A Fast Program Generator of Fast Fourier Transforms

Michael Clausen and Meinard Müller

Universität Bonn, Institut für Informatik V,
Römerstr. 164, D-53117 Bonn, Germany
clausen@cs.uni-bonn.de, meinard@cs.uni-bonn.de

Abstract. Let G be a finite group of order n. By Wedderburn's Theorem, the complex group algebra $\mathbb{C}G$ is isomorphic to an algebra of block diagonal matrices: $\mathbb{C}G \simeq \oplus_{k=1}^{h} \mathbb{C}^{d_k \times d_k}$. Every such isomorphism D, a so-called discrete Fourier transform of $\mathbb{C}G$, consists of a full set of pairwise inequivalent irreducible representations D_k of $\mathbb{C}G$. A result of Morgenstern combined with the well-known Schur relations in representation theory show that (under mild conditions) any straight line program for evaluating a DFT needs at least $\Omega(n \log n)$ operations. Thus in this model, every $O(n \log n)$ FFT of $\mathbb{C}G$ is optimal up to a constant factor. For the class of supersolvable groups we will discuss a program that from a pc-presentation of G constructs a DFT $D = \oplus D_k$ of $\mathbb{C}G$ and generates an $O(n \log n)$ FFT of $\mathbb{C}G$. The running time to construct D is essentially proportional to the time to write down all the monomial (!) twiddle factors $D_k(g_i)$ where the g_i are the generators corresponding to the pc-presentation. Finally, we sketch some applications.

1 Introduction

This paper is concerned with fast discrete Fourier transforms. From an engineering point of view, there are two types of domains: a signal domain and a spectral domain. Algebraically, both domains are finite dimensional vector spaces (over the complex numbers, say) and, in addition, they are equipped with a multiplication which turns both domains into associative $\mathbb{C}$-algebras. The multiplication in the signal domain, called convolution, comes from the multiplication in a finite group, whereas the multiplication in the spectral domain is closely related to matrix multiplication. Fourier transforms isomorphically link these two domains and thus link convolution and matrix multiplication.

To be more specific, let G be a finite group. The set $\mathbb{C}G := \{a | a : G \to \mathbb{C}\}$ of all $\mathbb{C}$-valued functions (signals) on G becomes a vector space over $\mathbb{C}$ by pointwise addition and scalar multiplication. A natural basis is given by the indicator functions $(G \ni h \mapsto \delta_{gh})$ of the group elements $g \in G$. Identifying each group element with its indicator function, $\mathbb{C}G$ can be viewed as the $\mathbb{C}$-linear span of G, i.e., the span of all formal sums $\sum_{g \in G} a_g g$ with complex coefficients. The multiplication in G extends to the so-called convolution in $\mathbb{C}G$:

$$\Big(\sum_{g \in G} a_g g\Big) \cdot \Big(\sum_{h \in G} b_h h\Big) = \sum_{k \in G} \Big(\sum_{g \in G} a_g b_{g^{-1}k}\Big) k.$$

In this way, $\mathbb{C}G$ becomes a $\mathbb{C}$-algebra, the so-called group algebra of G over $\mathbb{C}$. For example, if $G = C_n = \langle X \mid X^n = 1 \rangle$ is the cyclic group of order n, then $\mathbb{C}G$ can be identified with the polynomial ring $\mathbb{C}[X]$ modulo the ideal generated by $X^n - 1$. In this case, convolution in $\mathbb{C}G$ means ordinary polynomial multiplication modulo the relation $X^n = 1$. If ω is a primitive nth root of unity, then the factorization $X^n - 1 = \prod_{j=0}^{n-1}(X - \omega^j)$ combined with the chinese remainder theorem shows that $\mathbb{C}C_n$ is isomorphic to the algebra $\oplus_{j=0}^{n-1}\mathbb{C}^{1\times 1}$ of n-square diagonal matrices. With respect to natural bases in both spaces this isomorphism is described by the classical DFT matrix $\mathbf{D} = (\omega^{jk})_{0\leq j,k<n}$.

Wedderburn's structure theorem for split semisimple algebras yields the right generalization of the above situation: according to this theorem, the complex group algebra $\mathbb{C}G$ is isomorphic to an algebra of block diagonal matrices,

$$D = \oplus_{k=1}^{h} D_k : \mathbb{C}G \longrightarrow \oplus_{k=1}^{h} \mathbb{C}^{d_k \times d_k}.$$

Here, the number h of blocks equals the number of conjugacy classes of G and the projections $D_1, \ldots, D_h$ form a complete set of pairwise inequivalent irreducible representations of $\mathbb{C}G$. Recall that a representation of $\mathbb{C}G$ of degree f is an algebra morphism $F\colon \mathbb{C}G \to \mathbb{C}^{f\times f}$. It is irreducible iff F is surjective. Two representations F_1, F_2 of degree f are equivalent, $F_1 \sim F_2$, if an invertible matrix X exists such that for all $a \in \mathbb{C}G:\ F_1(a) = XF_2(a)X^{-1}$. Every isomorphism D is called a discrete Fourier transform (DFT). If G is non-abelian, there are infinitely many DFTs of $\mathbb{C}G$. However, according to the Skolem-Noether theorem, if D and Δ are DFTs of $\mathbb{C}G$, then there are invertible matrices X_k such that $\Delta(a) = \oplus_k X_i D_k(a) X_k^{-1}$ for all $a \in \mathbb{C}G$. In the sequel, $\mathrm{DFT}(G)$ denotes the set of all DFT matrices of G.

From an algebraic point of view, performing a DFT, $\mathbb{C}G \ni a \mapsto D(a)$, amounts to evaluating a full set of pairwise inequivalent irreducible representations. In matrix terminology this amounts to multiplying the corresponding DFT matrix $\mathbf{D}$ by an input vector $\mathbf{a} := (a_g)_{g\in G}$. The linear complexity $L(\mathbf{D})$ of the DFT matrix is the minimum number of additions, subtractions, and scalar multiplications[1] to compute the matrix-vector product $\mathbf{D}\cdot\mathbf{a}$ for arbitrary $\mathbf{a} \in \mathbb{C}^{|G|}$. If the program constants are restricted to be of absolute value ≤ 2, then the corresponding minimum number $L_2(\mathbf{D})$ is called the 2-linear complexity of $\mathbf{D}$. The linear complexity $L(G)$ of the finite group G is defined by

$$L(G) := \min\{L(\mathbf{D}) \mid \mathbf{D} \in \mathrm{DFT}(G)\}.$$

Similarly, one defines $L_2(G)$. Trivially, $|G| - 1 \leq L(G) \leq 2|G|\cdot(|G| - 1)$, and $L(G) \leq L_2(G)$. A theorem by Morgenstern [7] combined with the classical Schur relations yield [2]: $L_2(G) > \frac{1}{4}|G|\log|G|$. Thus performing a DFT with only $O(|G|\log|G|)$ additions, subtractions, and scalar multiplications by small program constants is almost optimal in the L_2-model. This justifies the name Fast Fourier Transform. (For more details on lower complexity bounds, see Chapter

[1] In the classical FFT algorithms these program constants are the so-called twiddle factors.

