THE DISCRETE FOURIER TRANSFORM:
A CANONICAL BASIS OF EIGENFUNCTIONS

Shamgar Gurevich[1], *Ronny Hadani*[2], *and Nir Sochen*[3]

[1]Department of Mathematics, University of California Berkeley
Berkeley, CA 94720, USA
E-mail: shamgar@math.berkeley.edu

[2]Department of Mathematics, University of Chicago
Chicago, IL, 60637, USA
E-mail: hadani@math.uchicago.edu

[3]School of Mathematical Sciences, Tel Aviv University
Tel Aviv, 69978, Israel
E-mail: sochen@math.tau.ac.il

ABSTRACT

We exhibit a canonical basis Φ of eigenvectors for the discrete Fourier transform (DFT). The transition matrix Θ from the standard basis to Φ defines a novel transform which we call the *discrete oscillator transform* (DOT for short). Finally, we describe a fast algorithm for computing Θ in certain cases.

1. INTRODUCTION

The discrete Fourier transform (DFT) is probably one of the most important operators in modern science. It is omnipresent in various fields of discrete mathematics and engineering, including combinatorics, number theory, computer science and, last but probably not least, digital signal processing. Formally, the DFT is a family $\{F_N\}$ of unitary operators, where each F_N acts on the Hilbert space $\mathscr{H}_N = \mathbb{C}(\mathbb{Z}/N\mathbb{Z})$ by the formula

$$F_N[f](w) = \frac{1}{\sqrt{N}} \sum_{t \in \mathbb{Z}/N\mathbb{Z}} e^{\frac{2\pi i}{N} wt} f(t).$$

Although, so widely used, the spectral properties of the DFT remains to some extent still mysterious. For example, the calculation of the multiplicities of its eigenvalues, which was first carried out by Gauss, is quite involved and requires a multiple of number theoretic manipulations [AT].

A primary motivation for studying the eigenvectors of the DFT comes from digital signal processing. Here, a function is considered in two basic realizations: The time realization and the frequency realization. Each realization, yields information on different attributes of the function. The DFT operator acts as a dictionary between these two realizations

$$\textbf{Time} \overset{F_N}{\rightleftarrows} \textbf{Frequency.}$$

From this point of view, it is natural to look for a diagonalization basis, namely, a basis of eigenvectors (eigen modes) for F_N. In this regard, the main conceptual difficulty comes from the fact that the diagonalization problem is ill-defined, since, F_N is an operator of order 4, i.e., $F_N^4 = Id$, which means that it has at most four eigenvalues $\pm 1, \pm i$, therefore each appears with large multiplicity (We assume $N \gg 4$).

An interesting approach to the resolution of this difficulty, motivated from results in continuous Fourier analysis, was developed by Grünbaum in [G]. In that approach, a tridiagonal operator S_N which commutes with F_N and admits a simple spectrum is introduced. This enable him to give a basis of eigenfunctions for the DFT. Specifically, S_N appears as a certain discrete analogue of the differential operator $D = \partial_t^2 - t^2$ which commutes with the continuous Fourier transform.

1.1 Main results of this paper

In this paper we describe a representation theoretic approach to the diagonalization problem of the DFT in the case when $N = p$ is an odd prime number. Our approach, puts to the forefront the Weil representation [W] of the finite symplectic group $Sp = SL_2(\mathbb{F}_p)$ as the fundamental object underlying harmonic analysis in the finite setting. Specifically, we exhibit a canonical basis Φ_p of eigenvectors for the DFT. We also describe the transition matrix Θ_p from the standard basis to Φ_p, which we call the *discrete oscillator transform* (DOT for short). In addition, in the case $p \equiv 1(\text{mod}4)$, we describe a fast algorithm for computing Θ_p (FOT for short).

It is our general feeling that the Weil representation yields a transparent explanation to many classical results in finite harmonic analysis. To justify this claim, we describe an alternative method for calculating the multiplicities of the eigenvalues for the DFT, a method we believe is more suggestive then the classical calculations.

The rest of the introduction is devoted to a more detailed account of the main ideas and results of this paper.

