Radboud University

Faculty of Science

Heisenberg groups

On applications using representation theory

Bachelor Thesis Mathematics

Author:
Krijn Reijnders (s4268202)

Supervisor:
Prof. Dr. Ben Moonen

June 24, 2015

Acknowledgements

First of all, I would like to thank my supervisor prof. dr. Ben Moonen. He has encouraged and inspired this mathematical quest through many fruitful discussions for over more than a year. Furthermore, his strictness has proven to be a great remedy to most of my sloppiness.

Secondly, I would like to thank dr. Shamgar Gurevich, from the University of Wisconsin in Madison. His work and help have been of great influence for this thesis and our discussions on the philosophy behind mathematics have been enlightening.

Finally, I would like to thank some of my fellow students, especially Wessel Martens and Tom Salet, for listening to countlessly many of my monologues on the topics of this thesis without boredom and for asking inspiring questions.

Contents

Introduction

This thesis is divided into four chapters. The first chapter consists of preliminaries. It introduces the basic notions of representation theory of finite groups and symplectic vector spaces. In fact, a course on representation theory of finite groups was the original source of inspiration for this thesis and is one of the most beautiful parts of mathematics that I have discovered so far. It was developed by Ferdinand Frobenius (1849-1917) in the beginning of the 20th century from a question that he was asked by Dedekind. It was further developed by William Burnside (1852-1927), Issai Schur (1875-1941) and Richard Brauer (1901-1977). Nowadays, representation theory is an area of mathematics with connections to harmonical analysis, geometry and number theory.

The second chapter is on Heisenberg groups and the Stone-von Neumann theorem. I was directed to this topic by Ben Moonen, and after some research on the topic we decided that this should become the core of the thesis. Heisenberg groups are named after the German physicist Werner Heisenberg (1901-1976), who used them to describe one-dimensional quantum mechanical states. Often, mathematicians and physicist speak of *the* Heisenberg group, with which they mean a certain group of upper-triangular matrices. This can however be generalized to more groups, as we will see in the second section of this chapter. We do however restrict ourselves to *finite* Heisenberg groups. Although Heisenberg groups are still often used in physics, I have chosen to focus on the mathematical aspect of the Heisenberg groups. The Stone-von Neumann theorem concerns both Heisenberg groups and representation theory and is therefore essential for chapters three and four. It was conjectured by Marshall Stone (1903-1989) in 1930 and proven by John von Neumann (1903-1957) in a year later. It has many different formulations, I have of course chosen for the one using representation theory.

Chapters three and four consist of applications of Heisenberg groups and representation theory. Shamgar Gurevich' expertise in this area of mathematics has been of influence. Chapter three describes the Fast Fourier Transform in an abstract mathematical sense. Nowadays, the Fast Fourier Transform is used in a lot of different sections of mathematics, science and engineering, for example in signal and image procession. This topic has been inspired by a presentation of Shamgar Gurevich at the IMA 2014 PI Summer Graduate Program in Chicago.

Chapter four describes my search for Mutually Unbiased Bases (MUBs) and their connection to Gauss sums using Heisenberg groups. MUBs are a relatively

new concept in mathematics with many uses in quantum information theory. Gauss sums however are almost two centuries old and were studied extensively by Carl Friedrich Gauss (1777-1855). They are very well studied in number theory. In chapter four we also discuss how they can be studied using representation theory. I was guided to this topic by Shamgar Gurevich, and I visited him in Madison (as part of the Radboud Honours Programme) to complete this chapter. Not only was this trip useful for the mathematical progression of chapter four, it was also a great experience to visit this foreign university and participate in some of the lectures given at the University of Wisconsin.

Chapter 1

Preliminaries

This chapter of preliminaries is split up in two parts. The first part introduces the basics of representation theory of finite groups. The second part introduces the basics of symplectic vector spaces. A reader who is familiar with these concepts may skip these sections, as they introduce only standard information.

1.1 Basic Notions from Representation Theory of Finite Groups

In this section we introduce the very basics of representation theory of finite groups. This is done as in [7], but we leave out most of what is not important to this thesis. We introduce (irreducible) representations of finite groups, characters and induced representations. This is supposed to be a short introduction to representation theory and we must therefore sometimes choose between simplicity, brevity and clarity on one side, and thoroughness, rigour and detail on the other. For a more elaborate introduction to representation theory of finite groups, we refer to [7].

For the rest of this section, G denotes a finite group and all vector spaces are vector spaces over $\mathbb{C}$.

1.1.1 Irreducible Representations

We start with the basic definition of representations and some examples. We then define what it means for representations to be isomorphic to each other and what it means to be irreducible. Furthermore, a fine result by Schur is explained in the end of the subsection, which is often used in representation theory.

Definition 1.1. A *representation* (π, V) of G is a homomorphism π of G to the space of linear transformations on V, or

$$\pi : G \to \mathrm{GL}(V)$$

$$\pi(gh) = \pi(g)\pi(h) \qquad \forall g, h \in G$$

Remark. Although a representation (π, V) consists of a homomorphism π *and* a vector space V, we will usually refer to a representation as just π or just V.

Example 1.2.

1. For every group G, we have the *trivial representation* $(\rho_{\mathrm{triv}}, \mathbb{C})$ which is defined by $g \mapsto 1$ for all $g \in G$.

2. An easy example of a representation of $C_3 = \{e, a, a^2\}$ is $(\rho, \mathbb{C}^2)$ with

$$\rho(e) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},\ \rho(a) = \begin{pmatrix} 1 & 0 \\ 0 & \zeta_3 \end{pmatrix},\ \rho(a^2) = \begin{pmatrix} 1 & 0 \\ 0 & {\zeta_3}^2 \end{pmatrix}$$

   where ζ_3 is a third root of unity.

3. Another example is the representation $(\pi, \mathbb{C}^3)$ of S_3 with

$$\pi(1,2) = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \pi(1,2,3) = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}$$

   and so on. We can extend this idea to a representation on S_n for all n. We call this representation of S_n the *standard representation*.

4. A simpler example of a representation of S_3 would be what is called the *sign representation*, $(\rho_{\mathrm{sgn}}, \mathbb{C})$, defined by

$$\rho_{\mathrm{sgn}} : S_3 \to \mathrm{GL}_1(\mathbb{C}) = \mathbb{C}^*$$

$$\rho_{\mathrm{sgn}}(\sigma) = \operatorname{sgn}\sigma \qquad \forall \sigma \in S_3$$

Definition 1.3. The *dimension of a representation* (π, V) is defined as the dimension of the vector space V. For example, the dimension of the representations in Example 1.2 are 1, 2, 3 and 1, respectively.

Although some representations may look very different, they may essentially be the same in the sense that they are *isomorphic representations*.

Definition 1.4. We call two representation (π, V) and (τ, W) of G *isomorphic representations* if there exists an isomorphism $f : V \xrightarrow{\sim} W$ that satisfies

$$f \circ \pi(g) = \tau(g) \circ f$$

for all $g \in G$. The function f is called an *intertwiner* between π and τ. We use the notation $\pi \cong \tau$ to indicate that two representations are isomorphic.

An important notion in representation theory is the *irreducibility* of representations. They can be considered the building blocks of all representations of a finite group. We define them as follows:

Definition 1.5. A representation (π, V) is called an *irreducible representation* if the only subspaces that are stable under the action of G are the trivial subspaces, $\{0\}$ and V. If π is not irreducible, we call it a *reducible representation*.

Example 1.6. We shall look at the irreducibility of the examples we saw in Example 1.2.

1. Every 1-dimensional vector space cannot have non-trivial subspaces, so every 1-dimensional representation is irreducible.
2. The subspace $\mathrm{Span}(1,0)$ is stable under $\rho(e), \rho(a)$ and $\rho(a^2)$. Therefore, ρ is a reducible representation.
3. The subspace $\mathrm{Span}(1,1,1)$ is stable under the image of π, because every permutation of (a,a,a) must be (a,a,a). Therefore, π is a reducible representation.
4. ρ_{sgn} is 1-dimensional, therefore ρ_{sgn} is irreducible.

We can also take the *direct sum* of two representations of the same group, and form a new representation. The direct sum of two representations (π, V_1) and (τ, V_2) is written as $(\pi + \tau, V_1 \oplus V_2)$ and defined, in matrix form, by

$$(\pi + \tau)(g) = \begin{pmatrix} \pi(g) & 0 \\ 0 & \tau(g) \end{pmatrix}.$$

For example, the representation ρ in Example 1.2.2 is the direct sum of the trivial representation and the irreducible representation $a \mapsto \zeta_3$, where ζ_3 is a third root of unity. It is easy to see that an irreducible representation cannot be the direct sum of two representations, as this would lead to stable subspaces. For a finite group, we have the following lemma:

Lemma 1.7. *Let G be a finite group and let (π, V) be a representation of G. If W is a non-trivial subspace of V stable under π, then there exists a non-trivial subspace W^0 that is also stable under π and*

$$W \bigoplus W^0 = V.$$

Proof. Let $V = W \bigoplus W'$ for some complement W'. Let P be the projection $V \twoheadrightarrow W$ such that $\ker P = W'$ and define P^0 as

$$P^0 = \frac{1}{|G|} \sum_{g \in G} \pi(g) P \pi(g)^{-1}.$$

We see that $P^0(V) \subset W$ as $P(V) \subset W$ and W is stable under π. For $x \in W$, we get $\pi(g)x \in W$, so $P\pi(g)(x) = \pi(g)x$ which implies that $\pi(g)P\pi(g)^{-1}x = x$. Therefore, P^0 is also a projection $V \twoheadrightarrow W$. We now need to prove that

$$W^0 := \ker P^0$$

is π-stable (we know that $V = W \bigoplus W^0$ as W^0 is the kernel of a projection on W). We will use that

$$\begin{aligned}\pi(h)P^0\pi(h)^{-1} &= \frac{1}{|G|}\sum_{g\in G}\pi(h)\pi(g)P\pi(g)^{-1}\pi(h)^{-1} \\ &= \frac{1}{|G|}\sum_{g\in G}\pi(hg)P\pi(hg)^{-1} \\ &= P^0\end{aligned}$$

and therefore

$$\pi(h)P^0 = P^0\pi(h).$$

Thus for $v \in W^0$, we get $P^0v = 0$ which implies $P^0\pi(h)v = \pi(h)P^0v = 0$. So $\pi(h)v \in \ker P^0 = W^0$. We conclude that W^0 is π-stable. □

In the following theorem we see why we would like to know all irreducible representations of a group G.

Theorem 1.8. *Every representation is a direct sum of irreducible representations.*

Proof. Let (π, V) be a representation of a group G. We will use induction on $\dim V$. If $\dim V = 1$, then π is irreducible and the theorem is obvious. If $\dim V > 1$, and π is not irreducible, then V is the direct sum of some non-trivial π-stable subspace W and its π-stable complement W^c by Lemma 1.7. Now $\dim W < \dim V$ and $\dim W^c < \dim V$ and therefore we can use the induction hypothesis on W and W^c. As V is the direct sum of W and W^c, we get that π is the direct sum of the sums of the irreducible representations on W and W^c. □

Because of this result, a goal in representation theory is often to find all irreducible representations of a group G. This is not always easy. In fact, the whole of chapter 2 is devoted to finding all irreducible representations of certain groups. We end this subsection with the following important result by Schur:

Lemma 1.9 (Schur's Lemma)**.** *Let (π, V) and (τ, W) be two irreducible representations of G and let $f : V \to W$ be a linear map such that*

$$f \circ \pi(g) = \tau(g) \circ f$$

for all $g \in G$. Then:

1. *if $\pi \ncong \tau$, then $f = 0$,*
2. *if $V = W$ and $\pi = \tau$ then $f = c \cdot \mathrm{Id}$ for some $c \in \mathbb{C}$.*

Proof. If $f = 0$, the lemma is trivially true. Therefore, consider the case where $f \neq 0$. Now look at $x \in \ker f$, then $f(\pi(g)(x)) = \tau(g)f(x) = 0$ and therefore we have $\pi(g)(x) \in \ker f$ for all $g \in G$. But this implies that the subspace $\ker f$ is

stable under the image of π, and by irreducibility, $\ker f = 0$ or $\ker f = V$. The second case is ruled out by our assumption $f \neq 0$, hence $\ker f = 0$. This implies $\operatorname{im} f = W$ as $\operatorname{im} f$ is a π-stable subspace of W stable under the image of τ and τ is irreducible. This proves that f is an isomorphism, so we can conclude now that if $\pi \ncong \tau$, then $f = 0$.