13 of [4].) As an example, the Cooley-Tukey FFT of $\mathbb{C}G$, G a cyclic 2-group, uses only roots of unity as program constants and is thus essentially 2-optimal in the L_2-model.

Prior to performing a DFT one has to solve a fundamental problem in representation theory: up to equivalence, all irreducible representations of $\mathbb{C}G$ have to be generated. In general not an easy task! Even worse: as we are interested in a fast evaluation of $D = \oplus D_k$ we should choose the representatives D_k in the equivalence classes very carefully. In other words, we have to choose the right X_k above.

At least for the class of supersolvable groups, it turns out that choosing the right bases is quite easy. Moreover both problems (generating a DFT, performing a DFT) can be solved in an essentially optimal way. Furthermore, the results lead to fast algorithms not only in theory but also in a practical sense. The aim of this paper is to present these results and some of its consequences without giving too much technical details. For more information we refer to [1, 2, 4, 5, 8, 12].

The rest of this paper is organized as follows. After some preparations, we describe in Section 2 geometrically a monomial DFT for supersolvable groups and indicate how this yields FFTs for supersolvable groups. Section 3 presents the main ideas of the algorithm that constructs monomial DFTs. Furthermore, some implementation details and running times are shown. Section 4 sketches some applications.

2 Monomial DFTs for Supersolvable Groups

A finite group G is called supersolvable iff there exists a chain

$$\mathcal{T} = (G = G_n \supset G_{n-1} \supset \ldots \supset G_1 \supset G_0 = \{1\})$$

such that each G_i is a normal subgroup in G and all indices $[G_i : G_{i-1}] =: p_i$ are prime. Thus $\mathcal{T}$ is a chief series of G with chief factors G_i/G_{i-1} of prime order. For example, all nilpotent groups (especially all groups of prime power order) are supersolvable.

In this section we are going to describe the irreducible representations of supersolvable groups in a geometric way. This approach will be the theoretical basis of the algorithmic approach shown in the next section. We need some preparations.

2.1 Basic Definitions and Tools

The character χ of a representation D of $\mathbb{C}G$ is defined by $\chi(g) := \mathrm{Trace}(D(g))$, for $g \in G$. Characters are constant on conjugacy classes and two representations are equivalent iff their characters are equal. Characters corresponding to irreducible representations are called irreducible characters. A character is called linear iff it corresponds to a representation of degree 1. By $\mathrm{Irr}(G)$

we denote the set of all irreducible characters of G. The space $CF(G, \mathbb{C})$ of all complex-valued class functions on G becomes an inner product space by $\langle \chi | \psi \rangle := |G|^{-1} \sum_{g \in G} \chi(g) \overline{\psi(g)}$. For a proof of the following facts we refer to [9]:

Theorem 1. *For a finite group G with h conjugacy classes the following is true:*

(1) $\mathrm{Irr}(G) = \{\chi_1, \ldots, \chi_h\}$ *is an orthonormal basis of $CF(G, \mathbb{C})$.*
(2) *Let F and D_k be representations of $\mathbb{C}G$ with characters χ and χ_k, respectively. If D_k is irreducible, then the multiplicity $\langle D_k | F \rangle$ with which D_k occurs in F equals $\langle \chi_k | \chi \rangle$.*
(3) *If $e_k := e_{\chi_k} := \frac{\chi_k(1)}{|G|} \sum_{g \in G} \chi_k(g^{-1}) g$, then $e_1, \ldots, e_h$ are a basis of the center of $\mathbb{C}G$. Moreover, $1 = e_1 + \ldots + e_h$ and $e_k e_j = \delta_{kj} e_k$. (The e_k are called central primitive idempotents in $\mathbb{C}G$.)*
(4) *If M is a left $\mathbb{C}G$-module affording the representation F with character χ, then $M = \oplus_{k=1}^h e_k M$ (isotypic decomposition). If M_k is a simple module affording the character χ_k, then $e_k M$, the isotypic component of type χ_k, is isomorphic to the direct sum of $\langle \chi_k | \chi \rangle$ copies of M_k. Every simple submodule of M affording the character χ_k is contained in $e_k M$.*

Let H be a subgroup of G. Then $\mathbb{C}H$ can be viewed as a subalgebra of $\mathbb{C}G$. If D is a representation of $\mathbb{C}G$, then its restriction to $\mathbb{C}H$ is a representation of $\mathbb{C}H$ denoted by $D \downarrow H =: F$. In turn, D is called an extension of F. Similarly, $\chi \downarrow H$ denotes the restriction of the character χ.

One important tool in constructing representations is the process of induction, where a representation of a group G is constructed from a representation of a subgroup H. In terms of modules, the construction is straightforward: let L be a left ideal in $\mathbb{C}H$. Then $\mathbb{C}GL$ is a left ideal in $\mathbb{C}G$ and with $G = \sqcup_{i=1}^r g_i H$ one obtains the decomposition $\mathbb{C}GL = \oplus_{i=1}^r g_i L$ as a $\mathbb{C}$-space. In particular, $\mathbb{C}GL$ has dimension $[G : H] \cdot \dim L$. The left $\mathbb{C}G$-module $\mathbb{C}GL$ is said to be induced by L. A look at the corresponding matrix representations leads to the following definition. Let H be a subgroup of the group G, $T := (g_1, \ldots, g_r)$ a transversal of the left cosets of H in G and let F be a representation of $\mathbb{C}H$ of degree f. The induced representation $F \uparrow_T G$ of $\mathbb{C}G$ of degree $f \cdot r$ is defined for $x \in G$ by

$$F \uparrow_T G(x) := (\dot{F}(g_i^{-1} x g_j))_{1 \leq i,j \leq r} \in (\mathbb{C}^{f \times f})^{r \times r},$$

where $\dot{F}(y) := F(y)$ if $y \in H$, and $\dot{F}(y)$ is the f-square all zero matrix, if $y \in G \setminus H$. It is easily checked that this defines a representation of $\mathbb{C}G$. Taking different transversals gives possibly different, but equivalent representations. Thus in non-critical situations we sometimes write $F \uparrow G$ instead of $F \uparrow_T G$. Note that $F \uparrow_T G(x)$, for $x \in G$, is a block matrix, with exactly one non-zero block in each block row and in each block column. In particular, if F is of degree 1, then, for all $x \in G$, the matrix $F \uparrow_T G(x)$ is monomial. (Recall that a matrix is called monomial iff it has exactly one non-zero entry in each row and in each column. A representation D of $\mathbb{C}G$ is said to be monomial iff $D(g)$ is monomial, for all $g \in G$.) A group G is called an M-group if every irreducible representation is equivalent to a monomial one. Below we will give an alternative proof to the well-known fact that supersolvable groups are M-groups. There is a close connection between restriction and induction. A more precise statement reads as follows.