1.2 Symmetries of the DFT

Let us fix an odd prime number p and for the rest of the introduction suppress the subscript p from all notations. Generally, when a (diagonalizable) linear operator A has eigenvalues admitting large multiplicities, it suggest that there exists a group $G = G_A \subset GL(\mathscr{H})$ of "hidden" symmetries consisting of operators which commute with A. Alas, usually the problem of computing the group G is formidable and, in fact, equivalent to the problem of diagonalizing A. If the operator

A arise "naturally", there is a chance that the group G can be effectively described. In preferred situations, G is commutative and large enough so that all degeneracies are resolved and the spaces of common eigenvectors with respect to G are one-dimensional. The basis of common eigenvectors with respect to G establishes a distinguish choice of eigenvectors for A. Philosophically, we can say that it is more correct to consider from start the group G instead of the single operator A.

Interestingly, the DFT operator $F = F_p$ admits a natural group of symmetries G_F, which, in addition, can be effectively described using the Weil representation. For the sake of the introduction, it is enough to know that the Weil representation in this setting is a unitary representation $\rho : Sp \rightarrow U(\mathscr{H})$ and the key observation is that F is proportional to a single operator $\rho(\mathrm{w})$. The group G_F is the image under ρ of the centralizer subgroup T_{w} of w in Sp.

1.3 The algebraic torus associated to the DFT

The subgroup T_{w} can be computed explicitly and is of a very "nice" type, it consists of rational points of a maximal algebraic torus in Sp which in plain language means that it is maximal commutative subgroup in Sp, consisting of elements which are diagonalizable over some field extension. Restricting the Weil representation to the subgroup T_{w} yields a collection $G_F = \{\rho(g) : g \in T_{\mathrm{w}}\}$ of commuting operators, each acts unitarily on the Hilbert space $\mathscr{H}$ and commutes with F. This, in turn, yields a decomposition, stable under Fourier transform, into character spaces

$$\mathscr{H} = \bigoplus \mathscr{H}_\chi, \tag{1}$$

where χ runs in the set of (complex valued) characters of T_{w}, namely, if $v \in \mathscr{H}_\chi$ then $\rho(g)v = \chi(g)v$. The main technical statement of this paper, Theorem 3, roughly says that $\dim \mathscr{H}_\chi = 1$ for every χ which appears in (1). Choosing a unit representative $\phi_\chi \in \mathscr{H}_\chi$ for every χ, gives the canonical basis $\Phi = \{\phi_\chi\}$ of eigenvectors for F. The oscillator transform Θ sends a function $f \in \mathscr{H}$ to the coefficients in the unique expansion

$$f = \sum a_\chi \phi_\chi.$$

The fine behavior of F and Θ is governed by the (split type) structure of T_{w}, which changes depending on the value of the prime p modulo 4. This phenomena has several consequences. In particular, it gives a transparent explanation to the precise way the multiplicities of the eigenvalues of F depend on the prime p. Another, algorithmic, consequence is related to the existence of a fast algorithm for computing Θ.

1.4 Properties of eigenvectors

The character vectors ϕ_χ satisfy many interesting properties and are objects of study in their own right. A comprehensive treatment of this aspect of the theory appears in [GHS].

1.5 Generalizations

1.5.1 Field extensions

All the results in this paper were stated for the basic finite field $\mathbb{F}_p$, for the reason of making the terminology more accessible. In fact, all the results can be stated and proved for any field extension of the form $\mathbb{F}_q$, $q = p^n$, one should only replace p by q in all appropriate places.

1.6 Structure of the paper

We begin by discussing the finite Heisenberg group and the Heisenberg representation. Next we introduce the Weil representation of the finite symplectic group, first it is described in abstract terms and then more explicitly invoking the idea of invariant presentation of an operator. We proceed to discuss the theory of tori in the one-dimensional Weil representation, we explain how to associate to a maximal torus $T \subset SL_2$, a transform Θ_T called the oscillator transform. We describe a fast algorithm for computing Θ_T in the case T is a split torus. The theory is then applied to the specific torus associated with the DFT operator. We finish with a treatment of the multiplicity problem for the DFT, from the representation theoretic perspective.