Now assume $V = W$ and $\pi = \tau$, and let λ be an eigenvalue of f. Now, $\ker(f - \lambda \cdot \mathrm{Id}) \neq 0$, and also

$$(f - \lambda \cdot \mathrm{Id}) \circ \pi(g) = \tau(g) \circ (f - \lambda \cdot \mathrm{Id}).$$

But as we have seen in the first part, this implies that $\ker(f - \lambda \cdot \mathrm{Id}) = V$ and therefore $f = \lambda \cdot \mathrm{Id}$. □

This shows also that intertwiners of irreducible representations are unique up to scalars. For example, take two intertwiners $f, g : V \to W$ for irreducible representations (π, V) and (τ, W), then $f \circ g^{-1} : W \to W$ fits the conditions of Schur's Lemma, thus $f \circ g^{-1} = \lambda \cdot \mathrm{Id}$ which results in $f = \lambda g$. This result is very useful in representation theory and is used in the following subsection on characters. We finish this subsection with a consequence of Schur's lemma:

Consequence 1.10. *Let (π, V) be an irreducible representation of a group G and let $Z \subset G$ denote the center. Then $\pi|_Z = \rho \cdot \mathrm{Id}_V$ for some 1-dimensional representation $\rho : Z \to \mathbb{C}^*$.*

Proof. Let (π, V) be an irreducible representation of a group G and let Z be the center of G. For $z \in Z$, we have $zg = gz$ and therefore $\pi(z)\pi(g) = \pi(g)\pi(z)$. As $\pi(z) \neq 0$ (because $\pi(z) \in \mathrm{GL}(V)$), we have a linear map that satisfies all the conditions of Schur's Lemma. We get that $\pi(z) = \rho(z) \cdot \mathrm{Id}$ for some function $\rho : Z \to \mathbb{C}^*$ and ρ is a representation because

$$\rho(zz') \cdot \mathrm{Id} = \pi(zz') = \pi(z)\pi(z') = \rho(z)\rho(z') \cdot \mathrm{Id}\,.$$

□

We call this ρ the *central character* of π.

Consequence 1.11. *Let A be an abelian group. Then all irreducible representations of A are 1-dimensional.*

Proof. Assume A is an abelian group. Then the center Z equals A and therefore every irreducible representation π acts as a 1-dimensional representation. If π was not 1-dimensional it would therefore be the direct sum of its central character and is therefore reducible. □

Example 1.12. The above consequence makes it very easy to calculate all irreducible representations of $(\mathbb{F}_p)^+$, because these are determined by the image of 1, as it generates $(\mathbb{F}_p)^+$. We need p 1-dimensional irreducible representations $\chi_1, ..., \chi_p$ and that obey $\chi_i(1)^p = 1$. Therefore, all irreducible representations are given by $\chi_i(1) = \zeta_p^i$, where ζ_p is a p-th root of unity.

1.1.2 Character theory

To make the search for irreducible representations easier, and to decompose reducible representations into irreducible representations more easily, we will introduce characters. These are functions associated with representations. Character theory is the study of these characters. In this section we will also explain the use of characters as a tool in representation theory of finite groups.

Definition 1.13. The *character* of a representation π, which we write as χ_π is the function defined by $\chi_\pi(g) = \operatorname{Tr} \pi(g)$.

Example 1.14. Again we will use the examples of Example 1.2 to demonstrate characters:

1. Every 1-dimensional representation is equal to its character, so $\chi_{\rho_{\text{triv}}} = \rho_{\text{triv}}$. For this reason we will often call 1-dimensional representations characters.

2. Taking the trace of the matrices, we get $\chi_\rho(e) = 2$, $\chi_\rho(a) = 1 + \zeta_3$ and $\chi_\rho(a^2) = 1 + \zeta_3^2$, where ζ_3 is a third root of unity.

3. In general, for the standard representation π of S_n, we get that $\chi_\pi(\sigma)$ equals the number of fixed points of $\sigma \in S_n$. Therefore, if $n = 3$ we get that $\chi_\pi(1, 2, 3) = 0$ and $\chi_\pi(2, 3) = 1$ and so on.

4. Again, this is a 1-dimensional representation and is therefore equal to its character.

The following lemma shows some basic properties of characters.

Lemma 1.15. *Let χ be a representation of an n-dimensional representation of G on the vector space V. Then the following holds:*

1. $\chi(e) = \dim V$,
2. $\chi(g^{-1}) = \overline{\chi(g)}$ *for all* $g \in G$,
3. $\chi(hgh^{-1}) = \chi(g)$ *for all* $g, h \in G$,
4. $\chi_{\pi+\tau} = \chi_\pi + \chi_\tau$.

Proof. 1. We get $\chi(e) = \operatorname{Tr} \mathrm{I}_n = n = \dim V$.

2. First, observe that $\pi(g)$ has finite order, and therefore also the eigenvalues of $\pi(g)$, which we denote with λ_i, have finite order, which means they have absolute value 1. We get that

$$\chi(g^{-1}) = \operatorname{Tr} \pi(g^{-1}) = \operatorname{Tr} \pi(g)^{-1} = \sum_i \lambda_i^{-1} = \sum_i \overline{\lambda_i} = \overline{\operatorname{Tr} \pi(g)} = \overline{\chi(g)}.$$

3. Using the known fact that $\mathrm{Tr}(AB) = \mathrm{Tr}(BA)$ for linear mappings A, B from V to V, we get that

$$\chi(hgh^{-1}) = \mathrm{Tr}\,\pi(hgh^{-1}) = \mathrm{Tr}\,\pi(gh^{-1}h) = \mathrm{Tr}\,\pi(g) = \chi(g).$$

4. We know that

$$(\pi + \tau)(g) = \begin{pmatrix} \pi(g) & 0 \\ 0 & \tau(g) \end{pmatrix}.$$

which implies that

$$\chi_{\pi+\tau} = \mathrm{Tr} \begin{pmatrix} \pi(g) & 0 \\ 0 & \tau(g) \end{pmatrix} = \mathrm{Tr}\,\pi(g) + \mathrm{Tr}\,\tau(g) = \chi_\pi + \chi_\tau.$$

□

We can define an inner product on characters of a group. This inner product will have some very useful properties.

Definition 1.16. Let ρ and χ be characters of representations of a group G. We define the *inner product* on G between ρ and χ as

$$\langle \rho, \chi \rangle_G := \frac{1}{|G|} \sum_{g \in G} \rho(g)\overline{\chi(g)}.$$

Remark. This is actually the natural inner product on the space of *class functions* on G. This is however beyond the scope of this introduction. For a more in-depth approach, see [7].

The usefulness of this inner product is made clear in the following important theorem:

Theorem 1.17. *Let χ and χ' be the characters of representations π and π' of a group G. Then*

1. $\langle \chi, \chi \rangle_G = 1$ *if and only if π is irreducible,*
2. $\langle \chi, \chi' \rangle_G = 0$ *if π and π' are non-isomorphic irreducible representations,*
3. $\langle \chi, \chi' \rangle_G$ *is the number of times that an irreducible representation π' appears in the direct sum decomposition of π,*
4. $\pi \cong \pi'$ *if and only if $\chi = \chi'$.*

We will not prove the theorem here, as it takes long. It can be found in [7]. A corollary of this theorem is the following result that limits the number of possible irreducible representations up to isomorphism.

Lemma 1.18. *The sum of the dimension squared of all irreducible representations (up to isomorphism) is equal to the order of G, or, if $\chi_1, ..., \chi_n$ are the characters of all these irreducible representations (up to isomorphism), we get*

$$\sum_{i=1}^{n} \chi_i(e)^2 = |G|.$$

Again we will not prove this lemma. The following theorem is also a very nice and useful result if one tries to find all irreducible representations of a group G. Again it comes without a proof, which can be found in [7].

Theorem 1.19. *The number of irreducible representations (up to isomorphism) of a group G is equal to the number of conjugacy classes of G.*

Example 1.20. The demonstrate the usefulness of these results let us look back at Example 1.2 (3), the standard representation π of S_3 and compute its decomposition into irreducible representations. As we know the character of π by Example 1.14 (3). we get

$$\begin{aligned}
\langle \rho_{\text{triv}}, \chi_\pi \rangle_{S_3} &= \frac{1}{6} \sum_{\sigma \in S_3} \rho_{\text{triv}}(\sigma) \overline{\chi_\pi(\sigma)} \\
&= \frac{1}{6} \sum_{\sigma \in S_3} \overline{\chi_\pi(\sigma)} \\
&= \frac{1}{6}(1 \cdot 3 + 3 \cdot 1 + 2 \cdot 0) \\
&= 1.
\end{aligned}$$

Therefore, $\pi \cong \rho_{\text{triv}} + \tau$ for some 2-dimensional represenation τ. Now, using Theorem 1.17, we get

$$\begin{aligned}
\langle \chi_\tau, \chi_\tau \rangle_{S_3} &= \frac{1}{6} \sum_{\sigma \in S_3} |\chi_\tau(\sigma)|^2 \\
&= \frac{1}{6} \sum_{\sigma \in S_3} |\chi_\pi(\sigma) - \rho_{\text{triv}}(\sigma)|^2 \\
&= \frac{1}{6}(1 \cdot 2^2 + 3 \cdot 0^2 + 2 \cdot (-1)^2) \\
&= 1.
\end{aligned}$$

Therefore, τ is irreducible and the decomposition of π into irreducible representations is indeed $\rho_{\text{triv}} + \tau$. Notice that there are only three irreducible representations of S_3 according to Theorem 1.19, so we have found them all: ρ_{triv}, ρ_{sgn} and τ

1.1.3 Induced Representations

The last subsection of this short introduction to representation theory of finite groups is devoted to induced representations. We will explain the use of induced

representations and end with the Frobenius reciprocity. We will not give proofs of the statements. The usefulness of the results lies in their applications and not in their proofs.

Let us imagine a situation where we have a finite group G, where we would like to find irreducible representations of G and we know a lot about the representations of a subgroup $H \subset G$. It would then be useful to be able to find representations of G from representations of H. It turns out that this is possible. This idea, where we create representations of G from representations of H is called *inducing representations* and will be used in chapter 2 to find irreducible representations of certain groups. There are a number of ways to define induced representations, I have chosen the following definition.

Definition 1.21. Let G be a group and let H be a subgroup of G with a represenation (τ, V). Then the *induced representation* of τ from H to G is the representation (π, W), where

$$W = \{f : G \to V \mid f(hx) = \tau(h)f(x)\ \forall h \in H,\ x \in G\},$$

and the action of π is defined as

$$(\pi(g)f)(x) = f(xg).$$

We will usuall write $\pi = \mathrm{Ind}_H^G(\tau)$

We see that the dimension of a representation induced from H to G is equal to the index $[G : H] \cdot \dim V$. The definition of an induced representation may not be very clear on first sight. However, a short example may provide some clarity.

Example 1.22. Let $G = S_3$ and let $H = \langle (1,2,3) \rangle$. We then have the (irreducible) representation $\rho : (1,2,3) \mapsto \zeta_3$ of H, where ζ_3 is a third root of unity. We can write $G/H = \{eH, (1,2)H\}$ and we see that $\mathrm{Ind}_H^G(\rho)$ is a 2-dimensional representation. We can then express $\mathrm{Ind}_H^G(\rho)$ in matrix form in the following way:

$$\mathrm{Ind}_H^G(\rho)(1,2) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

$$\mathrm{Ind}_H^G(\rho)(1,2,3) = \begin{pmatrix} \zeta_3 & 0 \\ 0 & \zeta_3^2 \end{pmatrix},$$

from which the action of $\mathrm{Ind}_H^G(\rho)$ follows for all $\sigma \in S_3$, as $(1,2)$ and $(1,2,3)$ generate S_3. It is not hard, although it requires quite a calculation, to check that this is isomorphic to the representation from Definition 1.21. Furthermore, this representation is irreducible as it is isomorphic to the τ we found in Example 1.20, which we can easily verify by comparing their characters and using the fact that identical characters imply isomorphic representations.

In the rest of the text, we will rely on the fact that we *can* induce a representation, while we will only rarely explicitly construct this representation. However, we will use the following result on the character of such an induced representation by George Mackey (1916-2006).

Theorem 1.23. *Let τ be a representation of a subgroup H of a group G and let χ be its character. Let $\pi = \mathrm{Ind}_H^G(\tau)$ and let χ_π be its character. Then*

$$\chi_\pi(g) = \sum_{s \in G/H} \hat{\chi}(sgs^{-1})$$

for all $g \in G$ where $s \in G/H$ denotes the representatives of the cosets in G/H and where

$$\hat{\chi}(g) = \begin{cases} \chi(g) & \text{if } g \in H, \\ 0 & \text{if } g \notin H. \end{cases}$$

We will conclude this subsection (and this section on representation theory) with an important result on induced representations, the Frobenius reciprocity. This result is frequently used in the rest of the text and is in general very useful in representation theory. We will first need to introduce the *restriction* of a representation.

Definition 1.24. Let $\pi : G \to \mathrm{GL}(V)$ be a representation of G and let H be a subgroup of G. Then we can restrict π to H and get the *restricted representation* which we denote with $\mathrm{Res}_H^G(\pi)$. It is defined by

$$\mathrm{Res}_H^G(\pi) : H \to \mathrm{GL}(V)$$

$$\mathrm{Res}_H^G(\pi)(h) = \pi(h)$$

for all $h \in H$.