Theorem 2 (Frobenius Reciprocity Theorem). *Let H be a subgroup of G. Furthermore, let F and D be irreducible representations of $\mathbb{C}H$ and $\mathbb{C}G$, respectively. Then the multiplicity $\langle F|D{\downarrow}H\rangle$ of F in $D{\downarrow}H$ equals the multiplicity $\langle D|F{\uparrow}G\rangle$ of D in $F{\uparrow}G$: $\langle F|D{\downarrow}H\rangle = \langle D|F{\uparrow}G\rangle$.*

If N is a normal subgroup of G and F a representation of $\mathbb{C}N$, then for $g \in G$ we define a new representation F^g of $\mathbb{C}N$ by $F^g(n) := F(g^{-1}ng)$ for all $n \in N$. F and F^g are called conjugate representations. As $\{f(n)|n \in N\} = \{F^g(n)|n \in N\}$, F is irreducible iff F^g is irreducible, and G acts on $\mathrm{Irr}(N)$ by conjugation via $g\chi := (N \ni n \mapsto \chi(g^{-1}ng))$. The last tool needed for our geometric approach is the following special case of Clifford theory.

Theorem 3 (Clifford's Theorem). *Let N be a normal subgroup in G of prime index p, and let F be an irreducible representation of $\mathbb{C}N$. For a fixed $g \in G\setminus N$, let T denote the transversal $(1, g, g^2, \ldots, g^{p-1})$ of the cosets of N in G. Then exactly one of the following two cases applies.*

(1) *All F^{g^i} are equivalent. Then there are exactly p irreducible representations $D_0, \ldots, D_{p-1}$ of $\mathbb{C}G$ extending F. The D_k are pairwise inequivalent and satisfy $F{\uparrow}G \sim D_0 \oplus \ldots \oplus D_{p-1}$. Moreover, if $\chi^0, \chi^1, \ldots, \chi^{p-1}$ are the linear characters of the cyclic group G/N in a suitable order, we have $D_k = \chi^k \otimes D_0$ for all k, i.e., $D_k(x) = \chi^k(xN)D_0(x)$, for all $x \in G$.*

(2) *The representations F^{g^i} are pairwise inequivalent. In this case, the induced representation $F{\uparrow}G$ is irreducible.*

To decide which case in Clifford's Theorem applies, we work with intertwining spaces. Recall that for two representations D and Δ of $\mathbb{C}G$, both of degree d, the intertwining space is defined by

$$\mathrm{Int}(D, \Delta) := \{X \in \mathbb{C}^{d\times d}|XD(g) = \Delta(g)X, \text{ for all } g \in G\}.$$

2.2 A Geometric Approach

Let $\mathcal{T} = (G_n \supset G_{n-1} \supset \ldots \supset G_0)$ be a chief series of the supersolvable group $G = G_n$ of exponent e and let $\chi \in \mathrm{Irr}(G)$ be a fixed irreducible character of G. We are going to associate to χ a simple left $\mathbb{C}G$-module M affording χ and a basis in M such that the resulting representation D is monomial, i.e., D is monomial and each non-zero entry of $D(g)$, $g \in G$, is an e-th root of unity. To this end, we consider all sequences $w = (\chi_0, \ldots, \chi_n = \chi)$, with irreducible characters $\chi_i \in \mathrm{Irr}(G_i)$. By Theorem 1 (4) and Frobenius Reciprocity the product

$$e(w) := e_{\chi_0} \cdot \ldots \cdot e_{\chi_n}$$

of the corresponding central primitive idempotents is non-zero iff all multiplicities $\langle \chi_i | \chi_{i+1} \rangle := \langle \chi_i | \chi_{i+1} {\downarrow} G_i \rangle$ are positive. The set of all those sequences will be denoted by $W(\chi)$:

$$W(\chi) := \{(\chi_0, \ldots, \chi_n = \chi) \mid \chi_i \in \mathrm{Irr}(G_i),\ \langle \chi_i | \chi_{i+1} \rangle > 0\}.$$

Theorem 4 ([6]). *Let G be a supersolvable group of exponent e, χ an irreducible character of G, and $\mathcal{T}$ a chief series of G. Then the following holds.*

(1) *Let $w \in W(\chi)$. Then $e(w)$ is a primitive idempotent in $\mathbb{C}G$ and $\mathbb{C}Ge(w)$ is a simple $\mathbb{C}G$-module affording the character χ. The dimension of $\mathbb{C}Ge(w)$ equals $\chi(1) = |W(\chi)|$.*

(2) *Let $v \in W(\chi)$ and set $M := \mathbb{C}Ge(v)$. Then $M = \oplus_{w \in W(\chi)} e(w)M$ is a decomposition of M into 1-dimensional linear subspaces $e(w)M$.*

(3) *G acts transitively on $W(\chi)$ by $g(\chi_0, \ldots, \chi_n) := (g\chi_0, \ldots, g\chi_n)$.*

(4) *For all $g \in G$ and all $w \in W(\chi)$ we have $ge(w)g^{-1} = e(gw)$.*

(5) *Let U be the stabilizer of $v \in W(\chi)$ and L a transversal of the left cosets of U in G. Then $\{\ell e(v) | \ell \in L\}$ is a $\mathbb{C}$-basis of $\mathbb{C}Ge(v)$ and the corresponding representation D of $\mathbb{C}G$ is e-monomial. More precisely, if the 1-dimensional $\mathbb{C}U$-module $e(v)\mathbb{C}Ge(v)$ affords the linear character λ, then $e(v)$ equals the central primitive idempotent corresponding to λ: $e(v) = e_\lambda$, and $D = \lambda{\uparrow}_L G$.*

Proof. (1). $e_\chi = (\prod_{i<n} 1)e_\chi = \prod_{i<n}(\sum_{\chi_i \in \mathrm{Irr}(G_i)} e_{\chi_i})e_\chi = \sum_{w \in W(\chi)} e(w)$. Thus e_χ is the sum of $\chi(1)$ pairwise orthogonal idempotents $e(w)$; thus all $e(w)$ are primitive.