1.7 Acknowledgements

It is a pleasure to thank our teacher J. Bernstein for his interest and guidance. We acknowledge R. Howe for sharing with us some unpublished results of his. We appreciate several discussions we had with A. Grünbaum, W. Kahan, B. Parlett, B. Porat, A. Sahai on the DFT. We thank A. Weinstein, J. Wolf and M. Zworski for interesting conversations related to the metaplectic representation. Finally, we thank K. Ribet, B. Poonen, M. Haiman and M. Gu for the opportunities to present this work in the number theory, representation theory and scientific computing seminars at Berkeley during February 2008.

2. THE OSCILLATOR TRANSFORM

2.1 The Heisenberg group

Let (V, ω) be a two-dimensional symplectic vector space over the finite field $\mathbb{F}_p$. The reader should think of V as $\mathbb{F}_p \times \mathbb{F}_p$ with the standard form $\omega((\tau, w), (\tau', w')) = \tau w' - w\tau'$. Considering V as an abelian group, it admits a non-trivial central extension called the *Heisenberg* group. Concretely, the group H can be presented as the set $H = V \times \mathbb{F}_p$ with the multiplication given by

$$(v,z) \cdot (v',z') = (v+v', z+z' + \tfrac{1}{2}\omega(v,v')).$$

The center of H is $Z = Z(H) = \{(0,z) : z \in \mathbb{F}_p\}$. The symplectic group $Sp = Sp(V, \omega)$, which in this case is isomorphic to $SL_2(\mathbb{F}_p)$, acts by automorphism of H through its action on the V-coordinate.

2.2 The Heisenberg representation

One of the most important attributes of the group H is that it admits, principally, a unique irreducible representation. The precise statement goes as follows. Let $\psi : Z \rightarrow \mathbb{C}^\times$ be a character of the center. For example we can take $\psi(z) = e^{\frac{2\pi i}{p} z}$.

Theorem 1.(Stone-von Neumann)

There exists a unique (up to isomorphism) irreducible unitary representation $(\pi, H, \mathscr{H})$ with the center acting by ψ, i.e., $\pi_{|Z} = \psi \cdot Id_{\mathscr{H}}$.

The representation π which appears in the above theorem will be called the *Heisenberg representation*.

2.2.1 Standard realization of the Heisenberg representation.

The Heisenberg representation $(\pi, H, \mathscr{H})$ can be realized as follows: $\mathscr{H}$ is the Hilbert space $\mathbb{C}(\mathbb{F}_p)$ of complex valued functions on the finite line, with the standard Hermitian product. The action π is given by $\pi(\tau,0) \rhd f(t) = f(t+\tau)$, $\pi(0,w) \rhd f(t) = \psi(wt) f(t)$ and $\pi(z) \rhd f(t) = \psi(z) f(t)$. We call this explicit realization the *standard realization*.

2.3 The Weil representation

A direct consequence of Theorem 1 is the existence of a projective representation $\widetilde{\rho} : Sp \to PU(\mathscr{H})$. The construction of $\widetilde{\rho}$ out of the Heisenberg representation π is due to Weil [W] and it goes as follows. Considering the Heisenberg representation π and an element $g \in Sp$, one can define a new representation π^g acting on the same Hilbert space via $\pi^g(h) = \pi(g(h))$. Clearly both π and π^g have the same central character ψ hence by Theorem 1 they are isomorphic. Since the space $\mathsf{Hom}_H(\pi,\pi^g)$ is one-dimensional, choosing for every $g \in Sp$ a non-zero representative $\widetilde{\rho}(g) \in \mathsf{Hom}_H(\pi,\pi^g)$ gives the required projective representation. In more concrete terms, the projective representation $\widetilde{\rho}$ is characterized by the formula

$$\rho(g)\,\pi(h)\,\rho\left(g^{-1}\right) = \pi(g(h)), \qquad (2)$$

for every $g \in Sp$ and $h \in H$. It is a peculiar phenomenon of the finite field setting that the projective representation $\widetilde{\rho}$ can be linearized into an honest representation.