Theorem 1.25 (Frobenius reciprocity)**.** *Let G be a group and let H be a subgroup of G. Let τ be a representation of H and let π be a representation of G. Then we have*

$$\langle \mathrm{Ind}_H^G(\tau), \pi \rangle_G = \langle \tau, \mathrm{Res}_H^G(\pi) \rangle_H.$$

Remark. For those who are familiar with basic category theory, the Frobenius reciprocity states that Ind_H^G and Res_H^G are adjoint functors between the categories $\mathrm{Rep}(H)$ of representations of H and $\mathrm{Rep}(G)$ of representations of G. To be precise, Ind_H^G is a left adjoint to Res_H^G for every G, but if G is finite also a right adjoint. This means we can write

$$\mathrm{Hom}_G(\mathrm{Ind}_H^G(\tau), \pi) \cong \mathrm{Hom}_H(\tau, \mathrm{Res}_H^G(\pi)),$$

which results in the simpler case of Theorem 1.25 since

$$\langle \pi, \pi' \rangle_G = \dim \mathrm{Hom}_G(\pi, \pi').$$

1.2 Symplectic spaces

In this section we will introduce basic theory on symplectic vector spaces. This is done along the lines of [4]. We will first discuss symplectic forms, then subspaces of symplectic spaces and we will finish with the symplectic group.

1.2.1 Symplectic forms

We start with the basic definition of a form, and some different types of forms. We then give an example of a symplectic form and prove an important lemma on the existence of symplectic forms. For the rest of the section, let V be a vector space over ground field K.

Definition 1.26. A *bilinear form* φ is a function $\varphi : V \times V \to K$ satisfying

1. $\varphi(u + v, w) = \varphi(u, w) + \varphi(v, w)$,
2. $\varphi(u, v + w) = \varphi(u, v) + \varphi(u, w)$,
3. $\varphi(\lambda u, v) = \varphi(u, \lambda v) = \lambda\varphi(u, v)$,

for all vectors $u, v, w \in V$ and scalars $\lambda \in K$.

Furthermore,

Definition 1.27. A bilinear form is called:

1. *symmetric*, if $\varphi(u, v) = \varphi(v, u)$ for all vectors $u, v \in V$,
2. *alternating*, if $\varphi(v, v) = 0$ for all vectors $v \in V$,
3. *skew-symmetric*, if $\varphi(u, v) = -\varphi(v, u)$ for all vectors $u, v \in V$,
4. *non-degenerate*, if $\varphi(u, v) = 0$ for all $v \in V$ implies that $u = 0$,
5. *symplectic*, if φ is both *alternating* and *non-degenerate*.

Lemma 1.28. *An alternating form φ is also skew-symmetric. The converse is true if* $\mathrm{char}(K) \neq 2$.

Proof. Let φ be an alternating form. Then

$$0 = \varphi(u + v, u + v) = \varphi(u, u) + \varphi(u, v) + \varphi(v, u) + \varphi(v, v).$$

Therefore $\varphi(u, v) = -\varphi(v, u)$. Now let $\mathrm{char}(K) \neq 2$ and let φ be a skew-symmetric form. Then

$$\varphi(u, u) = -\varphi(u, u) \;\Rightarrow\; \varphi(u, u) = 0.$$

□

Remark. If $\mathrm{char}(K) = 2$ and if φ is a skew-symmetric form, then φ is also a symmetric form.

First of all, let us look at an easy example of a symplectic form:

Example 1.29. Let $V = \mathbb{R}^2$ and $\varphi\left(\begin{pmatrix} a_1 \\ a_2 \end{pmatrix}, \begin{pmatrix} b_1 \\ b_2 \end{pmatrix}\right) = a_1b_2 - a_2b_1$. Then φ is a symplectic form.

We will now look at subspaces of V defined by the symplectic form on V. This will lead to a nice result on the existence of symplectic forms for vector spaces.

Definition 1.30. Let W be a subspace of V. Then the *symplectic subspace* $W^\perp$ is defined as

$$\{v \in V \mid \varphi(v, w) = 0 \ \forall w \in W\}$$

Lemma 1.31. *These properties hold for all linear subspaces W of V:*

a) $\dim(W) + \dim(W^\perp) = \dim(V)$

b) $(W^\perp)^\perp = W$

Proof. Let $W \subset V$ be a vector space. For a v in V, we can look at $\varphi(v, -) \in V^\vee$, where $V^\vee$ denotes the dual space of V. Because φ is symplectic, this gives us an isomorphism $V \xrightarrow{\sim} V^\vee$. We can project $V^\vee$ onto $W^\vee$, which has the same dimension as W. Now $W^\perp$ is exactly the kernel of the composition

$$V \xrightarrow{\sim} V^\vee \twoheadrightarrow W^\vee$$

because $w^\perp \in W^\perp$ if and only if $\varphi(w^\perp, -) = 0$ on W. This gives us a), since now $\dim W^\perp + \dim W^\vee = \dim V$ by the rank-nullity theorem. For part b), observe that $W \subset (W^\perp)^\perp$ and use part a) to get

$$\dim W^\perp + \dim W = \dim V = \dim W^\perp + \dim (W^\perp)^\perp$$

which implies $\dim W = \dim (W^\perp)^\perp$ and therefore $W = (W^\perp)^\perp$.

□

From now on we will call (V, φ) a *symplectic vector space* if φ is a symplectic form on V. Not all vector spaces can have a symplectic form and we have an easy criterion:

Lemma 1.32. *If a symplectic form $\varphi : V \times V \to K$ exists on V then the dimension of V is even and there exists a base of V relative to which φ is given by the matrix*

$$\begin{pmatrix} 0 & J \\ -J & 0 \end{pmatrix}, \text{where } J = \begin{pmatrix} 0 & \cdots & 0 & 0 & 1 \\ 0 & \cdots & 0 & 1 & 0 \\ 0 & \cdots & 1 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \cdots & 0 & 0 & 0 \end{pmatrix}.$$

Proof. Let $V \neq \{0\}$ (otherwise the result is trivial) and take a non-zero $v_1 \in V$. Because φ is non-degenerate, there exists a non-zero $w_1' \in V$ such that $\varphi(v_1, w_1') = \lambda_1 \neq 0$. Therefore, $\varphi(v_1, w_1) = -1$ for $w_1 = -\frac{w_1'}{\lambda_1}$, which gives $v_1 \neq \alpha w_1$, as otherwise $\varphi(v_1, w_1) = 0$. So v_1 and w_1 are linearly independent. Define $W := \mathrm{Span}(v_1, w_1)$ and consider its symplectic complement $W^\perp$:

$$W^\perp = \{w \in V \mid \varphi(w, v_1) = 0,\ \varphi(w, w_1) = 0\}.$$

By Lemma 1.31 (a) we get that $\dim W^\perp = \dim V - 2$. We claim that φ is a symplectic form on $W^\perp$. Our proof of the first part will then follow from induction.

$\varphi_{|W}$ is a symplectic form on $W^\perp$, since:

- $\varphi_{|W^\perp}$ is alternating, because $\varphi_{|W^\perp}(w, w) = \varphi(w, w) = 0$ for all $w \in W^\perp$,
- $\varphi_{|W^\perp}$ is non-degenerate because if $\varphi_{|W^\perp}(w, w') = 0$ for all $w' \in W^\perp$ then $\varphi(w, v) = 0$ for all $v \in V$, as $\varphi(w, v) = \varphi(w, v_W) + \varphi(w, v_{W^\perp}) = 0 + 0$. Then, by non-degeneracy of φ, we get $w = 0$.

By induction, it now follows that V is of dimension $2n$ for some $n \in \mathbb{N}$. For the construction of our base we repeat the first part of the proof for $W^\perp$ instead of V to find v_2 and w_2 such that $\varphi(v_2, w_2) = -1$. By induction, we find a basis $\{v_1, ..., v_n, w_n, ..., w_1\}$ for V with the following properties:

1. $\varphi(v_i, v_j) = 0$,
2. $\varphi(v_i, w_j) = -\delta_{ij}$,
3. $\varphi(w_i, v_j) = \delta_{ij}$,
4. $\varphi(w_i, w_j) = 0$.

We call this basis a *Darboux basis* or *symplectic basis*, and it is a base of V relative to which φ is given by the matrix $\begin{pmatrix} 0 & J \\ -J & 0 \end{pmatrix}$.

□

1.2.2 Subspaces of Symplectic Vector spaces

From now on, let $\dim V = 2n$. We will discuss subspaces of symplectic vector spaces with several different properties. We start with an illustrating example.

Example 1.33. We see that the symplectic form of example 1.29 is the only possible symplectic form on a 2-dimensional symplectic vector space. If we pick base $\{(1,0),(0,1)\}$, we must have

$$\varphi\left(\begin{pmatrix}1\\0\end{pmatrix}, \begin{pmatrix}0\\1\end{pmatrix}\right) = 1, \quad \varphi\left(\begin{pmatrix}0\\1\end{pmatrix}, \begin{pmatrix}1\\0\end{pmatrix}\right) = -1,$$

and the rest follows from linearity:

$$\begin{aligned}\varphi\left(\begin{pmatrix}a\\b\end{pmatrix},\begin{pmatrix}c\\d\end{pmatrix}\right) &= \varphi\left(\begin{pmatrix}a\\0\end{pmatrix},\begin{pmatrix}0\\d\end{pmatrix}\right)+\varphi\left(\begin{pmatrix}0\\b\end{pmatrix},\begin{pmatrix}0\\d\end{pmatrix}\right)\\ &\quad+\varphi\left(\begin{pmatrix}a\\0\end{pmatrix},\begin{pmatrix}c\\0\end{pmatrix}\right)+\varphi\left(\begin{pmatrix}0\\b\end{pmatrix},\begin{pmatrix}c\\0\end{pmatrix}\right)\\ &= ad-bc.\end{aligned}$$

Unlike in vector spaces with an inner product, we could easily have $W \cap W^\perp \neq \{0\}$. For example, take $W = \mathbb{C}\cdot v$ for some non-zero v. Then $\varphi(\lambda_1 \cdot v, \lambda_2 \cdot v) = 0$ which gives us $W \subset W^\perp$. We will define the following instances:

Definition 1.34. We call a linear subspace W of V

- *symplectic*, if $W \cap W^\perp = \{0\}$,
- *isotropic*, if $W \subset W^\perp$,
- *co-isotropic*, if $W^\perp \subset W$,
- *Lagrangian*, if $W = W^\perp$.

Example 1.35. Some easy examples are found if we let V be a symplectic vector space and let $\{e_1, ..., e_n, f_n, ..., f_1\}$ be a Darboux base. Then

- $\mathrm{Span}(e_i, f_i)$ is symplectic for all $1 \leq i \leq n$,
- $\mathrm{Span}(e_1, e_2, ..., e_k)$ is isotropic ($k \leq n$),
- $\mathrm{Span}(e_1, e_2, ..., e_n, f_i)$ is co-isotropic,
- $\mathrm{Span}(e_1, e_2, ..., e_n)$ is Lagrangian.

Using Lemma 1.31, we can prove the following lemma on these kind of subspaces:

Lemma 1.36. *Let $M \subset V$ be an isotropic subspace, then $\dim(M) \leq n$. Furthermore, if M is Lagrangian, $\dim(M) = n$.*

Proof. We know that $\dim(M) + \dim(M^\perp) = 2n$. Thus if $M \subset M^\perp$ we get $\dim(M) \leq \dim(M^\perp)$ and thus $2 \cdot \dim(M) \leq 2n$. Furthermore, if $M = M^\perp$ we get $2 \cdot \dim(M) = 2n$. □

We can use these Lagrangian subspaces and their connection to dual spaces to construct a decomposition of V, as the following example shows.

Example 1.37. If we have a Lagrangian subspace $L = \mathrm{Span}(e_1, e_2, ..., e_n)$ then also $L' = \mathrm{Span}(f_1, f_2, ..., f_n)$ is Lagrangian and we get $V = L \oplus L'$. Now, we can use φ to construct an isomorphism from $L' \xrightarrow{\sim} L^\vee$, where $l' \mapsto \varphi(l', -)$. We find that $V \cong L \oplus L^\vee$, but notice that this is not a canonical decomposition, as it depends on the choice of our Darboux basis. Furthermore, if we write $v = l + \gamma$ for $v \in V$, $l \in L$ and $\gamma \in L^*$ then we can write $\varphi(l_1+\gamma_1, l_2+\gamma_2) = \gamma_2(l_1) - \gamma_1(l_2)$.

1.2.3 The Symplectic Group $\mathbf{Sp}(V, \varphi)$

Let V be a symplectic vector space again, with φ a symplectic form. There is a group of transformations associated with (V, φ) called the symplectic group. It is found by looking at transformations that preserve our form φ.

Definition 1.38. The *symplectic group* $\mathrm{Sp}(V, \varphi)$ is the group of invertible linear transformations $T : V \to V$ such that for every v, w in V we get $\varphi(Tv, Tw) = \varphi(v, w)$.