(2). Let $M = \mathbb{C}Ge(v)$ and $w \in W(\chi)$. Then $e(w)$ applied to M causes successive isotypic decompositions of M along $\mathcal{T}$: $e_{\chi_0} \cdot (e_{\chi_1} \cdot (\ldots \cdot (e_{\chi_{n-1}} M) \ldots))$. As $\langle \chi_i | \chi_{i+1} \rangle = 1$ (by Clifford's Theorem), $e(w)M$ is a simple $\mathbb{C}G_0$-module, hence one-dimensional.

(3). By Clifford's Theorem, G acts transitively on the irreducible constituents of $\chi \downarrow G_{n-1}$. Observing that G_{n-1} acts trivially on $\mathrm{Irr}(G_{n-1})$, an induction on n yields our claim.

(4). This follows from $ge_{\chi_i}g^{-1} = e_{g\chi_i}$, for all $\chi_i \in \mathrm{Irr}(G_i)$.

(5). By (3) and (4), G acts transitively on the set of lines $\{e(w)M | w \in W(\chi)\}$ according to $ge(w)M = e(gw)gM = e(gw)M$. Choosing any nonzero vector $x_w \in e(w)M$ yields a basis $(x_w)_{w \in W(\chi)}$ of M, and, by (2), the corresponding matrix representation is monomial. Now we choose the x_w in such a way that the non-zero entries in the representation matrices to the group elements are all eth roots of unity. To this end, let $U \leq G$ denote the stabilizer of $v \in W(\chi)$ and L a left coset transversal of U in G. As for $g \in G$, $0 \neq ge(v) = e(gv)ge(v) \in e(gv)M$ and G acts transitively on $W(\chi)$, the set $\{ge(v)|g \in L\}$ is a $\mathbb{C}$-basis of M. As U stabilizes the line $e(v)M = \mathbb{C}e(v)$, there exists a linear character λ of U such that $ue(v) = \lambda(u)e(v)$, for all $u \in U$. Now let $e_\lambda = |U|^{-1} \sum_{u \in U} \lambda(u^{-1})u \in \mathbb{C}U$ denote the central primitive idempotent corresponding to λ. Then $e_\lambda e(v) = e(v) = e(v)e_\lambda$, and hence $\mathbb{C}Ge(v) \leq \mathbb{C}Ge_\lambda$. Thus $e_\lambda = ae(v)$, for some $a \in \mathbb{C}G$. But then, $e(v) = e_\lambda e(v) = ae(v)e(v) = ae(v) = e_\lambda$. □

The above result suggests to introduce the $\mathcal{T}$-character graph of G. This graph has $n+1$ levels. The nodes of level i are the irreducible characters of G_i. Edges do exist at most between nodes of adjacent levels. More precisely, there is an edge between $\chi_i \in \mathrm{Irr}(G_i)$ and $\chi_{i+1} \in \mathrm{Irr}(G_{i+1})$ iff $\langle \chi_i | \chi_{i+1} \rangle > 0$ (note that $\langle \chi_i | \chi_{i+1} \rangle = 1$ for supersolvable groups). In addition, it is very convenient to know the action of G on each level. For this it suffices to take one element $g_j \in G_j \setminus G_{j-1}$ for each j and specify the action of $g_n, \ldots, g_{i+1}$ on $\mathrm{Irr}(G_i)$. (Note that all g_k, $k \leq i$, act trivially on $\mathrm{Irr}(G_i)$.) Figure 1 in Subsection 3.3 shows the character graph of a group of order 128.

2.3 FFTs for Supersolvable Groups

According to the last subsection we already know that a supersolvable group G of exponent e has an e-monomial DFT $D = \oplus_{k=1}^{h} D_k$. The construction of the D_k was along $\mathcal{T}$. Now we look at such a DFT D from a more algorithmic point of view.

Definition 1. *Let $\mathcal{T} = (G = G_n \supset \ldots G_0 = \{1\})$ be a chain of subgroups of the finite group G. A representation D of $\mathbb{C}G$ is called $\mathcal{T}$-adapted iff for all $0 \leq i \leq n$ the following conditions hold:*

(1) *The restriction $D{\downarrow}G_i$ is equal to the direct sum of irreducible representations of $\mathbb{C}G_i$, i.e., $D{\downarrow}G_i = \oplus_q F_{iq}$, with irreducible representations F_{iq}.*
(2) *Equivalent irreducible constituents of $D{\downarrow}G_i$ are equal, i.e., if $F_{iq} \sim F_{it}$ then $F_{iq} = F_{it}$ (but not necessarily $q = t$).*

If D is $\mathcal{T}$-adapted then for all $i \leq n$, $D{\downarrow}G_i$ is $\mathcal{T}_i$-adapted, where $\mathcal{T}_i$ denotes the chain $(G_i \supset \ldots \supset G_0)$. It is not hard to show that the above constructed monomial DFTs for supersolvable groups are in fact $\mathcal{T}$-adapted, see, e.g. [BCS,p.337]. Now we can state the main result of this subsection.

Theorem 5 (Baum [1]). *If G is a supersolvable group with chief series $\mathcal{T}$, then any $\mathcal{T}$-adapted DFT of $\mathbb{C}G$ is monomial and can be evaluated with at most $\gamma \cdot |G| \cdot \log |G|$ operations, where $1.5 \leq \gamma \leq 8.5$ depends on the prime divisors of $|G|$.*

Proof. (Sketch) Let $[G_n : G_{n-1}] = p$ and $D^n = \oplus D_k$ a $\mathcal{T}$-adapted monomial DFT of $\mathbb{C}G_n$. Let $F_1, \ldots, F_r$ be the distinct irreducible constituents of $D{\downarrow}G_{n-1}$. Then $D^{n-1} = \oplus_{\ell=1}^{r} F_\ell$ is a monomial, $\mathcal{T}_{n-1}$-adapted DFT of $\mathbb{C}G_{n-1}$. As copying is free in our model, $L(D^n{\downarrow}G_{n-1}) = L(D^{n-1})$.