Theorem 2 *There exists a unique*[1] *linear representation*

$$\rho : Sp \longrightarrow GL(\mathscr{H}),$$

which satisfies equation (2).

2.3.1 Invariant presentation of the Weil representation

Let us denote by $\mathbb{C}(H,\psi)$ the space of (complex valued) functions on H which are ψ-equivariant with respect to the action of the center, namely, a function $f \in \mathbb{C}(H,\psi)$ satisfies $f(zh) = \psi(z) f(h)$ for every $z \in Z$, $h \in H$. Given an operator $A \in \mathsf{End}(\mathscr{H})$, it can be written in a unique way as $A = \pi(K_A)$, where $K_A \in \mathbb{C}\left(H,\psi^{-1}\right)$ and π denotes the extended action $\pi(K_A) = \sum_{h\in H} K_A(h)\,\pi(h)$. The function K_A is called the *kernel of A* and it is given by the *matrix coefficient*

$$K_A(h) = \frac{1}{\dim \mathscr{H}} Tr\left(A\pi\left(h^{-1}\right)\right). \qquad (3)$$

In the context of the Heisenberg representation, formula (3) is usually referred to as the *Weyl transform*. Using the Weyl transform one is able to give an explicit description of the Weil representation. The idea [GH1] is to write each operator $\rho(g)$, $g \in Sp$ in terms of its kernel function $K_g = K_{\rho(g)} \in \mathbb{C}\left(H,\psi^{-1}\right)$. The following formula is taken from [GH1]

$$K_g(v,z) = \frac{\sigma(-1)}{\dim \mathscr{H}} \sigma(\det(\kappa(g)+I))\,\psi\left(\tfrac{1}{4}\omega(\kappa(g)v,v)+z\right) \qquad (4)$$

for every $g \in Sp$ such that $g - I$ is invertible, where σ denotes the unique quadratic character (Legendre character) of the multiplicative group $\mathbb{F}_p^\times$ and κ is the Cayley transform $\kappa(g) = \frac{g+I}{g-I}$, $g \in Sp$.

2.4 The theory of tori

A maximal (algebraic) torus in Sp is a maximal commutative subgroup which becomes diagonalizable over some field extension. There exists two conjugacy classes of maximal (algebraic) tori in Sp. The first class consists of those tori which are diagonalizable already over $\mathbb{F}_p$ or equivalently those are the tori that are conjugated to the standard diagonal torus

$$A = \left\{ \begin{pmatrix} a & 0 \\ 0 & a^{-1} \end{pmatrix} : a \in \mathbb{F}_p \right\}.$$

A torus in this class is called a *split* torus. The second class consists of those tori which become diagonalizable over a quadratic extension $\mathbb{F}_{p^2}$ or equivalently those are tori which are not conjugated to A. A torus in this class is called a *non-split* torus (sometimes it is called inert torus).

Example 1(Example of a non-split torus)

It might be suggestive to explain further the notion of non-split torus by exploring, first, the analogue notion in the more familiar setting of the field $\mathbb{R}$. Here, the standard example of a maximal non-split torus is the circle group $SO(2) \subset SL_2(\mathbb{R})$. Indeed, it is a maximal commutative subgroup which becomes diagonalizable when considered over the extension field $\mathbb{C}$ of complex numbers. The above analogy suggests a way to construct an example of a maximal non-split torus in the finite field setting as well.

Let us identify the symplectic plane $V = \mathbb{F}_p \times \mathbb{F}_p$ with the quadratic extension $\mathbb{F}_{p^2}$. Under this identification, $\mathbb{F}_{p^2}$ acts on V and for every $g \in \mathbb{F}_{p^2}$ we have $\omega(gu,gv) = \det(g)\,\omega(u,v)$, which implies that the group

$$T_{ns} = \left\{ g \in \mathbb{F}_{p^2}^\times : \det(g) = 1 \right\}$$

naturally lies in Sp. The group T_{ns} is an example of a non-split torus which the reader might think of as the "finite circle".

