})$$

multiplications and fewer additions.

3.1 Morphisms into Bratteli Diagrams

Theorem 3.5 reduces complexity counts to determining $\# \operatorname{Hom}(Q; \mathcal{B})$ for Q a quiver as in Example 6 and $\mathcal{B}$ a Bratteli diagram. For Bratteli diagrams associated to group algebra chains we generalize Stanley's work on differential posets [Sta88] to deterime these counts.

Let $\mathbb{C}[V(\mathcal{B})]$ denote the space of finitely supported linear combinations of vertices of $\mathcal{B}$, let $\mathcal{B}^i$ denote the vertices $\alpha \in V(\mathcal{B})$ with $gr(\alpha) = i$, and let $\mathbb{C}[\mathcal{B}^i]$ denote the space of finitely supported linear combinations of vertices at level i in $\mathcal{B}$. Define an inner product $\langle \ , \ \rangle$ on $\mathbb{C}[V(\mathcal{B})]$ making the vertices orthonormal. For $\alpha, \beta \in V(B)$, let $M_B(\alpha, \beta)$ denote the number of paths from β to α in B. Note that because $\mathcal{B}$ is a Bratteli diagram, $\alpha, \beta \in V(\mathcal{B})$ correspond to irreducible representations γ, ρ and $M_{\mathcal{B}}(\alpha, \beta) = M(\gamma, \rho)$, as in Definition 2.9.

As in [Sta88] define linear operators U and D on $\mathbb{C}[V(\mathcal{B})]$ by linearly extending the action on $\alpha \in \mathcal{B}^i$:

$$U\alpha = \sum_{\gamma \in \mathcal{B}^{i+1}} M_{\mathcal{B}}(\gamma, \alpha)\gamma,$$

$$D\alpha = \sum_{\beta \in \mathcal{B}^{i-1}} M_{\mathcal{B}}(\alpha, \beta)\beta,$$

where, by convention, if $\mathcal{B}$ has highest grading n, $\mathcal{B}^{-1} = \emptyset = \mathcal{B}^{n+1} = \mathcal{B}^{n+2} = \cdots$.

Note 3.6 *As the vertices of $\mathcal{B}$ are labeled by the irreducible representations of $\mathbb{C}[\mathcal{B}_i]$, elements of $\mathbb{C}[\mathcal{B}^i]$ correspond to representations of the path algebra $\mathbb{C}[\mathcal{B}_i]$. In this context, U is induction and D restriction (see [GdlHJ89] Proposition 2.3.1).*

Let d_α denote the dimension of the representation corresponding to vertex α. Theorem 3.8 below generalizes Theorem 3.7 of [Sta88] and Theorem 2.3 of [Sta90].

Definition 3.7 *Let $w = w_l \cdots w_1$ be a word in U and D and let $\mathcal{S} = \{i \mid w_i = D\}$. For each $i \in \mathcal{S}$, let $a_i = \#\{D$'s in w to the right of $w_i\}$, and similarly let $b_i = \#\{U$'s in w to the right of $w_i\}$. If $b_i - a_i \geq 0$ for all $i \in \mathcal{S}$, we call w an **admissible word**.*

Theorem 3.8 *Let $\mathcal{B}$ be the Bratteli diagram associated to a group algebra chain and let $\lambda_i = |G_i/G_{i-1}|$. Then for $w = D^{d_n} U^{u_n} \cdots D^{d_1} U^{u_1}$ an admissible word in U and D, $s = \sum_{i=1}^n u_i - d_i$, and $\alpha \in \mathcal{B}^s$,*

$$\langle w\hat{0}, \alpha \rangle = d_\alpha \prod_{i \in \mathcal{S}} \lambda_{b_i - a_i}.$$

Definition 3.9 *A quiver Q is **n-toothed** if it consists of $2n+1$ (not necessarily distinct) vertices $\gamma_0, \ldots, \gamma_n$, $\beta_1, \ldots, \beta_n$ and distinct arrows connecting γ_{i-1} to β_i and γ_i to β_i.*
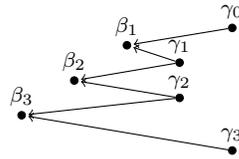
Fig. 7

Example 3.10 *The quiver of Figure 7 is an example of a 3-toothed quiver.*

Theorem 3.11 *Let $\mathcal{B}$ be a locally free Bratteli diagram, Q an n-toothed quiver with vertices γ_i at level l_i, β_i at level m_i. Then for $w = D^{m_n - l_n} U^{m_n - l_{n-1}} \cdots D^{m_1 - l_1} U^{m_1 - l_0}$,*

$$\# \operatorname{Hom}(Q; \mathcal{B}) = \sum_{\alpha \in \mathcal{B}^{l_n - l_0}} \langle w\hat{0}, \alpha \rangle = \prod_{i \in \mathcal{S}} \lambda_{b_i - a_i} \sum_{\alpha \in \mathcal{B}^{l_n - l_0}} d_\alpha.$$

Proof: Follows from Theorem 3.8 and induction. □

Theorem 3.11 together with Theorem 3.5 provides the complexity results of Theorems 1.1, 1.2, 1.3.