Instead of evaluating D^n at $a \in \mathbb{C}G_n$ directly, we rewrite a according to the coset decomposition $G_n = \sqcup_{j<p} g^j G_{n-1}$. Then for suitable $a_j \in \mathbb{C}G_{n-1}$ we have $a = \sum_{j<p} g^j a_j$. Hence $D^n(a) = \sum_{j<p} D^n(g^j)(D^n{\downarrow}G_{n-1})(a_j)$. This formula suggests a divide-and-conquer strategy. In the divide-step, we evaluate the p "smaller" DFTs $D^{n-1}(a_j)$. By a tricky application of Clifford's Theorem combined with local FFTs of size p to handle simultaneously the cases of p extensions, the conquer-step is managed in such a way that altogether an $O(|G| \log |G|)$ upper bound is obtained. □

According to Theorem 4 and Theorem 5, a $\mathcal{T}$-adapted DFT $D = \oplus_{k=1}^{h} D_k$ of $\mathbb{C}G$ is essentially unique. In the sequel, we sometimes write $\mathrm{Irrep}(G, \mathcal{T}) = \{D_1, \ldots, D_h\}$ and similarly $\mathrm{Irrep}(G_i, \mathcal{T}_i)$.

3 Efficient Construction of monomial DFTs

In this section we give a summary of the algorithm in [3] which constructs a monomial DFT of a supersolvable group G given by a pc-presentation with $O(|G| \log |G|)$ operations. One can even show, that the running time is essentially proportional to the output length. For a detailed description and analysis of this algorithm we refer to [3].

3.1 PC-Presentations

Let G be a supersolvable group with chief series $\mathcal{T}$ as above. For $1 \leq i \leq n$ let g_i be an element in G_i not in G_{i-1}. With respect to $(g_1, \ldots, g_n)$ each element $g \in G$ can be expressed uniquely in normal form

$$g = g_n^{e_n} \cdot g_{n-1}^{e_{n-1}} \cdot \ldots \cdot g_1^{e_1} \quad (0 \leq e_i < p_i).$$

The multiplication in G is completely described, if the normal forms of all powers $g_i^{p_i}$ and all commutators $[g_i, g_j] := g_i^{-1} g_j^{-1} g_i g_j$ are known. More formally, every supersolvable group has a power-commutator presentation (pc-presentation) of the form

$$G = \langle g_1, \ldots, g_n \mid g_i^{p_i} = u_i \ (1 \leq i \leq n), [g_i, g_j] = w_{ij} \ (1 \leq i < j \leq n) \rangle,$$

with primes p_i as well as words $u_i \in G_{i-1}$ and $w_{ij} \in G_i$, all given in normal form. Moreover, we require the presentation to be consistent, i.e., that every word in the generators has a unique normal form. Consistent pc-presentations of this kind exactly describe the class of supersolvable groups.

With respect to such a pc-presentation, an irreducible representation of the group G_i is fully described by the representing matrices of the generators $g_i, \ldots, g_1$.

As an example, we give a consistent pc-presentation of a supersolvable group with 128 elements denoted by G_{128}. In the presentation, trivial commutator relations are omitted.

$$\begin{aligned} G_{128} = \langle g_7, g_6, g_5, g_4, g_3, g_2, g_1 \mid & g_1^2 = g_2^2 = g_4^2 = g_5^2 = g_6^2 = 1, g_3^2 = g_1, g_7^2 = g_4, \\ & [g_2, g_6] = [g_2, g_7] = [g_3, g_4] = [g_3, g_5] = [g_3, g_6] = g_1, [g_3, g_7] = g_2, \\ & [g_4, g_5] = g_2 \cdot g_1, [g_4, g_6] = g_3 \cdot g_1, [g_5, g_7] = g_3, [g_6, g_7] = g_5 \rangle \end{aligned}$$

3.2 The Algorithm

Before describing the algorithm, we want to mention the following important points. First, the pc-presentation of G already contains all the information on the group needed in the algorithm, so no group operations are required at all.

Second, even though the irreducible representations are computed over $\mathbb{C}$, it turns out that the algorithm uses just integer arithmetic. Hence, we never run into numerical problems! More precisely, all matrices to be processed by the algorithm are e-monomial and all matrix manipulations are multiplications. Therefore, we can compute in the additive group $\mathbb{Z}_e$, which is isomorphic to the group of eth roots of unity in $\mathbb{C}$. (One can show that the algorithm works over any field K containing such a primitive eth root of unity, but, for simplicity, we just consider the case $K = \mathbb{C}$.)

The central idea of the algorithm is based on Clifford's Theorem. In our notation it says that given an irreducible representation F of $\mathbb{C}G_{i-1}$, $0 < i \leq n$, then there are two cases:

Case 1. F extends to $p_i = [G_i : G_{i-1}]$ pairwise inequivalent irreducible representations of $\mathbb{C}G_i$ of the same degree $\deg(F)$.

Case 2. The induction of F is an irreducible representation of $\mathbb{C}G_i$ of degree $p_i \cdot \deg(F)$.

Furthermore, up to equivalence all irreducible representations of $\mathbb{C}G_i$ can be obtained this way. This allows us to construct the irreducible representations of $\mathbb{C}G$ iteratively in a bottom-up fashion along the chief series $\mathcal{T}$. However, constructing an arbitrary DFT is not what we want. We are interested in the construction of a very special set of irreducible representations - namely representations resulting in e-monomial matrices when evaluated at group elements. Suppose, we already have constructed a full set of nonequivalent irreducible e-monomial representations of $\mathbb{C}G_{i-1}$ denoted by $\mathcal{F}$. In order to construct an e-monomial $D \in \mathrm{Irrep}(G_i, \mathcal{T}_i)$ of level i from a given e-monomial $F \in \mathcal{F}$ of level $i-1$, we need to know the relation between the conjugate representation F^{g_i} and the corresponding $\tilde{F} \in \mathcal{F}$ with $F^{g_i} \sim \tilde{F}$. That is the reason, why the intertwining spaces $\mathrm{Int}(F^{g_i}, \tilde{F})$ come into play. It turns out that all intertwining matrices in $\mathrm{Int}(F^{g_i}, \tilde{F})$ are scalar multiples of an e-monomial matrix. In a second phase, the algorithm computes such intertwining matrices. To cut a long story short, we now give a summary of the algorithm. At level i the algorithm takes the following input:

Phase 1. $\mathcal{F} = \mathrm{Irrep}(G_{i-1}, \mathcal{T}_{i-1})$, i.e., a full set of nonequivalent irreducible e-monomial representations of $\mathbb{C}G_{i-1}$ such that $\bigoplus_{F \in \mathcal{F}} F$ is $\mathcal{T}_{i-1}$-adapted.