2.4.1 Decompositions with respect to a maximal torus

Restricting the Weil representation to a maximal torus $T \subset Sp$ yields a decomposition

$$\mathscr{H} = \bigoplus_{\chi} \mathscr{H}_\chi, \qquad (5)$$

where χ runs in the set $T^\vee$ of complex valued characters of the torus T. More concretely, choosing a generator[2] $t \in T$, the decomposition (5) naturally corresponds to the eigenspaces of the linear operator $\rho(t)$. The decomposition (5) depends on the split type of T. Let σ_T denote the unique quadratic character of T.

Theorem 3 ([GH3])*If T is a split torus*

$$\dim \mathscr{H}_\chi = \begin{cases} 1 & \chi \neq \sigma_T, \\ 2 & \chi = \sigma_T. \end{cases}$$

If T is a non-split torus

$$\dim \mathscr{H}_\chi = \begin{cases} 1 & \chi \neq \sigma_T, \\ 0 & \chi = \sigma_T. \end{cases}$$

[1] Unique, except in the case the finite field is $\mathbb{F}_3$. For the canonical choice in the latter case see [GH1].

[2] A maximal torus T in $SL_2(\mathbb{F}_p)$ is a cyclic group, thus there exists a generator.

2.5 The discrete oscillator transform associated to a maximal torus

Let us fix a maximal torus T. Every vector $v\in\mathscr{H}$ can be written uniquely as a direct sum $v=\sum v_{\chi}$ with $v_{\chi}\in\mathscr{H}_{\chi}$ and χ runs in $I=\mathsf{Spec}_{T}\left(\mathscr{H}\right)$ - the spectral support of $\mathscr{H}$ with respect to T consisting of all characters $\chi\in T^{\vee}$ such that $\dim\mathscr{H}_{\chi}\neq 0$. Let us choose, in addition, a collection of unit vectors $\phi_{\chi}\in\mathscr{H}_{\chi}$, $\chi\in I$ and let $\phi=\sum\phi_{\chi}$. We define the transform $\Theta_{T}=\Theta_{T,\phi}:\mathscr{H}\rightarrow\mathbb{C}\left(I\right)$ by $\Theta_{T}\left[v\right]\left(\chi\right)=\left\langle v,\phi_{\chi}\right\rangle$. We will call the transform Θ_{T} the *discrete oscillator transform* (DOT for short) with respect to the torus T and the test vector ϕ.

Remark 1 We note that in the case T is a non-split torus, Θ_{T} maps $\mathscr{H}$ isomorphically to $\mathbb{C}\left(I\right)$. In the case T is a split torus, Θ_{T} has a kernel consisting of $v\in\mathscr{H}$ such that $\left\langle v,\phi_{\sigma_{T}}\right\rangle=0$.

2.5.1 The oscillator transform (integral form)

Let $\mathscr{M}_{T}:\mathbb{C}(T)\rightarrow\mathbb{C}(T^{\vee})$ denote the Mellin transform

$$\mathscr{M}_{T}\left[f\right]\left(\chi\right)=\frac{1}{\#T}\sum_{g\in T}\overline{\chi}\left(g\right)f\left(g\right),$$

for $f\in\mathbb{C}\left(T\right)$. Let us denote by $m_{T}:\mathscr{H}\rightarrow\mathbb{C}\left(T\right)$ the matrix coefficient $m_{T}\left[v\right]\left(g\right)=\left\langle v,\rho\left(g^{-1}\right)\phi\right\rangle$ for $v\in\mathscr{H}$.

Lemma 1 ([GH3]) *We have*

$$\Theta_{T}=\mathscr{M}_{T}\circ m_{T}.$$

2.5.2 Fast oscillator transforms

In practice, it is desirable to have a "fast" algorithm for computing the oscillator transform (FOT for short). We work in the following setting. The vector v is considered in the standard realization $\mathscr{H}=\mathbb{C}\left(\mathbb{F}_{p}\right)$ (see 2.2.1), in this context the oscillator transform gives the transition matrix between the basis of delta functions and the basis $\left\{\phi_{\chi}\right\}$ of character vectors. We will show that when T is a split torus and for an appropriate choice of ϕ, the oscillator transform can be computed in $O\left(p\log\left(p\right)\right)$ arithmetic operations. Principally, what we will show is that the computation reduces to an application of DFT followed by an application of the standard Mellin transform, both transforms admit a fast algorithm [CT].