Example 1.39. Recall our example $V = \mathbb{R}^2$, $v = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}, w = \begin{pmatrix} b_1 \\ b_2 \end{pmatrix}$ and $\varphi(v, w) = a_1 b_2 - a_2 b_1$. Let $A = \begin{pmatrix} x & y \\ z & w \end{pmatrix}$. Then

$$\begin{aligned}
\varphi(A\begin{pmatrix} a_1 \\ a_2 \end{pmatrix}, A\begin{pmatrix} b_1 \\ b_2 \end{pmatrix}) &= \varphi(\begin{pmatrix} xa_1 + ya_2 \\ za_1 + wa_2 \end{pmatrix}, \begin{pmatrix} xb_1 + yb_2 \\ zb_1 + wb_2 \end{pmatrix}) \\
&= (xa_1 + ya_2)(zb_1 + wb_2) - (za_1 + wa_2)(xb_1 + yb_2) \\
&= xza_1b_1 + yza_2b_1 + xwa_1b_2 + ywa_2b_2 \\
&\quad - xza_1b_1 - xwa_2b_1 - yza_1b_2 - ywa_2b_2 \\
&= (xw - yz)(a_1b_2 - a_2b_1) = \det(A) \cdot \varphi(v, w).
\end{aligned}$$

We find that $\mathrm{Sp}(\mathbb{R}^2, \varphi) = \mathrm{Sl}_2(\mathbb{R})$.

Chapter 2

Representation Theory and the Heisenberg Groups

This chapter is split into three sections. The first section introduces nilpotent groups, the second section introduces Heisenberg groups, which are nilpotent, and the third section combines the results from the first two sections to prove the Stone-von Neumann theorem, which is the goal of this chapter.

2.1 Nilpotent Groups

In this section we define nilpotent groups, show some examples of nilpotent groups and prove some results about nilpotent groups in representation theory. These results will be used later on in the sections about Heisenberg groups and representation theory.

To define nilpotent groups, we first define normal series.

Definition 2.1. Let G be a group. A *normal series* of G is a sequence of subgroups $G_i \neq G_j$ for $i \neq j$ such that

$$G = G_0 \trianglerighteq G_1 \trianglerighteq ... \trianglerighteq G_{n-1} \trianglerighteq G_n = 1.$$

We call n the length of the normal series.

An easy example of a normal series is $S_n \trianglerighteq A_n \trianglerighteq 1$. Every group G has the normal series $G \trianglerighteq Z(G) \trianglerighteq 1$, where $Z(G)$ is the center of G.

Definition 2.2. A *nilpotent group* is a group G such that it has a normal series of finite length n for which G_i/G_{i+1} is in the center of G/G_{i+1} for all $i < n$. We call such a series a *central series*. The smallest n for which such a central series exists is called the *nilpotency class* of G and we will then call G nilpotent of class n.

Example 2.3. There is of course only one group of nilpotency class 0, the trivial group. Furthermore, the only groups of nilpotency class 1 are the non-trivial Abelian groups. A more interesting example of a nilpotent group is Q_8, the quaternion group. It is non-abelian and has the central series

$$Q_8 \trianglerighteq \{1, -1\} \trianglerighteq 1.$$

It is nilpotent because -1 commutes with every element of Q_8 and therefore $\{1, -1\}/1$ is in the center of $Q_8/1$. Furthermore $Q_8/\{1, -1\}$ is of order 4 and therefore abelian, which implies it is contained in its own center. Its nilpotency class is 2.

The following properties of nilpotent groups are in general useful and will be used later on.

Lemma 2.4. *1. Every subgroup of a nilpotent group is nilpotent.*

2. Every quotient group of a nilpotent group is nilpotent.

Proof. Let G be a nilpotent group, H be a subgroup of G and N be a normal subgroup of G.

1. We need to find a central series for H. Take $H_i = G_i \cap H$. It is of course possible that $H_j = 1$ for some $j < n$, but then we just truncate our series. We get the central series

$$H = G \cap H = G_0 \cap H \trianglerighteq G_1 \cap H \trianglerighteq ... \trianglerighteq G_n \cap H = 1 \cap H = 1.$$

   We also need to show that $H_i/H_{i+1} \subset Z(H/H_{i+1})$. This follows because $H/H_{i+1} \subset G/G_{i+1}$ and $G_i/G_{i+1} \subset Z(G/G_{i+1})$ and therefore $H_i/H_{i+1} = (H/H_{i+1}) \cap (G_i/G_{i+1})$. Combining these results we get that $H_i/H_{i+1} \subset Z(H/H_{i+1})$. So, H is nilpotent.

2. We need to find a central series for G/N. The image of G_i under the canonical map $G \to G/N$ will give us this central series. Again it is possible that $G_j \mapsto 1$ for some $j < n$, and again we will then truncate our series. So, N is nilpotent.

□

A consequence of this lemma is that if G contains a simple non-abelian subgroup, then G is *not* nilpotent. This follows because a simple non-abelian group is not nilpotent, as it has no non-trivial normal subgroups, and therefore it has no central series. For example, the symmetric group S_n for $n \geq 5$ is not nilpotent, because A_n is simple for $n \geq 5$ and therefore not nilpotent.

Lemma 2.5. *If G is a non-abelian nilpotent group, then there is a normal subgroup A that is nilpotent, non-central and abelian.*

Proof. Let G be a non-abelian nilpotent group. Define

$$Z_2(G) = \{g \in G \mid \forall y \in G\ [g, y] \in Z(G)\}$$

and take $a \in Z_2(G) - Z(G)$, which is non-empty as G is nilpotent.

Now define $A := \langle Z(G), a\rangle$. We get that A is nilpotent by Lemma 2.4 and non-central because $a \notin Z(G)$. Clearly A is abelian. We are left to prove that A is normal. Take $g \in G$. Now $[a, g] = aga^{-1}g^{-1} = z$ for some $z \in Z(G)$, so $ga^{-1}g^{-1} = a^{-1}z \in A$. We conclude that A is the subgroup we were looking for. □

Lemma 2.6. *Let V be a complex vector space. If (π, V) is a finite-dimensional irreducible representation of a nilpotent group G, then (π, V) is induced from a one dimensional representation of some subgroup of G.*

Remark. This lemma holds for all nilpotent groups (see, for example [3]). However, we will only prove it for finite nilpotent groups, because the focus of this thesis is on those groups.

Proof. If G is abelian, the result of Lemma 2.6 is trivial . Thus, let G be a finite nilpotent group which is not abelian and let (π, V) be an irreducible representation of G. We will proof this by induction on $\dim V$ and $|G|$. If $\dim V = 1$, it is trivial. Now let $\dim V > 1$.

First, assume π is not injective. Let us write $K = \ker \pi$. Then π is an irreducible representation on G/K, which is a nilpotent group of a smaller order than G. By the inductive hypothesis, $\pi = \mathrm{Ind}_{H/K}^{G/K}(\rho)$ for some $H \subset G$ and some representation ρ of H. But then we can lift π to G with $\pi \cong \mathrm{Ind}_H^G(\rho)$, because

$$G/H \cong (G/K)/(H/K).$$

Now consider the case where π is injective. G is nilpotent and not abelian, so by Lemma 2.5, there is an abelian normal non-central subgroup $A \subsetneq G$. If we now look at $\pi|_A$, we see that it must be a sum of characters, as A is abelian. We can pick one of these characters, χ, which acts on a subspace $W \subset V$ and we can then look at the subgroup $F = \{g \in G \mid \chi^g = \chi\}$. Because A is not central, there must be a $g' \in G$ such that $ag' \neq g'a$ and, because π is injective, this gives $\pi(g'ag'^{-1}) \neq \pi(a)$, which implies $\chi^{g'} \neq \chi$. We find that F must be a strict subset of G. Because F preserves W, we can look at the representation ρ of F on W. We claim that

$$\pi \cong \mathrm{Ind}_F^G(\rho).$$

By construction of $\mathrm{Ind}_F^G(\rho)$, we know that it acts on a function space of functions from $G \to W$, which we call $\mathcal{H}$. Let g_j be the representatives of G/F and define $f_i^v \in \mathcal{H}$ by $f_i^v(g_j) = \delta_{ij}v$ for $v \in W$, where δ_{ij} is the Kroneker delta. These f_i^v

span $\mathcal{H}$, because a function in $\mathcal{H}$ is fixed by its values on g_j. We now construct an intertwiner h between π and $\mathrm{Ind}_F^G(\rho)$ by

$$h : \mathcal{H} \to V$$
$$h(f_i^v) = \pi(g_i^{-1})v$$

and it follows therefore that $\pi \cong \mathrm{Ind}_F^G(\rho)$. Thereby the claim is proven, and by proving the claim, our proof is complete using the inductive argument on $\dim V$ and the order of G. □

The following theorem on finite nilpotent groups is a nice general result on these groups. A very similar result holds for finitely generated nilpotent groups and can be found in [3]. That theorem is however beyond the focus of this text.

Theorem 2.7 (Brown's Criterion)**.** *Let G be a finite nilpotent group. Then a finite-dimensional representation π is irreducible if and only if there is no $s \in G - H$ such that $\rho^s = \rho$ on $H^s \cap H$, where $\rho : H \to \mathbb{C}^*$ is the character such that $\pi = \mathrm{Ind}_H^G(\rho)$.*

Proof. Let π be a finite dimensional representation of a finite nilpotent group G. Let ρ be the character such that $\pi = \mathrm{Ind}_H^G(\rho)$. By Lemma 1.23 we get that

$$\chi_\pi(g) = \sum_{s \in G/H} \hat{\rho}(sgs^{-1}),$$

where $\hat{\rho} = \rho$ on H and $\hat{\rho} = 0$ on $G - H$. Using Frobenius Reciprocity, we get that

$$\begin{aligned}
\langle \chi_\pi, \chi_\pi \rangle_G &= \langle \rho, \chi_\pi \rangle_H \\
&= \frac{1}{|H|} \sum_{g \in H} \rho(g)\overline{\chi_\pi(g)} \\
&= \frac{1}{|H|} \sum_{g \in H} \rho(g) \sum_{s \in G/H} \overline{\hat{\rho}(sgs^{-1})} \\
&= \frac{1}{|H|} \sum_{g \in H} \sum_{s \in G/H} \rho(g)\overline{\hat{\rho}(sgs^{-1})} \\
&= \sum_{s \in G/H} \frac{1}{|H|} \sum_{g \in H} \rho(g)\overline{\hat{\rho}(sgs^{-1})} \\
&= \sum_{s \in G/H} \frac{1}{|H|} \sum_{h \in H \cap H^s} \rho(h)\overline{\rho^s(h)} \\
&= \sum_{s \in G/H} \frac{|H \cap H^s|}{|H|} \left(\frac{1}{|H \cap H^s|} \sum_{h \in H \cap H^s} \rho(h)\overline{\rho^s(h)} \right) \\
&= \sum_{s \in G/H} \frac{|H \cap H^s|}{|H|} \langle \rho, \rho^s \rangle_{H \cap H^s}
\end{aligned}$$

We know that in this case, since ρ and ρ^s are both characters of $H \cap H^s$, that

$$\langle \rho, \rho^s \rangle_{H \cap H^s} = \begin{cases} 1 & \text{if } \rho = \rho^s, \\ 0 & \text{if } \rho \neq \rho^s. \end{cases}$$

This then concludes the proof, because for $e \in G/H$ we get that $\rho = \rho^e$ and $\frac{|H \cap H^e|}{|H|} = 1$. Therefore, π is irreducible if and only if $\langle \chi_\pi, \chi_\pi \rangle_G = 1$ if and only if there is no other $s \in G/H$ such that $\rho^s = \rho$ on $H^s \cap H$.

□

2.2 Heisenberg Groups

The goal of this section is to introduce a family of groups known as the *Heisenberg groups*. We will first define Heisenberg groups and give an easy example of a Heisenberg group. After that, we compute its conjugacy classes and show a subfamily of Heisenberg groups that will be important in the following chapters.

Let H be a group such that $V := H/Z(H)$ is abelian and let π be the projection on V. This gives us the following exact sequence

$$0 \to Z(H) \hookrightarrow H \xrightarrow{\pi} V \to 0,$$

and a pairing $e : V \times V \to Z(H)$, which we call the *commutator pairing*.

For v in V, we denote by $\tilde{v}$ an element in H such that $\tilde{v}$ modulo the center equals v, or, $\pi(\tilde{v}) = v$. The commutator pairing e maps (v, w) to the projection of $[\tilde{v}, \tilde{w}]$ on the center. It is therefore *alternating*, that is $e(v, v) = 0$ for all $v \in V$. This follows because $[\tilde{v}, \tilde{v}]$ is always 0. It is also symplectic, because it is non-degenerate as the following lemma will show:

Lemma 2.8. *The commutator pairing e of a Heisenberg group is non-degenerate.*

Proof. We will prove that if $e(v, w) = 0$ for all $w \in V$, then $v = 0$. Assume $e(v, w) = 0$ for all $w \in V$, then $[\tilde{v}, \tilde{w}] = 0$ for all $w \in V$. Therefore $\tilde{v}$ commutes with all $h \in H$ and thus $\tilde{v} \in Z(H)$ which means $v = 0$. □

Definition 2.9. A *Heisenberg group* is a group H such that $H/Z(H)$ is abelian and for every $v \neq 0$ in V the function $e(v, -) : V \to Z(H)$ is surjective.