Phase 2. For every $i - 1 < j \leq n$ a permutation π_j of $\mathcal{F}$ such that $F^{g_j} \sim \pi_j F$ for all $F \in \mathcal{F}$ as well as e-monomial matrices $X_{jF} \in \mathrm{Int}(F^{g_j}, \pi_j F)$, $F \in \mathcal{F}$.

The following output is computed:

Phase 1. $\mathcal{D} = \mathrm{Irrep}(G_i, \mathcal{T}_i)$, i.e., a full set of nonequivalent irreducible e-monomial representations of $\mathbb{C}G_i$ such that $\bigoplus_{D \in \mathcal{D}} D$ is $\mathcal{T}_i$-adapted.

Phase 2. For every $i < j \leq n$ a permutation τ_j of $\mathcal{D}$ such that $D^{g_j} \sim \tau_j D$ for all $D \in \mathcal{D}$ as well as e-monomial matrices $Y_{jD} \in \mathrm{Int}(D^{g_j}, \tau_j D)$, $D \in \mathcal{D}$.

Note that the input of level 0 is trivial, all intertwining matrices being set to 1. Level i of the algorithm consists of two phases.

Phase 1 (Computation of $\mathcal{D}$). Consider $F \in \mathcal{F}$ and its g_i-conjugate representation F^{g_i}.
Case 1. $F \sim F^{g_i}$, i.e., $\pi_i F = F$. Then, by Clifford's Theorem, there are exactly $p := p_i$ pairwise nonequivalent irreducible extensions $D_0, \ldots, D_{p-1}$ of F to $\mathbb{C}G_i$ satisfying $D_k{=}\chi^k \otimes D_0$, where $\chi^0, \chi^1, \ldots, \chi^{p-1}$ are the irreducible characters of the cyclic group G_i/G_{i-1}. Since $D_k \downarrow G_{i-1} = F$, $k = 1, \ldots, p-1$, in this step only the $D_k(g_i)$ have to be computed. One can show that $D_k(g_i) \in \mathrm{Int}(F^{g_j}, F)$ and $c^p X_{iF}^p = F(g_i^p)$ with a constant $c \in \mathbb{C}^*$. The last equation has p distinct solutions $c_0, \ldots, c_{p-1} \in \mathbb{C}^*$, which can be proven to be even eth roots of unity. Thus the desired e-monomial matrices $D_k(g_i)$, $0 \leq k < p$, just differ by a factor, which is a power of a p_ith root of unity, and are given by $D_k(g_i) := c_k X_{iF}$.

Case 2. $F \not\sim F^{g_i}$, i.e., $\pi_i F \neq F$. Again, by Clifford's Theorem, the induced representation $F \uparrow G_i$ is irreducible and $(F \uparrow G_i) \downarrow G_{i-1} = \bigoplus_{k=0}^{p-1} F^{g_i^k}$. As $F^{g_i^k} \sim \pi_i^k F$, we know the existence of a unique irreducible representation D of $\mathbb{C}G$, sucht that $D \downarrow G_{i-1} = \bigoplus_{k=0}^{p-1} \pi_i^k F$. This $\mathcal{T}_i$-adapted representation is now to be computed. We already know $D(g_1), \ldots, D(g_{i-1})$ from level $i-1$. Thus it remains to specify $D(g_i)$. Here, the intertwining spaces constructed in level $i-1$ are to be used. If $X_k := X_{i\pi_i^{k-1}F} \cdot \ldots \cdot X_{iF}$, $0 \leq k < p$, then

$$D(g_i) = \begin{bmatrix} & & & & Z \\ X_1 X_0^{-1} & & & & \\ & X_2 X_1^{-1} & & & \\ & & \ddots & & \\ & & & X_{p-1} X_{p-2}^{-1} & \end{bmatrix}$$

where $Z := X_0 F(g_i^p) X_{p-1}^{-1}$, as shown in [3].

By these two constructions, all irreducible representations of G_i up to isomorphism are obtained, and Phase 1 is complete. In addition, during the construction in Phase 1 a bipartite graph is built up in which $F \in \mathcal{F}$ and $D \in \mathcal{D}$ are linked if and only if F is a constituent of $D \downarrow G_{i-1}$. This "traceback" information is needed in the next phase. Furthermore, this information, collected over all levels $i = 1, \ldots, n$, is nothing else but the $\mathcal{T}$-character graph of the group G.

Phase 2 (Computation of τ_j and Y_{jD}). Let $F \in \mathcal{F}$ and $i < j \leq n$. We have to consider the same two cases as in Phase 1.
Case 1. $\pi_i F = F$. In Phase 1, the p extensions $D_0, \ldots, D_{p-1}$ have been computed. As D_k is an extension of F, one can show that $\tau_j D$ must be an extension of $\pi_j F$. Let $\varDelta_0, \ldots, \varDelta_{p-1}$ be the extension of $\pi_j F$. Then it can be shown that $Y_{jD_k} := X_{jF}$ must be set for all k and $\tau_j(\{D_0, \ldots, D_{p-1}\}) = \{\varDelta_0, \ldots, \varDelta_{p-1}\}$. Using Y_{jD_k} one can determine τ_j as is explained in [3].
Case 2. $\pi_i F \neq F$. In this case, $\tau_j(D)$ can be immediately determined, since it equals the unique $\varDelta \in \mathcal{D}$ such that $\varDelta \downarrow G_{i-1}$ contains $\pi_j F$ (this information is encoded in the bipartite graph built up in Phase 1). We don't want to discuss here the construction of Y_{jD}, which is a bit delicate, but refer to [3].

3.3 An Example: G_{128}

Figure 1 shows the character graph of the group G_{128} given by its pc-presentation in Subsection 3.1. Each node represents an irreducible character and its corresponding irreducible representation. The numbers on the left hand side indicate the levels and the numbers on the top are the degrees of the corresponding reprentations of the top level. To illustrate the above algorithm, we describe the construction of the irreducible representation of level 7, denoted by D, corresponding to the circled node in Figure 1.