Assume T is a split torus. Since all split tori are conjugated to one another, there exists, in particular, an element $s\in Sp$ conjugating T with the standard diagonal torus A. In more details, we have a homomorphism of groups $Ad_{s}:T\longrightarrow A$ sending $g\in T$ to $Ad_{s}\left(g\right)=sgs^{-1}\in A$. Dually, we have a homomorphism $Ad_{s}^{\vee}:A^{\vee}\longrightarrow T^{\vee}$ between the corresponding groups of characters.

The main idea is to relate the oscillator transform with respect to T with the oscillator transform with respect to A. The relation is specified in the following simple lemma.

Lemma 2 ([GH3]) *We have*

$$\left(Ad_{s}^{\vee}\right)^{*}\circ\Theta_{T,\phi}=\Theta_{A,\rho(s)\phi}\circ\rho\left(s\right). \qquad (6)$$

Remark 2 Roughly speaking, (6) means that (up to a "reparametrization" of $T^{\vee}$ by $A^{\vee}$ using $Ad_{s}^{\vee}$) the oscillator transform of a vector $v\in\mathscr{H}$ with respect to the torus T is the same as the oscillator transform of the vector $\rho\left(s\right)v$ with respect to the diagonal torus A.

In order to finish the construction we need to specify two basic facts about the Weil representation in the standard realization.

- The standard torus A acts by (normalized) scalings, the precise formula of $\rho\left(g\right)$ for $g=\begin{pmatrix}a & 0\\ 0 & a^{-1}\end{pmatrix}\in A$ is
$$\left(\rho\left(g\right)\triangleright f\right)\left(x\right)=\sigma\left(a\right)f\left(ax\right).$$
- Every operator $\rho\left(g\right)$, $g\in Sp$, can be written in the form $\rho\left(g\right)=M_{g_{1}}\circ F\circ M_{g_{2}}\circ S_{a}$ where $M_{g_{1}},M_{g_{2}}$ are the operators of multiplication by some functions $g_{1},g_{2}\in\mathbb{C}\left(\mathbb{F}_{p}\right)$, S_{a} is the operator of scaling by $a\in\mathbb{F}_{p}^{\times}$ and F is the DFT.
$$F\left[f\right]\left(y\right)=\frac{1}{\sqrt{p}}\sum_{x\in\mathbb{F}_{p}}\psi\left(yx\right)f\left(x\right).$$

Given a function $f\in\mathbb{C}\left(\mathbb{F}_{p}\right)$, applying formula (6) with $\phi=\rho\left(s\right)^{-1}\delta_{1}$ yields

$$\Theta_{T,\phi}\left[f\right]\left(Ad_{s}^{\vee}\left(\chi\right)\right)=\frac{1}{p-1}\sum_{a\in\mathbb{F}_{p}^{\times}}\sigma\left(a\right)\overline{\chi}\left(a\right)\left(\rho\left(s\right)\triangleright f\right)\left(a\right), \qquad (7)$$

for every $\chi\in A^{\vee}$. In conclusion, formula (7) implies that $\Theta_{T,\phi}\left[f\right]$ can be computed by, first, applying the operator $\rho\left(s\right)$ to f and then applying Mellin transform to the result.

Problem 1

Does there exists a fast algorithm for computing the oscillator transform associated to a ***non-split*** *torus?*

2.6 Diagonalization of the discrete Fourier transform

In this subsection we apply the previous development in order to exhibit a canonical basis of eigenvectors for the DFT. We will show that the DFT can be naturally identified (up to a normalization scalar) with an operator $\rho\left(\mathrm{w}\right)$ in the Weil representation, where w is an element in a maximal torus $T_{\mathrm{w}}\subset Sp$. We take w$\in Sp=SL_{2}\left(\mathbb{F}_{p}\right)$ to be the Weyl element

$$\mathrm{w}=\begin{pmatrix}0 & 1\\ -1 & 0\end{pmatrix}.$$

Lemma 3([H, GH3]) *We have*

$$F=C\cdot\rho\left(\mathrm{w}\right),$$

where $C=i^{\frac{p-1}{2}}$.