The reason we introduced nilpotent groups and proved some lemmas on representation theory in the previous section is the following lemma:

Lemma 2.10. *A Heisenberg group is nilpotent of class 2.*

Proof. Because $H/Z(H)$ is abelian (and therefore equal to $Z(H/Z(H))$) we get the central series

$$H \trianglerighteq Z(H) \trianglerighteq 1.$$

□

Example 2.11. An example of a Heisenberg group, which we will also use in further sections, is the Heisenberg group

$$H(\mathbb{F}_q) := \left\{ \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} \;\middle|\; a,\, b,\, c \in \mathbb{F}_q \right\} \subset \mathrm{GL}_3(\mathbb{F}_q).$$

The center of $H(\mathbb{F}_q)$ is

$$Z := \left\{ \begin{pmatrix} 1 & 0 & c \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \;\middle|\; c \in \mathbb{F}_q \right\}.$$

We see that H/Z is abelian, as it is isomorphic to $\mathbb{F}_q \times \mathbb{F}_q$. Furthermore the function $e((x,y),-)$ is surjective, because if we compute $e((x,y),(a,b))$, we get

$$\begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & a & c \\ 0 & 1 & b \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & -x & -z+xy \\ 0 & 1 & -y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & -a & -c+ab \\ 0 & 1 & -b \\ 0 & 0 & 1 \end{pmatrix}$$
$$=$$
$$\begin{pmatrix} 1 & 0 & bx-ya \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

and as $\mathbb{F}_q$ is a field, for every $\alpha \in \mathbb{F}_q$ there exist $a, b \in \mathbb{F}_q$ such that $bx - ya = \alpha$. We see that $e((x,y),-)$ is surjective and conclude that $H(\mathbb{F}_q)$ is a Heisenberg group.

The simplicity of its conjugacy classes of a Heisenberg group is one of the properties that will be useful to us in the next section on representation theory and the Heisenberg groups.

Lemma 2.12. *The conjugacy classes of a Heisenberg group H are $\{z\}$ for elements in the center, and the cosets $hZ(H)$ for elements $h \in H - Z(H)$.*

Proof. The conjugacy classes of the center are obvious. Now take $h \in H$ such that $h \notin Z(H)$ and write $\overline{h}$ as the projection of h on V. We need to prove that for every $z \in Z(H)$ there is an $x \in H$ such that $xhx^{-1} = zh$. We know that $e(\overline{h}, w) \in Z(H)$ for all $w \in V$, therefore the conjugacy class of $\overline{h}$ is contained in $hZ(H)$. Now because $e(\overline{h}, -)$ is surjective, we find that such an $x \in H$ must exist and therefore the conjugacy class of h equals $hZ(H)$. □

Lastly, we will identify two subfamilies of Heisenberg groups. The first family will be used in chapter 4.

Lemma 2.13. *The group $H = V \times K$ where K is a field and V a vector space with a symplectic form φ, and group law*

$$(v,k)(v',k') = (v+v', k+k'+\frac{1}{2}\varphi(v,v')),$$

is a Heisenberg group.

Proof. It is easy to see that $K \cong Z(H)$ (by $k \mapsto (0,k)$) and therefore that $H/Z(H) = V$ is abelian. We see by calculation that $e(v,w) = \varphi(v,w)$ therefore we will need to prove that $\varphi(v,-)$ is surjective, for $v \neq 0$. However, because φ is non-degenerate there exists a w such that $\varphi(v,w) \neq 0$. Therefore, for a $k \in K$ we get that

$$\varphi\left(v, \frac{k}{\varphi(v,w)} \cdot w\right) = \frac{k}{\varphi(v,w)} \cdot \varphi(v,w) = k,$$

and therefore $\varphi(v,-)$ is surjective. □

The second family of Heisenberg groups is based on the (non-canonical) decomposition of a vector space V into $L \bigoplus L^\vee$ for a Lagrangian subspace L, as we saw in Example 1.37. We can then construct the Heisenberg group $H = V \times K$, where K is the ground field of V, with group law

$$(l_1, \gamma_1, k_1) \cdot (l_2, \gamma_2, k_2) = (l_1 + l_2, \gamma_1 + \gamma_2, k_1 + k_2 + \gamma_1(l_2).$$

Actually, Example 2.11 is of this type, as we can see the matrices with $b, c = 0$ as a Lagrangian subspace L. It is easy to see that K is the center and that $H/K \cong V$, which is abelian. We have left out the proof that its commutator pairing is surjective.

2.3 Representation Theory of Heisenberg Groups

In this section we prove the Stone-von Neumann theorem for finite Heisenberg groups, which will conclude this chapter. The Stone-von Neumann theorem is important for the following chapters and combines the knowledge from the previous sections on the Heisenberg groups and representation theory. For the rest of this section, assume all Heisenberg groups are non-abelian; otherwise the results are trivial.

We will need the next lemma.

Lemma 2.14. *If an irreducible representation π of a finite Heisenberg group H is induced from some subgroup B and character ρ, then $Z(H) \subsetneq B \subset H$.*

Proof. First, assume $z \in Z(H)$ and $z \notin B$. But then $\rho^z(g) = \rho(zgz^{-1}) = \rho(g)$ and therefore π cannot be irreducible according to Brown's criterion. We see that we must have $Z(H) \subset B$. Now let us assume $Z(H) = B$, then we see that, for $z \in Z(H)$,

$$\chi_\pi(z) = \sum_{v \in V = H/Z(H)} \rho(v^{-1}zv) = \sum_{v \in V} \rho(zvv^{-1}) = |V|\rho(z)$$

and we see that π is not irreducible because

$$\begin{aligned}\langle \chi_\pi, \chi_\pi \rangle &= \frac{1}{|H|} \sum_{h \in H} \chi_\pi(h)\overline{\chi_\pi(h)} \\ &\geq \frac{1}{|Z(H)| \cdot |V|} \sum_{z \in Z(H)} |V|^2 \rho(z)\overline{\rho(z)} \\ &\geq |V| \cdot \left(\frac{1}{|Z(H)|} \sum_{z \in Z(H)} \rho(z)\overline{\rho(z)} \right) \\ &\geq |V| \\ &> 1.\end{aligned}$$

Hence we must have $Z(H) \subsetneq B$. □

With this result we can prove the Stone-von Neumann theorem. Recall that $H/Z(H) = V$ and that an irreducible representation has a central character (Lemma 2.14).

Theorem 2.15 (Stone-von Neumann)**.**
Let H be a finite Heisenberg group and let $\psi \neq 1$ be a character of $Z(H)$. Then there exists a unique (up to isomorphism) irreducible representation (π, W) with central character ψ for some vector space W.

Proof. We will first find the representation π. Secondly we calculate its character, which will prove that π is irreducible and thirdly we prove that π is uniquely determined by ψ up to isomorphism. Let $\psi \neq 1$ be a character of $Z(H)$.

I. Finding π. We know (by counting) that there must be an irreducible representation π' of H with dimension greater than 1. By Brown's Criterion, π' is induced from some character ρ of a subset B. We have already shown (in Lemma 2.14) that $Z(H) \subsetneq B$. Therefore, the dimension of such an irreducible representation π' is $[H : B] < [H : Z(H)] = V$. The problem, however, is that we do not know whether or not the central character of π' is ψ or some other character of $Z(H)$, but we can change this central character (and our representation) in the following way. Write z_h as the projection on $Z(H)$ of an element h of H. Now project B onto $Z(H)$ with P and look at the composition $\psi \circ P$:

$$\begin{aligned} B \twoheadrightarrow Z(H) &\xrightarrow{\psi} \mathbb{C} \\ h &\to \psi(z_h). \end{aligned}$$

We then claim to have found the representation we were looking for by

$$\pi = \mathrm{Ind}_B^H(\psi \circ P),$$

as this will certainly have central character ψ and by construction via π' and B be irreducible.

II. The character and irreducibility of π. We will first calculate the character of π and we will use that result to show that π is irreducible. We start with the following claim.

Claim. *The character of π, for which we write χ_π, is given by*

$$\chi_\pi(h) = \begin{cases} \sqrt{|V|} \cdot \psi(h) & \text{if } h \in Z(H) \\ 0 & \text{if } h \notin Z(H) \end{cases}$$

Proof. We prove this in two parts. First we prove that $\chi_\pi(z) = \sqrt{|V|} \cdot \psi(z)$ for $z \in Z(H)$. Secondly, we prove that $\chi_\pi(h) = 0$ for $h \notin Z(H)$.

For the first part we only need to prove that $\dim W = \sqrt{|V|}$. This follows from Frobenius reciprocity:

$$\dim W = \langle \mathrm{Res}^H_{Z(H)}(\pi), \psi \rangle_{Z(H)} = \langle \pi, \mathrm{Ind}^H_{Z(H)}(\psi) \rangle_H.$$

Notice that $V = H/Z(H)$, which means that the dimension of $\mathrm{Ind}^H_{Z(H)}(\psi)$ equals $|V|$. Now, $\dim \pi = \dim W$ and π appears $\dim W$ times in $\mathrm{Ind}^H_{Z(H)}(\psi)$. Therefore, $\dim W = \sqrt{|V|}$.

For the second part, we know that $\chi_\pi(h) = \chi_\pi(h')$ if h and h' are conjugated. We know, by Lemma 2.12, that the conjugacy class of h is $hZ(H)$. Thus,

$$\chi_\pi(h) = \chi_\pi(zh) = \chi_\pi(z)\chi_\pi(h)$$

for all $z \in Z(H)$. As $\psi \neq 1$, there must be some z such that $\psi(z) \neq 1$, which implies that $\chi_\pi(h) = 0$. □

The fact that π is irreducible is now an easy calculation:

$$\begin{aligned} \langle \chi_\pi, \chi_\pi \rangle &= \frac{1}{|H|} \sum_{h \in H} \chi_\pi(h)\overline{\chi_\pi(h)} \\ &= \frac{1}{|V| \cdot |Z(H)|} \sum_{z \in Z(H)} \sqrt{|V|}^2 \psi(z)\overline{\psi(z)} \\ &= \frac{|V|}{|V|} \frac{1}{|Z(H)|} \sum_{z \in Z(H)} \psi(z)\overline{\psi(z)} \\ &= 1. \end{aligned}$$

Here in the second step, we use that the character χ_π is zero outside of the center and in the last step, we used the fact that ψ is a character of $Z(H)$ and therefore irreducible, which implies that $\langle \psi, \psi \rangle = 1$.

III. Uniqueness. We continue our proof of uniqueness. Let τ be an irreducible representation on W_τ such that $\tau(z) = \psi(z)\,\mathrm{Id}_{W_\tau}$ (or equivalently, τ has central character ψ).

If we look at the character of τ, we write χ_τ, then again we get that $\chi_\tau(h) = 0$ if $h \notin Z(H)$. All that is left to proof is that $\dim W_\tau = \sqrt{|V|}$. This follows from the following:

$$\begin{aligned}
1 &= \langle \chi_\tau, \chi_\tau \rangle \\
&= \frac{1}{|H|} \sum_{h \in H} \chi_\tau(h) \overline{\chi_\tau(h)} \\
&= \frac{1}{|V| \cdot |Z(H)|} \sum_{z \in Z(H)} |\dim W_\tau|^2 \rho(z) \overline{\psi(z)} \\
&= \frac{|\dim W_\tau|^2}{|V|} \frac{1}{|Z(H)|} \sum_{z \in Z(H)} \rho(z) \overline{\psi(z)} \\
&= \frac{|\dim W_\tau|^2}{|V|}
\end{aligned}$$

where we used that ψ is irreducible in the fourth step. We see that $\chi_\pi = \chi_{\pi'}$ which implies that π' is isomorphic to π. □

Remark. We see that this theorem works if $|V|$ is a square. It is not clear why this should always be the case. For the examples in this text, it follows because in our cases V is a vector space with a symplectic form and therefore is even-dimensional, which implies that its order is a square. Also, in the Heisenberg groups constructed over a Lagrangian subspace L and its dual space $L^\vee$, we saw that $V \cong L \bigoplus L^\vee$, which implies that the order of V is a square.

However, we do not always have such a decomposition into (Lagrangian) subspaces. Take for example the Quaternion group Q_8, which is a Heisenberg group, as $Q_8/Z(Q_8) \cong V_4$, the (abelian) Klein Vier group and its commutator pairing is surjective. Q_8 cannot be decomposed (Q_8 is not a semi-direct product) and therefore does not fit into these general cases. Still, of course, the order of V_4 is a square.

The following consequence concludes that we have in fact found all the representations of a (finite) Heisenberg group with the Stone-von Neumann theorem.