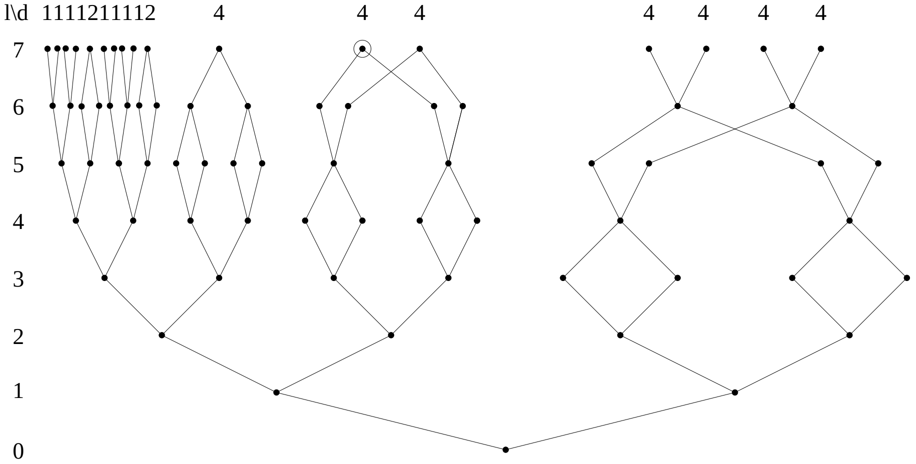

Fig. 1. Character graph of G_{128}.

The representation D is induced, let's say by the representation F of level 6, and is constructed in Phase 1 Case 2 of the algorithm. Suppose the algorithm has already constructed all data up to level 6 including F and the intertwining matrix X_{7F}. Since $p = p_7 = 2$, we have $D \downarrow G_6 = F \oplus \pi_7 F$ and D is known at the generators $g_1, \ldots, g_6$. We can compute $D(g_7)$ by the formula

$$D(g_7) = \begin{bmatrix} & F(g_7^2)X_1^{-1} \\ X_1 X_0^{-1} & \end{bmatrix} = \begin{bmatrix} & F(g_4)X_{7F}^{-1} \\ X_{7F} & \end{bmatrix} = \begin{bmatrix} & & 1 & \\ & & & \omega^4 \\ 1 & & & \\ & 1 & & \end{bmatrix},$$

where ω denotes an 8th root of unity, $e = 8$ being the exponent of G_{128}. Here we have used that X_0 is always the identity, $X_1 = X_{7F}$ is in our specific example also the identity matrix of dimension 2. $F(g_7^2) = F(g_4) = \begin{bmatrix} 1 & \\ & \omega^4 \end{bmatrix}$ has to be computed using the power-relation $g_7^2 = g_4$ of the pc-presentation.

3.4 Implementation

The presented algorithm has been implemented in the programming language C. The tests were run on an Intel Pentium II with 300 MHz. As we have mentioned previously, no field arithmetic is needed but only computations in the additve group $\mathbb{Z}_e$. For simplicity, we have assumed e to be known, even though one can show that this is not necessary.

The efficiency of the implementation is based on the fact, that e-monomial matrices of size N can be multiplied or inverted with only N operations in $\mathbb{Z}_e$. Since any e-monomial matrix $M \in \mathbb{C}^{N\times N}$ can be written in the form

$$M = \pi \mathrm{diag}(\omega^{a_1}, \ldots, \omega^{a_n})$$

with a permutation $\pi \in S_N$ and non-zero coefficients $\omega^{a_1}, \ldots, \omega^{a_n}$, just the $2N$ integers $\pi(1), \ldots, \pi(N)$ and $a_1, \ldots, a_N$ have to be stored for M. For the group G and any $r \in \mathbb{N}$ define $d^r(G) := \sum_{k=1}^{h} d_k^r$, where h denotes the number of conjugacy classes of G and $d_1, \ldots, d_h$ the degrees of the irreducible characters of G. One easily checks, that the running time to write out the result of the algorithm, i.e., all matrices $D_{i,k}(g_l)$, $1 \le i \le n$, $1 \le k \le h_i$ (h_i denoting the number of conjugacy classes of G_i), $1 \le l \le i$, is proportional to $\sum_{i=1}^{n} i \cdot d^1(G_i)$, which is bounded from above by $\sum_{i=1}^{n} \log(|G_i|) \cdot d^1(G_i)$.

One can show that the number of operations of the algorithm is of this magnitude $O(\sum_{i=1}^{n} \log(|G_i|) \cdot d^1(G_i))$ with a moderate constant ≤ 20. In this sense the algorithm is nearly optimal. The following table shows the running times for some small supersolvable groups to construct all the matrices $D_{i,k}(g_l)$ as above. Here $|G|$ is the order of G, h the number of conjugacy classes of G, o.l. the output length of the algorithm (i.e., $\sum_{i=1}^{n} \log(|G_i|) \cdot d^1(G_i)$), r.t. the running time in milliseconds and r.t./o.l. the quotient of the last two quantities. The groups in the first three examples are direct products of the symmetric group S_3, the group in the forth example is G_{128} from subsection 3.3 and the last example is concerned with a Sylow 2-subgroup of the symmetric group S_{16}.

G	$\|G\|$	h	o.l.	r.t. (ms)	r.t/o.l.
$(S_3)^5$	7776	243	13235	266	0.020
$(S_3)^6$	46656	729	63528	1125	0.018
$(S_3)^7$	279936	2187	296464	4250	0.014
G_{128}	128	40	280	15	0.054
$\mathrm{Syl}_2(S_{16})$	32768	230	30960	2156	0.069

Of course, the first three groups are of a very simple nature. However, the running time of the algorithm does not essentially depend on the complexity of the pc-presentation, but mainly on the number and degrees of the irreducible representations constituting the DFT. This is verified by the more complex example $\mathrm{Syl}_2(S_{16})$. Therefore, the actual running times for constructing a monomial DFT of $\mathbb{C}G$ reflect very well the theoretical result concerning the output length.

4 Applications

4.1 Related Work

The fast DFT-generation algorithm has been used as a subroutine to solve other computational problems. Thümmel [12] has designed an algorithm that computes from a pc-presentation of a finite p-group G in time $O(p \cdot h \cdot |G|)$ its h conjugacy classes as well as the character table. Omrani and Shokrollahi [8] have combined the fast DFT-generation algorithm with Galois theory to construct a full set of irreducible representations of a supersolvable group G over a finite field K, $\operatorname{char} K \not| \, |G|$, which is not assumed to contain a primitive eth root of unity.