Lemma 3 implies that the diagonalization problems of the operators F and $\rho\left(\mathrm{w}\right)$ are equivalent. The second problem can be approached using representation theory, which is what we are going to do next.

Let us denote by T_{w} the centralizer of w in Sp, namely T_{w} consists of all elements $g\in Sp$ such that gw=wg, in particular we have that w$\in T_{\mathrm{w}}$.

Proposition 1 ([GH3]) *The group* T_{w} *is a maximal torus. Moreover the split type of* T_{w} *depends on the prime* p *in the following way:* T_{w} *is a split torus when* $p\equiv 1\left(mod4\right)$ *and is a non-split torus when* $p\equiv 3\left(mod4\right)$.

Proposition 1 has several consequences. First consequence is that choosing a unit character vector $\phi_{\chi}\in\mathscr{H}_{\chi}$ for every $\chi \in \mathsf{Spec}_{T_{\mathrm{w}}}(\mathscr{H})$ gives a canonical (up to normalizing unitary constants) choice of eigenvectors for the DFT [3]. Second, more mysterious consequence is that although the formula of the DFT is uniform in p, its qualitative behavior changes dramatically between the cases when $p \equiv 1\,(\mathrm{mod}4)$ and $p \equiv 3\,(\mathrm{mod}4)$. This is manifested in the structure of the group of symmetries: In the first case, the group of symmetries is a split torus consisting of $p-1$ elements and in the second case it is a non-split torus consisting of $p+1$ elements. It also seems that the structure of the symmetry group is important from the algorithmic perspective, in the case $p \equiv 1\,(\mathrm{mod}4)$ we built a fast algorithm for computing Θ, while in the case $p \equiv 3\,(\mathrm{mod}4)$ the existence of such an algorithm remains open (see Problem 1).

2.6.1 *Multiplicities of eigenvalues of the DFT*

Considering the group T_{w} we can give a transparent computation of the eigenvalues multiplicities for the operator $\rho(\mathrm{w})$. First we note that, since w is an element of order 4, the eigenvalues of $\rho(\mathrm{w})$ lies in the set $\{\pm1, \pm i\}$. For $\lambda \in \{\pm1, \pm i\}$, let m_λ denote the multiplicity of the eigenvalue λ. We observe that

$$m_\lambda = \bigoplus_{\chi \in I_\lambda} \dim \mathscr{H}_\chi,$$

where I_λ consists of all characters $\chi \in \mathsf{Spec}_{T_{\mathrm{w}}}(\mathscr{H})$ such that $\chi(\mathrm{w}) = \lambda$. The result now follows easily from Theorem 3, applied to the torus T_{w}. We treat separately the split an non-split cases.