References

[AT79]   L. Auslander and R. Tolimieri. Is computing with the finite Fourier transform pure or applied mathematics? *Bull. Amer. Math. Soc. (N.S.)*, 1(6):847–897, 1979.

[Cla89]   M. Clausen. Fast generalized Fourier transforms. *Theoret. Comput. Sci.*, 67(1):55–63, 1989.

[CT65]   J. Cooley and J. Tukey. An algorithm for the machine calculation of complex Fourier series. *Math. Comp.*, 19:297–301, 1965.

[Dia80]   P. Diaconis. Average running time of the fast Fourier transform. *J. Algorithms*, 1:187–208, 1980.

[DMT13]   A. Danelakis, M. Mitrouli, and D. Triantafyllou. Blind image deconvolution using a banded matrix method. *Numer. Algorithms*, 64(1):43–72, 2013.

[DR90]   P. Diaconis and D. Rockmore. Efficient computation of the Fourier transform on finite groups. *J. Amer. Math. Soc.*, 3(2):297–332, 1990.

[GdlHJ89] F. Goodman, P. de la Harpe, and V. Jones. *Coxeter Graphs and Towers of Algebras*, volume 14 of *Mathematical Sciences Research Institute Publications*. Springer-Verlag, New York, 1989.

[JF07]   S. Johnson and M. Frigo. A modified split-radix FFT with fewer arithmetic operations. *IEEE Trans. Signal Process.*, 55(1):111–119, 2007.

[LVB07]   T. Lundy and J. Van Buskirk. A new matrix approach to real FFTs and convolutions of length 2^k. *Computing*, 80(1):23–45, 2007.

[Mas98]   D. Maslen. The efficient computation of Fourier transforms on the symmetric group. *Math. Comp.*, 67(223):1121–1147, 1998.

[MR95]   D. Maslen and D. Rockmore. Adapted diameters and FFTs on groups. In *Proc. 6th ACM-SIAM SODA*, pages 253–262. ACM, 1995.

[MR97a]   D. Maslen and D. Rockmore. Generalized FFTs—a survey of some recent results. In *Groups and computation, II (New Brunswick, NJ, 1995)*, volume 28 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 183–237. Amer. Math. Soc., Providence, RI, 1997.

[MR97b]   D. Maslen and D. Rockmore. Separation of variables and the computation of Fourier transforms on finite groups. I. *J. Amer. Math. Soc.*, 10(1):169–214, 1997.

[MR00]   D. Maslen and D. Rockmore. Double coset decompositions and computational harmonic analysis on groups. *Journal of Fourier Analysis and Applications*, 6(4):349–388, 2000.

[MRW]   David K. Maslen, Daniel N. Rockmore, and Sarah Wolff. Separation of variables and the computation of Fourier transforms on finite groups. II. preprint.

[Ram97]   A. Ram. Seminormal representations of Weyl groups and Iwahori-Hecke algebras. *Proc. London Math. Soc. (3)*, 75(1):99–133, 1997.

[Roc00]   D. Rockmore. The FFT: An algorithm the whole family can use. *Computing in Science and Eng.*, 2(1):60–64, January 2000.

[Sta88]   R. Stanley. Differential posets. *J. Amer. Math. Soc.*, 1(4):919–961, 1988.

[Sta90]   R. Stanley. Variations on differential posets. In *Invariant theory and tableaux (Minneapolis, MN, 1988)*, volume 19 of *IMA Vol. Math. Appl.*, pages 145–165. Springer, New York, 1990.

[TAL97]   R. Tolimieri, M. An, and C. Lu. *Algorithms for Discrete Fourier Transform and Convolution*. Signal Processing and Digital Filtering. Springer-Verlag, New York, second edition, 1997.

[Yav68]   R. Yavne. An economical method for calculating the discrete Fourier transform. *Proc. AFIPS Fall Joint Computer Conf.*, 33:115–125, 1968.