Consequence 2.16. *The irreducible representations of H are:*

- *the $|V|$ characters χ of $V = H/Z(H)$ lifted to characters χ' of H (where $\chi'(h) = \chi(v)$ if v is a representative of the coset $hZ(H)$),*
- *the $|Z(H)| - 1$ irreducible representations $\pi_\rho = \mathrm{Ind}_B^H(\rho)$, where ρ is a non-trivial character of $Z(H)$.*

Proof. This is easily proven by counting. We know that

$$|H| = \sum_{\tau \in \mathrm{Irr}(H)} \dim \tau^2,$$

and if we count the irreducible representations we have found, we find that

$$\begin{aligned}\sum_{\tau\in\mathrm{Irr}(H)} \dim \tau^2 &= |V|\cdot(\dim\chi_v)^2 + (|Z(H)|-1)\cdot(\dim\pi_\rho)^2\\ &= |V|\cdot 1 + (|Z(H)|-1)\cdot|V|\\ &= |Z(H)|\cdot|V|\\ &= |H|.\end{aligned}$$

We see that we have indeed found all the irreducible representations of H. □

Remark. It is known that the characters of the irreducible representations of Q_8 and $D_4 \cong H(\mathbb{F}_2)$ (from Example 2.11) are the same. This is now clear, because both groups are Heisenberg groups with $Q_8/Z(Q_8) \cong V_4 \cong D_4/Z(D_4)$ and $Z(Q_8) \cong Z(D_4)$, which implies that the characters will be the same by the previous consequence.

Chapter 3

The Fast Fourier Transform

In the 1950s and '60s, the Discrete Fourier Transform (DFT) was used in a lot of calculations, ranging from spectral analysis and data compression to polynomial multiplication and even multiplying large integers. However, the speed of the DFT was low, as it required $\mathcal{O}(N^2)$ operations, where N is the number of points on which the DFT acts. A solution came from J. W. Cooley and J. W. Tukey in 1965, when they published an algorithm that could perform the DFT on N points in only $\mathcal{O}(N \log N)$ operations. It was named the Cooley-Tukey algorithm and is still the most commonly used algorithm to compute the DFT. This faster Fourier Transform was named the Fast Fourier Transform (FFT). Later it was discovered that the Cooley-Tukey algorithm was not new, and that it had been discovered by Carl Friedrich Gauss as early as 1805, but was never widely recognized.

In this chapter we will explain the abstract mathematics behind the algorithm, which is essentially an application of Heisenberg groups and representation theory. This was discovered by Auslander and Tolimieri in [1].

We will first explain the remarkable connection between a Heisenberg group over some finite ring and the DFT. Furthermore, the application of the Fourier Transform in the popular music identification application Shazam, that runs on iOS and Android. We then show a construction of the FFT, and compute the speed of the FFT.

3.1 A Model of a Heisenberg Representation

In this section we will introduce a model of an irreducible representation of the Heisenberg group over the finite ring $\mathbb{Z}/N\mathbb{Z}$, where N is the number of points on which the DFT will act. We start with $N = p^2$, where p is prime and odd and will later generalize this to $N = p^k$, where p is prime and odd and k a positive integer. We will use the notation $\zeta := e^{\frac{2\pi i}{N}}$ for the rest of the chapter.

Definition 3.1. The Heisenberg group over the finite ring $\mathbb{Z}/N\mathbb{Z}$ is the group $H(\mathbb{Z}/N\mathbb{Z}) := V \times Z$, where V is the symplectic vector space $\mathbb{Z}/N\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z}$ and $Z = \mathbb{Z}/N\mathbb{Z}$ is the center. V is a symplectic vector space with symplectic form

$$\varphi : V \times V \to \mathbb{Z}/N\mathbb{Z}$$

$$\varphi\left(\begin{pmatrix} a_1 \\ b_1 \end{pmatrix}, \begin{pmatrix} a_2 \\ b_2 \end{pmatrix}\right) = a_1 b_2 - b_1 a_2.$$

The group law on $H(\mathbb{Z}/N\mathbb{Z})$ is

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

Remark. There is an important distinction between this Heisenberg group and the Heisenberg groups as defined in chapter 2. We see for example that $\varphi(p \cdot v, -)$ is *not* surjective for some $v \in V$ and therefore the commutator pairing $e(p \cdot v, -)$ is not surjective. We find that $H(\mathbb{Z}/N\mathbb{Z})$ is not a proper Heisenberg group! This means that our proof of the Stone-von Neumann theorem would not work for every character ψ of the center $\mathbb{Z}/N\mathbb{Z}$, because the conjugacy classes of $H(\mathbb{Z}/N\mathbb{Z})$ are different from an actual Heisenberg group. However, if we choose $\psi : n \mapsto \zeta^n$ as our character, we *can* construct an irreducible representation π with central character ψ, because for this specific case our proof of the Stone-von Neumann theorem still holds. Because of this reason, we will refer to $H(\mathbb{Z}/N\mathbb{Z})$ as a Heisenberg group, although technically this is incorrect.

We shall now realize this representation π. To do this, we need to introduce the following space and functions on this space.

Let $\mathcal{H}$ be the space of functions $f : \mathbb{Z}/N\mathbb{Z} \to \mathbb{C}$.

Definition 3.2. For every $\tau \in \mathbb{Z}/N\mathbb{Z}$, we define the translating function $L_\tau : \mathcal{H} \to \mathcal{H}$ by

$$(L_\tau f)(n) = f(n + \tau).$$

For every $\omega \in \mathbb{Z}/N\mathbb{Z}$, we define the rotating function $M_\omega : \mathcal{H} \to \mathcal{H}$ by

$$(M_\omega f)(n) = \zeta^{\omega n} f(n).$$

For physicists, the following relation between L_τ and M_ω may be enlightening, as it is the *Heisenberg Commutation Relation*.

Lemma 3.3. *We have the relation*

$$M_\omega \circ L_\tau = \zeta^{-\omega\tau} \cdot L_\tau \circ M_\omega$$

for all $\tau, \omega \in \mathbb{Z}/N\mathbb{Z}$.

Proof. Just apply the left hand side and the right hand side to a function $f \in \mathcal{H}$:

$$(M_\omega \circ L_\tau(f))(n) = M_\omega(f)(n+\tau) = \zeta^{\omega n} f(n+\tau)$$

and if we write $h(n) = \zeta^{\omega n} f(n)$:

$$(L_\tau \circ M_\omega(f))(n) = L_\tau(h)(n) = h(n+\tau) = \zeta^{\omega(n+\tau)} f(n+\tau).$$

□

We can now realize π as

$$\pi : H(\mathbb{Z}/N\mathbb{Z}) \to GL(\mathcal{H})$$

$$\pi((\tau,\omega),z) = \zeta^{(\frac{1}{2}\tau\omega+z)} \cdot M_\omega \circ L_\tau \qquad \tau,\omega,z \in \mathbb{Z}/N\mathbb{Z}$$

and we can show that this is indeed a representation.

Lemma 3.4. *The function* π *is a representation.*

Proof. We will check that π is a homomorphism. We know that

$$(\tau,\omega,z)\cdot(\tau',\omega',z') = (\tau+\tau',\omega+\omega',z+z'+\frac{1}{2}(\tau\omega'-\tau'\omega)).$$

Now we apply π to both sides and rewrite using Lemma 3.3

$$\begin{aligned}
\pi(\tau,\omega,z)\cdot\pi(\tau',\omega',z') &= \zeta^{\frac{1}{2}\tau\omega+z} \cdot M_\omega \circ L_\tau \circ \zeta^{\frac{1}{2}\tau'\omega'+z'} \cdot M_{\omega'} \circ L_{\tau'} \\
&= \zeta^{\frac{1}{2}\tau\omega+\frac{1}{2}\tau'\omega'+z+z'} \cdot M_\omega \circ (L_\tau \circ M'_\omega) \circ L_{\tau'} \\
&= \zeta^{\frac{1}{2}\tau\omega+\frac{1}{2}\tau'\omega'+z+z'} \cdot \zeta^{\tau\omega'} M_\omega \circ M'_\omega \circ L_\tau \circ L_{\tau'} \\
&= \zeta^{\frac{1}{2}\tau\omega+\frac{1}{2}\tau'\omega'+\tau\omega'+z+z'} \cdot M_{\omega+\omega'} \circ L_{\tau+\tau'} \\
&= \pi(\tau+\tau',\omega+\omega',z+z'+\frac{1}{2}(\tau\omega'-\tau'\omega))
\end{aligned}$$

□

We find that π indeed acts on the center as ψ. To introduce the DFT, we will need to define two different, although quite similar, representations of $H(\mathbb{Z}/N\mathbb{Z})$. To realize these representations, we introduce the spaces $\mathcal{H}_T$, of time domain functions, and $\mathcal{H}_W$, of wave (or frequency) domain functions, which will be isomorphic to $\mathcal{H}$.

Definition 3.5. Let $T \subset V$ be the subspace $\{(t,0) \in V \mid t \in \mathbb{Z}/N\mathbb{Z}\}$. We define $\mathcal{H}_T$ as the space of functions $f : T \to \mathbb{C}$. Then we have the following representation π_T of $H(\mathbb{Z}/N\mathbb{Z})$:

$$\pi_T : H(\mathbb{Z}/N\mathbb{Z}) \to GL(\mathcal{H}_T)$$

$$\pi_T((\tau,\omega),z) = \zeta^{(\frac{1}{2}\tau\omega+z)} \cdot M_\omega \circ L_\tau \qquad \tau,\omega,z \in \mathbb{Z}/N\mathbb{Z},$$

which is a representation because it is isomorphic to π from Lemma 3.4.

Let $W \subset V$ be the subspace $\{(0,w) \in V \mid w \in \mathbb{Z}/N\mathbb{Z}\}$. We define $\mathcal{H}_W$ as the space of functions $f : W \to \mathbb{C}$. Then we have the following representation π_W of $H(\mathbb{Z}/N\mathbb{Z})$:

$$\pi_W : H(\mathbb{Z}/N\mathbb{Z}) \to GL(\mathcal{H}_W)$$

$$\pi_W((\tau,\omega),z) = \zeta^{(-\frac{1}{2}\tau\omega+z)} \cdot M_{-\tau} \circ L_\omega \qquad \tau,\omega,z \in \mathbb{Z}/N\mathbb{Z},$$

which is a representation using Lemma 3.4 (and the mapping $-\tau \mapsto \omega, \omega \mapsto \tau$).

Notice that π_T and π_W both act on the center as ψ. By the Stone-von Neumann theorem, they must therefore be isomorphic, e.g. there exists an intertwiner between these representations. We are now ready to realize the DFT.

3.2 The Discrete Fourier Transform

Definition 3.6. The DFT on N points is the function

$$\mathrm{DFT} : \mathcal{H}_T \to \mathcal{H}_W$$

$$f(n) \mapsto \frac{1}{\sqrt{N}} \sum_{x \in \mathbb{Z}/N\mathbb{Z}} f(x) \cdot \zeta^{nx},$$

where we identify functions in $\mathcal{H}_T$ and $\mathcal{H}_W$ as functions from $\mathcal{H}$, because we know that $T \cong \mathbb{Z}/N\mathbb{Z} \cong W$.

3.2.1 Application of the Fourier Transform: Shazam

In this subsection we will explain one of the applications of the Fourier Transform that is used in the popular music identification application *Shazam*.

Shazam is used in the following way: imagine yourself standing in a café or crowded place with a song playing. You recognize the song but do not know the name of the artist or the name of the song. You tap a button in the Shazam app, Shazam records the sound for a couple of seconds and tells you what song is playing.

We will not explain in detail all the exact algorithms used by Shazam, but will focus on the use of the Fourier Transform. First of all, Shazam records a small part of the song. We can then look at the signal.

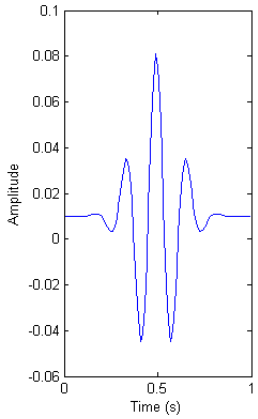

Figure 3.1: A time domain signal

We call this a time domain signal. In reality, this will not often be a nice as this signal, as there is more noise in the recording. If we now apply the Fourier Transform on N points (this is usually a large N, say $N = 1024$), this will transform this time domain signal into a frequency domain signal. For example, the previous signal will turn into the following frequency domain signal.

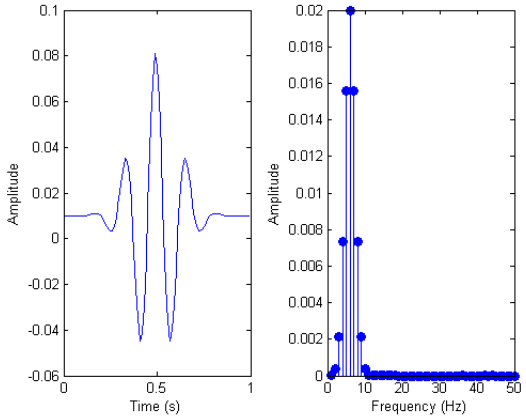

Left: The time domain signal. Right: The frequency domain signal

This gives us a *signature* of our signal. We see a large peak around 8 Hz, which makes our signature (8). For more complex signals we will get more peaks and we get a signature such as $(a_1, ..., a_k)$, if there is a peak at a_i Hz. Even if the recording is noisy, these peaks will appear, which is why Shazam works in crowded places.