- Assume T_{w} is a split torus, which happens when $p = 1\,(\mathrm{mod}4)$, namely, $p = 4l+1$, $l \in \mathbb{N}$. Since $\dim \mathscr{H}_\chi = 1$ for $\chi \neq \sigma_{T_{\mathrm{w}}}$ it follows that $m_{\pm i} = \frac{p-1}{4} = l$. We are left to determine the values of $m_{\pm 1}$, which depend on whether $\sigma_{T_{\mathrm{w}}}(\mathrm{w}) = \mathrm{w}^{\frac{p-1}{2}}$ is 1 or -1.
  Since w is an element of order 4 in T_{w} we get that
  $$\sigma_{T_{\mathrm{w}}}(\mathrm{w}) = \begin{cases} 1, & p \equiv 1\,(\mathrm{mod}8), \\ -1, & p \equiv 5\,(\mathrm{mod}8), \end{cases}$$
  which implies that when $p \equiv 1\,(\mathrm{mod}8)$ then $m_1 = l+1$ and $m_{-1} = l$ and when $p \equiv 5\,(\mathrm{mod}8)$ then $m_1 = l$ and $m_{-1} = l+1$.
- Assume T_{w} is a non-split torus, which happens when $p \equiv 3\,(\mathrm{mod}4)$, namely, $p = 4l+3$, $l \in \mathbb{N}$. Since $\dim \mathscr{H}_\chi = 1$ for $\chi \neq \sigma_{T_{\mathrm{w}}}$ it follows that $m_{\pm i} = \frac{p+1}{4} = l+1$. The values of $m_{\pm 1}$ depend on whether $\sigma_{T_{\mathrm{w}}}(\mathrm{w}) = \mathrm{w}^{\frac{p+1}{2}}$ is 1 or -1. Since w is an element of order 4 in T_{w} we get that
  $$\sigma_{T_{\mathrm{w}}}(\mathrm{w}) = \begin{cases} 1, & p \equiv 7\,(\mathrm{mod}8), \\ -1, & p \equiv 3\,(\mathrm{mod}8), \end{cases}$$
  which implies that when $p \equiv 7\,(\mathrm{mod}8)$ then $m_1 = l$ and $m_{-1} = l+1$ and when $p \equiv 3\,(\mathrm{mod}8)$ then $m_1 = l+1$ and $m_{-1} = l$.

Summarizing, the multiplicities of the operator $\rho(\mathrm{w})$ are

	m_1	m_{-1}	m_i	m_{-i}
$p = 8k+1$	$2k+1$	$2k$	$2k$	$2k$
$p = 8k+3$	$2k$	$2k+1$	$2k+1$	$2k+1$
$p = 8k+5$	$2k+1$	$2k+2$	$2k+1$	$2k+1$
$p = 8k+7$	$2k+2$	$2k+1$	$2k+2$	$2k+2$

(8)

Considering now the DFT operator F. If we denote by n_μ, $\mu \in \{\pm1, \pm i\}$ the multiplicity of the eigenvalue μ of F then the values of n_μ can be deduced from table 8 by invoking the relation $n_\mu = m_\lambda$ where $\lambda = i^{\frac{p-1}{2}} \cdot \mu$ (see Lemma 3). Summarizing, the multiplicities of the DFT are

	n_1	n_{-1}	n_i	n_{-i}
$p = 4l+1$	$l+1$	l	l	l
$p = 4l+3$	$l+1$	$l+1$	$l+1$	l

For a comprehensive treatment of the multiplicity problem from a more classical point of view see [AT]. Other applications of Lemma 3 appear in [GH2].

REFERENCES

[AT] Auslander L. and Tolimieri R. Is computing with the finite Fourier transform pure or applied mathematics? *Bull. Amer. Math. Soc. (N.S.)* **1** (1979), no. 6, 847-897.

[CT] Cooley J.W. and Tukey J.W., An algorithm for the machine calculation of complex Fourier series. Math. Comput. 19, 297–301 (1965)

[G] Grünbaum F. A. The eigenvectors of the discrete Fourier transform. *J. Math. Anal. Appl. 88, 355-363.* (1982).

[GH1] Gurevich S. and Hadani R., The Geometric Weil representation. *Selecta Mathematica* (accepted: 2006).

[GH2] Gurevich S. and Hadani R., Quadratic reciprocity and sign of Gauss sum via the finite Weil representation. Submitted (Feb. 2008).

[GH3] Gurevich S. and Hadani R., On the diagonalization of the discrete Fourier transform. In preparation .

[GHS] Gurevich S., Hadani R. and Sochen N., The finite harmonic oscillator and its associated sequences. *Proceedings of the National Academy of Sciences of the United States of America* (accepted: Feb. 2008).

[H] Howe R., Private communication (Sep. 2006).

[W] Weil A., Sur certains groupes d'operateurs unitaires. *Acta Math. 111* (1964) 143-211.

[3] In the case T_w is a split torus there is a slight ambiguity in the choice of a character vector with respect to σ_{T_w}. This ambiguity can be resolved by further investigation which we will not discuss here.