4.2 Fast Convolution

FFT-algorithms allow a fast convolution in the group algebra $\mathbb{C}G$ along the formula: $a \cdot b = D^{-1}(D(a) \cdot D(b))$, for $a, b \in \mathbb{C}G$. (Note that the linear complexities of a DFT D and its inverse do not differ substantially, for $|L(D) - L(D^{-1})| \leq |G|$, see [2].) Let $D = \oplus_{k=1}^{h} D_k$ be a DFT of $\mathbb{C}G$ and d_k the degree of D_k. Then the convolution in $\mathbb{C}G$ can be performed with at most $2L(D) + L(D^{-1}) + 2\sum_{k=1}^{h} d_k^3$ arithmetic operations. Thus if D, and hence D^{-1}, allows a fast Fourier transform, and $d := \max_k d_k$, then convolution can be done in time $O(|G| \log |G| + d|G|)$. As $1 \leq d \leq |G|^{1/2}$, this constitutes a substantial improvement over the naive convolution algorithm, which performs this task in time quadratic in the order of G. Even in a very special case, a variant of this FFT-based fast convolution in $\mathbb{C}G$ might shed new light onto a classical problem in computational group theory. A sketch of this will be the last topic of this paper.

4.3 DFT-based Collection

As already mentioned, every element a in a pc-presented supersolvable group G can be expressed as a normal word: $a = g^{\alpha} := g_n^{\alpha_n} \cdot g_{n-1}^{\alpha_{n-1}} \cdot \ldots \cdot g_1^{\alpha_1}$. The normal form problem is to compute on input (α, β) the unique γ with $g^{\alpha} \cdot g^{\beta} = g^{\gamma}$. Classical techniques for solving this problem involve various kinds of collection processes (see, e.g., [10]) or Hall polynomials combined with interpolation techniques (see, e.g., [11]). To the best of our knowledge, there is no strategy that is always superior to all other strategies.

As an alternative to classical collection strategies we propose DFT-based normalization. To simplify our notation, we start with a pc-presented p-group G with corresponding chief series $\mathcal{T} = (G_n \supset \ldots \supset G_0)$ and complete lists $\operatorname{Irrep}(G_i, \mathcal{T}_i)$ of $\mathcal{T}_i$-adapted e-monomial irreducible representations of $\mathbb{C}G_i$. $D_{i,0}$ always denotes the trivial representation of $\mathbb{C}G_i$, $D_{i,1}$ always a non-trivial extension of $D_{i-1,0}$ satisfying $D_{i,1}(g_i) = \zeta$, where ζ is a primitive p-th root of unity. On input (α, β), the algorithm proceeds in n steps (n downto 1) to compute γ. After Step $i+1$, the numbers $\gamma_n, \ldots, \gamma_{i+1}$ are known. To get γ_i in Step i, we work in G/G_{i-1} by replacing g_j by 1, for all $j < i$. Consider the word

$$w_i := g_{i+1}^{-\gamma_{i+1}} \cdots g_n^{-\gamma_n} \cdot g_n^{\alpha_n} \cdots g_i^{\alpha_i} \cdot g_n^{\beta_n} \cdots g_i^{\beta_i}.$$

By definition, $w_i \in G_i$ and $w_i \equiv g_i^{\gamma_i} \bmod G_{i-1}$. We want to compute $D_{i,1}(w_i)$, since γ_i is determined by $D_{i,1}(w_i) = D_{i,1}(g_i^{\gamma_i}) = \zeta^{\gamma_i}$. However, since w_i is expressed in all generators $g_1, \ldots, g_n$, this cannot be done directly at level i. To this end we choose a suitable representation $F \in \mathrm{Irrep}(G_n, \mathcal{T}_n)$ whose restriction to $\mathbb{C}G_i$ contains $D_{i,1}$ as its first irreducible constituent. Then all what remains to do is to compute the first position of the diagonal matrix $F(w_i)$, which equals $D_{i,1}(w_i) = D_{i,1}(g_i^{\gamma_i}) = \zeta^{\gamma_i}$. As

$$F(w_i) = F(g_{i+1})^{-\gamma_{i+1}} \cdots F(g_n)^{-\gamma_n} \cdot F(g_n)^{\alpha_n} \cdots F(g_i)^{\alpha_i} \cdot F(g_n)^{\beta_n} \cdots F(g_i)^{\beta_i}$$

is a product of monomial matrices and we are interested in only one entry of the final result, each factor $F(g_j)$ causes only one addition in $\mathbb{Z}_e$. Altogether, we obtain the following.

Theorem 6. *Let G be a pc-presented p-group of order p^n and exponent e, with corresponding chief series $\mathcal{T}$. Then, given (suitable parts of) the $\mathcal{T}$-adapted DFT, normalization of the product of two normal words in G can be done with at most $2 \cdot p \cdot n^2$ additions in $\mathbb{Z}_e$.*

Finally, we want to remark that a similar result holds for the normalization of any formula in the generators $g_1, \ldots, g_n$ of G.

References

1. Baum, U.: Existence and efficient construction of fast Fourier transforms for supersolvable groups. Computational Complexity, **1** (1991), 235–256.
2. Baum, U., Clausen, M.: Some lower and upper complexity bounds for generalized Fourier transforms and their inverses. SIAM J. Comput., **20** (1991), 451–459.
3. Baum, U., Clausen, M.: Computing irreducible representations of supersolvable groups. Mathematics of Computation, Volume **63**, Number 207 (1994) 351–359.
4. Bürgisser, P., Clausen, M., Shokrollahi, M.A.: Algebraic Complexity Theory. Grundlehren der mathematischen Wissenschaften, Volume **315**, Springer Verlag, Berlin, 1997.
5. Clausen, M., Baum, U.: Fast Fourier Transforms. BI-Wissenschaftsverlag, Mannheim, 1993.
6. Clausen, M., Baum, U.: Ein kombinatorischer Zugang zur Darstellungstheorie überauflösbarer Gruppen. Bayreuther Mathematische Schriften, **44** (1993), 99–107.
7. Morgenstern, J.: Complexité linéaire de calcul. Thèse, Univ. de Nice, 1978.
8. Omrani, A., Shokrollahi, M.A.: Computing irreducible representations of supersolvable groups over small finite fields. Mathematics of Computation, Volume **66**, Number 218 (1997) 779–786.
9. Serre, J.P.: Linear Representations of Finite Groups. Graduate Texts in Mathematics, Springer, 1986.
10. Sims, C.C.: Computation with finitely presented groups. Cambridge University Press, 1994.
11. Sims, C.C.: Fast multiplication and growth in groups. ISSAC'98, Rostock, Germany (1998), 165–170.
12. Thümmel, A.: Computing character tables of p-groups. pp. 150–154 in Proceedings ISSAC'96, Zürich, Switzerland.