For a long recording, this would heap up all the frequencies in that recording into one signature. This would be inconvenient, if we want to identify the song, as we lose information about the time signal. Therefore, Shazam splits up a recording in smaller parts and creates a row of signatures. All of these signatures combined makes a *fingerprint* of your recording. Shazam has a huge library of these fingerprint for all songs in its database and compares you fingerprint to the fingerprints in the database. It will then show you the songs that have the most matching fingerprints to your recording.

The fact that Shazam works so quickly is partly because of their fast algorithms to search their database for identical fingerprints, but also because it relies on the Fast Fourier Transform, instead of the Discrete Fourier Transform. Because it has to calculate the Fourier Transform so often, the computational difference is significant.

It may not be clear why the Fourier Transform is connected to our Heisenberg group from the first section. We will explain this in the following subsection.

3.2.2 The connection between the DFT and Heisenberg groups

The following lemma is the crucial connection between the DFT on N points and the Heisenberg group $H(\mathbb{Z}/N\mathbb{Z})$.

Lemma 3.7. *The* DFT *is an intertwiner between* π_T *and* π_W*.*

Proof. We will need to prove for all $h \in H(\mathbb{Z}/N\mathbb{Z})$ that the following holds:

$$\mathrm{DFT} \circ \pi_T(h) = \pi_W(h) \circ \mathrm{DFT}\,.$$

To do this, we will apply both sides to $f(n)$.

First we calculate $\pi_T(h)(f)$,

$$\pi_T(h)(f)(n) = \zeta^{(\frac{1}{2}\tau\omega+z)} \cdot M_\omega \circ L_\tau(f)(n) = \zeta^{(\frac{1}{2}\tau\omega+z)} \cdot \zeta^{\omega n} f(n+\tau).$$

Then we apply the DFT,

$$\mathrm{DFT} \circ \pi_T(h)(f)(n) = \frac{1}{\sqrt{N}} \zeta^{(\frac{1}{2}\tau\omega+z)} \sum_{x \in \mathbb{Z}/N\mathbb{Z}} \zeta^{nx} \cdot \zeta^{\omega x} f(x+\tau).$$

As we sum x over $\mathbb{Z}/N\mathbb{Z}$, we can substitute $x \mapsto x - \tau$, to get

$$\frac{1}{\sqrt{N}} \zeta^{(\frac{1}{2}\tau\omega+z)} \sum_{x \in \mathbb{Z}/N\mathbb{Z}} \zeta^{n(x-\tau)} \cdot \zeta^{\omega(x-\tau)} f(x),$$

or rearranged

$$\frac{1}{\sqrt{N}} \zeta^{(-\frac{1}{2}\tau\omega+z)} \cdot \zeta^{-\tau n} \sum_{x \in \mathbb{Z}/N\mathbb{Z}} \zeta^{(n+\omega)x} f(x).$$

Now we calculate $\pi_W(h) \circ \mathrm{DFT}(f)(n)$. For clarity, write $g := \mathrm{DFT}(f)$. Then

$\pi_W(h)(g)(n) = \zeta^{(-\frac{1}{2}\tau\omega+z)} \cdot M_{-\tau} \circ L_\omega(g)(n) = \zeta^{(-\frac{1}{2}\tau\omega+z)} \cdot \zeta^{-\tau n} g(n+\omega)$.

Putting this together, we find

$$\pi_W(h) \circ \mathrm{DFT}(f)(n) = \frac{1}{\sqrt{N}} \zeta^{(-\frac{1}{2}\tau\omega+z)} \cdot \zeta^{-\tau n} \sum_{x \in \mathbb{Z}/N\mathbb{Z}} \zeta^{(n+\omega)x} f(x).$$

□

The problem with the DFT is that we need to execute N^2 calculations, because for every $n \in \{0, ..., N-1\}$ we sum N values. In the following section we will show that some of these calculations are unnecessary and that the speed can be improved.

3.3 The Fast Fourier Transform

In this section we will introduce new representations of $H(\mathbb{Z}/N\mathbb{Z})$. Then we find intertwiners between these new representations and the representations of the previous section. After that we calculate the speed of these intertwiners composed and we will find that we have increased the speed of the calculations. This transform will be called the Fast Fourier Transform (FFT).

3.3.1 Constructing the FFT

We will first introduce two new spaces that will be isomorphic to the previous spaces $\mathcal{H}_T$ and $\mathcal{H}_W$.

Definition 3.8. Let $\mathcal{H}^W$ be the space of W-invariant functions from H to $\mathbb{C}$. This means that for all $f \in \mathcal{H}^W$ we have:

- $f : H \to \mathbb{C}$
- $f(w \cdot h) = f(h) \qquad \forall\ w \in W,\ h \in H$
- $f(z \cdot h) = \zeta^z \cdot f(h) \qquad \forall\ z \in Z,\ h \in H$

Analogously, we define $\mathcal{H}^T$ as the space of T-invariant functions, so again:

- $f : H \to \mathbb{C}$
- $f(t \cdot h) = f(h) \qquad \forall\ t \in T,\ h \in H$
- $f(z \cdot h) = \zeta^z \cdot f(h) \qquad \forall\ z \in Z,\ h \in H$

We define a representation $\pi^W : H \to GL(\mathcal{H}^W)$ by

$$\pi^W(h)[f(h')] = f(h' \cdot h) \qquad \forall\ h, h' \in H$$

and we define a representation $\pi^T : H \to GL(\mathcal{H}^T)$ by

$$\pi^T(h)[f(h')] = f(h' \cdot h) \qquad \forall\ h, h' \in H$$

We see that $\mathcal{H}_T \cong \mathcal{H}^W$ by $f \mapsto \tilde{f}(wtz) = \zeta^z \cdot f(t)$ and that $\mathcal{H}_w \cong \mathcal{H}^T$ by $f \mapsto \tilde{f}(twz) = \zeta^z \cdot f(w)$. Notice also that $\pi^W(z) = \zeta^z \cdot \mathrm{Id}_{\mathcal{H}^W}$ and $\pi^T(z) = \zeta^z \cdot \mathrm{Id}_{\mathcal{H}^T}$, which implies that they are isomorphic using the Stone-von Neumann theorem.

Now we shall define a new Lagrangian subspace and a representation onto the linear transforms of that space. Recall that $N = p^2$ and $V = \mathbb{Z}/N\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z}$.

Definition 3.9. Let $\Lambda = \{p \cdot v \mid v \in V\} \subset V$. We see Λ as a subgroup of H by $\lambda \to (\lambda, 0)$. For $(\lambda, 0) \in H$, we will write $\lambda \in H$. We can define $\mathcal{H}^\Lambda$ as the space of Λ-invariant functions $f : H \to \mathbb{C}$. This means that $f \in \mathcal{H}^\Lambda$ if:

- $f : H \to \mathbb{C}$
- $f(\lambda \cdot h) = f(h) \qquad \forall\ \lambda \in \Lambda,\ h \in H$
- $f(z \cdot h) = \zeta^z \cdot f(h) \qquad \forall z \in Z,\ h \in H$

Again, we find the representation $\pi^\Lambda : H \to GL(\mathcal{H}^\Lambda)$

$$\pi^\Lambda(h)[f(h')] = f(h' \cdot h) \qquad \forall\ h, h' \in H$$

and again we notice that

$$\pi^\Lambda(z) = \zeta^z \cdot Id_{\mathcal{H}^\Lambda}.$$

We see that all three representations π^W, π^T and π^Λ act the same on the center, therefore the Stone-von Neumann theorem tells us that all three representations are isomorphic, therefore there exist intertwiners between their vector spaces. We shall call these intertwiners $\mathcal{F}^\Lambda_W$ and $\mathcal{F}^T_\Lambda$, respectively for the intertwiners $\mathcal{H}^W \to \mathcal{H}^\Lambda$ and $\mathcal{H}^\Lambda \to \mathcal{H}^T$. Figure 3.2 shows us a diagram of the construction so far.

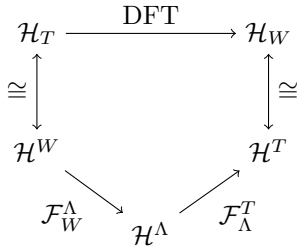

Figure 3.2: A diagram of the construction for $N = p^2$.

We can give explicit formulas for these intertwiners $\mathcal{F}^\Lambda_W$ and $\mathcal{F}^T_\Lambda$, because we can just average over Λ and T respectively. This gives us

$$(\mathcal{F}^\Lambda_W f)(h) = \sum_{\lambda \in \Lambda} f(\lambda h)$$

and

$$(\mathcal{F}_\Lambda^T f)(h) = \sum_{t\in T} f(th).$$

It now follows that $\mathcal{F}_\Lambda^T \circ \mathcal{F}_W^\Lambda$ is an intertwiner between π^W and π^T just as the DFT. By Schur's Lemma, they are therefore equal for some scalar c which is equal to the coefficient $\frac{1}{\sqrt{N}}$ of the DFT, or,

$$\text{DFT} = \frac{1}{\sqrt{N}} \cdot \mathcal{F}_\Lambda^T \circ \mathcal{F}_W^\Lambda.$$

3.3.2 Comparing Complexities

We have seen that we can find another expression for the DFT, however there is yet no reason to assume that it is quicker. But we know that for $f \in \mathcal{H}^W$ we have $f(w\lambda) = f(\lambda)$ and therefore

$$\sum_{\lambda\in\Lambda} f(\lambda h) = |\Lambda \cap W| \cdot \sum_{\lambda\in\Lambda/[\Lambda\cap W]} f(\lambda h)$$

with $\Lambda/[\Lambda \cap W] = \{(0,0),\ p(1,0),\ p(2,0), ...,\ p(p-1,0)\}$. This gives us the difference we need, because we only need to sum $|\Lambda/[\Lambda \cap W]| = p$ terms to calculate $(\mathcal{F}_W^\Lambda f)(h)$, instead of $|\Lambda| = p^2$. Equivalently, we only need to sum p terms to calculate $(\mathcal{F}_\Lambda^T f)(h)$.

The speed of $c \cdot \mathcal{F}_\Lambda^T \circ \mathcal{F}_W^\Lambda$ is therefore

$$\underbrace{\underbrace{\underbrace{p^2}_{h\text{-values}} \cdot \underbrace{p}_{\text{sums}}}_{\mathcal{F}_W^\Lambda} + \underbrace{\underbrace{p^2}_{h\text{-values}} \cdot \underbrace{p}_{\text{sums}}}_{\mathcal{F}_\Lambda^T}}_{} = \underbrace{p}_{\text{constant}} \cdot \underbrace{p^2}_{N} \cdot \underbrace{2}_{\log(N)}$$

In the case where $N = p^2$ it is not obvious why p may be viewed as a constant, but this will follow from the results of the following section.

3.4 Generalisation of the FFT

In this section we generalize the FFT to $N = p^k$ for some odd prime p and $k > 1$. The construction of the DFT does not change. However, to increase the speed of the FFT we will have to find *more* Lagrangian subspaces. On these Lagrangian subspaces we will again define representations. Then we repeat the procedure of composing the intertwiners, which will give us the FFT for $N = p^k$. Let us therefore start by defining subspaces, prove that they are Lagrangian and construct the FFT. Afterwards we will compute the speed of this transform.

Definition 3.10. Let $N = p^k$ for some odd prime p and $k > 1$. We define the following $k + 1$ subspaces of $V = \mathbb{Z}/N\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z}$,

$$\Lambda_i := \{p^{k-i} \cdot a, p^i \cdot b) \mid a, b \in \mathbb{Z}/N\mathbb{Z}\}.$$

Lemma 3.11. *For every i, the subspace Λ_i is Lagrangian.*

Proof. We need to show that the symplectic form φ is zero on $\Lambda_i \times \Lambda_i$:

$$\begin{aligned} \forall\ \lambda, \mu \in \Lambda_i \quad \varphi(\lambda, \mu) &= p^{k-i}\lambda_1 \cdot p^i \mu_2 - p^i \lambda_2 \cdot p^{k-i}\mu_1 \\ &= p^k(\lambda_1\mu_2 - \lambda_2\mu_1) \\ &= 0. \end{aligned}$$

□

There is again overlap between the Lagrangian subspaces Λ_i and Λ_{i+1}. To see this more clearly, we write

$$\Lambda_i = \{(p^{k-i} \cdot a, p^i \cdot b \mid a \in \{0, ..., p^i\}, b \in \{0, ..., p^{k-i}\}\}$$

and

$$\Lambda_{i+1} = \{(p^{k-i-1} \cdot a, p^{i+1} \cdot b \mid a \in \{0, ..., p^{i+1}\}, b \in \{0, ..., p^{k-i-1}\}\}.$$

We now see that

$$|\Lambda_i \cap \Lambda_{i+1}| = |\{(a, b) \mid a \in \{0, ..., p^i\}, b \in \{0, ..., p^{k-i-1}\}\}| = p^{k-1}.$$

Notice also that $\Lambda_0 = W$ and $\Lambda_k = W$. In the same way as before, we can define the space $\mathcal{H}^{\Lambda_i}$, find the representation π^{Λ_i} and find the intertwiner $\mathcal{F}^{\Lambda_{i+1}}_{\Lambda_i}$. This gives us the following diagram:

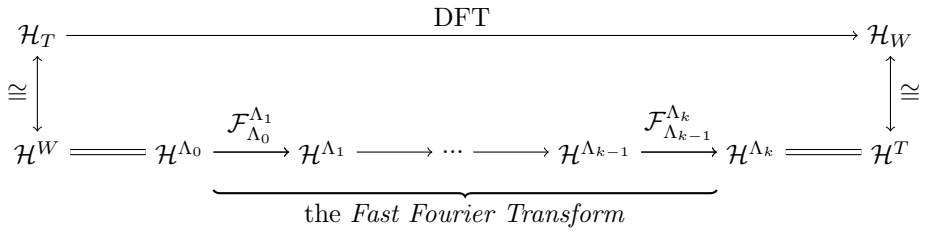

Figure 3.3: A diagram of the construction for $N = p^k$

Again we see that the composition of all these $\mathcal{F}^{\Lambda_i}_{\Lambda_{i-1}}$ is an intertwiner between $\mathcal{H}^W$ and $\mathcal{H}^T$. Therefore, for some $c \in \mathbb{C}$, this must be equal to the DFT. Again, this is the scalar $\frac{1}{\sqrt{N}}$.

Definition 3.12. The composition of all $\mathcal{F}^{\Lambda_i}_{\Lambda_{i-1}}$ multiplied with $c = \frac{1}{\sqrt{p^k}}$ is called the *Fast Fourier Transform* and will be denoted by $\mathcal{F}$, or,

$$\mathcal{F} := \frac{1}{\sqrt{p^k}} \cdot \mathcal{F}^{\Lambda_k}_{\Lambda_{k-1}} \circ ... \circ \mathcal{F}^{\Lambda_1}_{\Lambda_0}.$$

We see again why this is faster, because we do not have to sum every p^k terms of $\mathcal{F}^{\Lambda_i}_{\Lambda_{i-1}}$, but only $|\Lambda_i/(\Lambda_i \cap \Lambda_{i+1})| = p^k/p^{k-1} = p$ terms. Therefore we can now calculate its speed:

$$\underbrace{\underbrace{\underbrace{p^k}_{h\text{-values}} \cdot \underbrace{p}_{\text{sums}}}_{\mathcal{F}^{\Lambda_1}_{\Lambda_0}} + \ldots + \underbrace{\underbrace{p^k}_{h\text{-values}} \cdot \underbrace{p}_{\text{sums}}}_{\mathcal{F}^{\Lambda_k}_{\Lambda_{k-1}}}}_{k \text{ times}} = \underbrace{p}_{\text{constant}} \cdot \underbrace{p^k}_{N} \cdot \underbrace{k}_{\log(N)}$$

We see that we have constructed the Fast Fourier Transform on p^k points, where p is an odd prime, in only $\mathcal{O}(N \log N)$ operations instead of $\mathcal{O}(N^2)$ operations. This is the essence of the Cooley-Tukey algorithm and in some sense the abstract version of the calculations performed by Gauss in 1805.

Chapter 4

Mutually Unbiased Bases and Gauss sums

In this chapter we will look at yet another application of Heisenberg groups using representation theory. We will explain what *Mutually Unbiased Bases* (MUBs) are and why we try to find them. Furthermore, we will look at a Gauss sum,

$$\sum_{x \in \mathbb{F}_p} e^{\frac{2\pi i}{p} x^2},$$

and use some of the results we obtained in our search for MUBs to explain known facts of this Gauss sum. We conclude the chapter with a discussion on our results and how they can be improved.

4.1 Mutually Unbiased Bases

We introduce mutually unbiased bases of the complex vector space $\mathbb{C}^n$, with the standard inner product.

Definition 4.1. Let $\mathcal{B}_1, ..., \mathcal{B}_k$ be k orthonormal bases for $\mathbb{C}^n$. We call these bases *mutually unbiased* if they obey

$$\forall v \in \mathcal{B}_i, w \in \mathcal{B}_j, i \neq j \qquad |\langle v, w \rangle|^2 = \frac{1}{n}.$$

We will refer to these bases as *MUBs*.

MUBs are used in quantum information theory, which is a theory that tries to generalize classical information theory to quantum information. To do this, one needs to determine quantum states. For example as said in the so-called *'The Mean King's Problem'* (see [2]):

"A mean king challenges a physicist, who got stranded on the remote island ruled by the king, to prepare a spin-$\frac{1}{2}$ atom in any state of her choosing and to perform a control measurement of her liking. Between her preparation and her measurement, the kings men determine the value of either σ_x, or σ_y, or σ_z. Only after she completed her control measurement, the physicist is told which spin component has been measured, and she must then state the result of that intermediate measurement correctly. How does she do it?"

Here MUBs are used because the outcome of a measurement in one basis will be random if the state is prepared in another basis *if* the two bases are unbiased. Therefore it would be useful to have many mutually unbiased bases in higher dimensional states.

Furthermore, MUBs are used in quantum cryptography, in [5] for example, in the field of secure quantum key exchange. Until now, quantum cryptography protocols rely on 2-dimensional quantum states, but security against eavesdropping has been found to increase significantly if higher dimensional quantum states are used, which requires higher dimensional MUBs, see also [6].

One of the problems on MUBs is their existence. An open question about MUBs is:

"For $n \in \mathbb{N}$, what is the largest integer k such that there are MUBs $\mathcal{B}_1, ..., \mathcal{B}_k$ for $\mathbb{C}^n$?"

In this chapter we will prove and find $p + 1$ MUBs for $\mathbb{C}^p$ with p odd and prime. It is known that one can find $p^k + 1$ MUBs for $\mathbb{C}^{p^k}$. However, if the dimension is not a prime power ($n = 6, 10, 12, ...$), it is unknown what the largest number of MUBs is. For example, for $n = 6$, the techniques used so far can only find 3 MUBs (as does the technique explained in this chapter). The general consensus therefore seems to be that 3 MUBs is indeed the maximum number for $n = 6$, but this is of course not proved or disproved and the available evidence

is rather weak. We do however know an upper bound for all dimensions by the following lemma.

Lemma 4.2. *The largest integer k such that there are MUBs $\mathcal{B}_1, ..., \mathcal{B}_k$ for $\mathbb{C}^n$ is less than or equal to $n+1$.*

Remark. A proof of this lemma can be found in [8]. It is not difficult, but rather long and beyond the focus of this text.

It follows that we will have answered the question in the case where $n = p$ with p prime and odd, namely, we can find $p+1$ MUBs. Let us therefore continue our search for MUBs. To do this, we will use a specific representation of a Heisenberg group, which will help us in proving the existence of these MUBs.

4.1.1 A Model of a Representation of a Heisenberg Group

The Heisenberg group we will use in this section is $H = V \times Z$ with $V = \mathbb{F}_p \times \mathbb{F}_p$, a symplectic vector space with symplectic form ω, and the center $Z = \mathbb{F}_p$, where p is odd and prime. The group law is

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

The 2-dimensional vector space V has 1-dimensional subspaces which we call *lines*, as we can see V as a finite plane. If we pick 2 different lines L and M, we can write $V = L \bigoplus M$. The vector $v \in V$ will sometimes be written as $v = l_v + m_v$ where $l_v \in L$ and $m_v \in M$ are the projections on the lines L and M respectively.

Remark. For the rest of this section we will simplify our notation for the elements $((l,0),0) \in H$ as $'l \in H'$ and $((m,0),0) \in H$ as $'m \in H'$. Furthermore, we will write $z \in H$ for $(0,z) \in H$, where $z \in Z(H)$.

We know by the Stone-von Neumann theorem (Theorem 2.15) that the only irreducible representations, excluding characters, are the representations π with a central non-trivial character ψ. The representation will work on the function space $\mathcal{H}$ of functions $f : L \to \mathbb{C}$, which is equipped with the inner product

$$\langle f, g \rangle = \sum_{l \in L} f(l)\overline{g(l)}.$$

The representation $(\pi, \mathcal{H})$ is

$$\begin{aligned}
\pi &: H \to \mathrm{GL}(\mathcal{H}), \\
\pi(l')f(l) &= f(l+l') && \forall l' \in L, \\
\pi(m)f(l) &= \psi \circ \omega \left(\begin{pmatrix} 0 \\ m \end{pmatrix}, \begin{pmatrix} l \\ 0 \end{pmatrix} \right) f(l) && \forall m \in M, \\
\pi(z)f(l) &= \psi(z)f(l) && \forall z \in Z.
\end{aligned}$$

Notice that we have not yet picked our lines L and M, nor have we picked a specific non-trivial central character ψ. This is done on purpose as we will try to keep this approach as general as possible. Only after we have found the formulas for our MUBs will we pick our lines and character, which will give us specific bases. We are now ready to prove the existence of $p+1$ MUBs in the following subsection.

4.1.2 Existence of MUBs in $\mathbb{C}^p$

A quick summary of what will follow will improve the clarity on the whole. We will demonstrate the approach for p prime and odd, although the same approach works for prime powers of p prime and odd. We know that there are $p+1$ lines in the finite plane $\mathbb{F}_p \times \mathbb{F}_p$, see for an easy example Figure 4.1 where we take $p = 5$.

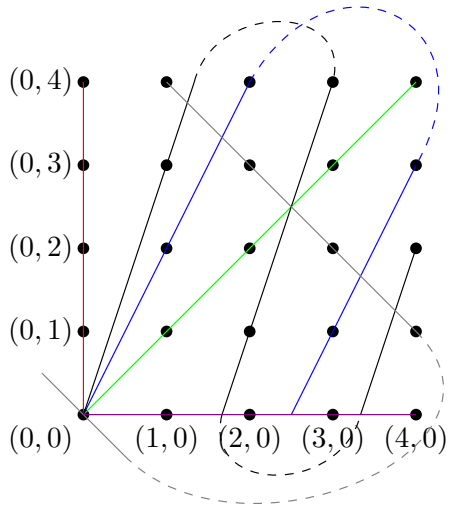

Figure 4.1: The finite plane $\mathbb{F}_5 \times \mathbb{F}_5$ and its 6 lines

The idea is that every line in $V = \mathbb{F}_p \times \mathbb{F}_p$ will give us one orthonormal basis, using the representation π we defined on our Heisenberg group $H = V \times \mathbb{F}_p$. Surprisingly, it will follow that these bases are mutually unbiased. This gives us also an idea why this approach will not work to find MUBs for $\mathbb{C}^n$ if n is not a prime power.

To begin, let us take a line $N \subset V$. Now look at the following decomposition of $\mathcal{H}$, where the $\psi_N : N \to \mathbb{C}^*$ are the p characters of N (using Lemma 1.11 and the fact that $N \cong \mathbb{F}_p$):

$$\mathcal{H} = \bigoplus_{\psi_N} \mathcal{H}_{\psi_N}.$$

Here $\mathcal{H}_{\psi_N}$ is the subspace we get by the projection P_{ψ_N}, associated to ψ_N,

defined by

$$P_{\psi_N} = \frac{1}{p} \sum_{n \in N} \psi_N(-n)\pi(n).$$

Now take functions f_{ψ_N} (of norm 1) that span the subspaces $\mathcal{H}_{\psi_N}$. These will form an orthonormal base for $\mathcal{H}$ if they are eigenfunctions, e.g.

$$\pi(n) f_{\psi_N}(l) = \psi_N(n) f_{\psi_N}(l).$$

We claim that for every line $N \subset V$, these functions f_{ψ_N} form an orthonormal basis of $\mathcal{H}$ and that the bases thus obtained will be mutually unbiased. To prove the claim we divide the claim into the following steps

Step 1. The dimension of $\mathcal{H}_{\psi_N}$ is 1 for every character ψ_N.

Step 2. For every pair of different lines N and N' and for every character ψ_N and $\psi_{N'}$ of N and N' respectively, we have

$$|\langle f_{\psi_N}, f_{\psi_{N'}} \rangle|^2 = \frac{1}{p}.$$

4.1.3 The steps

We first need to prove that P_{ψ_N} is indeed a projection.

Lemma 4.3. *The function $P_{\psi_N} : \mathcal{H} \to \mathcal{H}$ given by*

$$P_{\psi_N} = \frac{1}{p} \sum_{n \in N} \psi_N(-n)\pi(n)$$

is a projection.

Proof. We will check that $P_{\psi_N} \circ P_{\psi_N} = P_{\psi_N}$ by computation.

$$P_{\psi_N} \circ P_{\psi_N} = \frac{1}{p} \sum_{n \in N} \frac{1}{p} \sum_{n' \in N} \psi_N(-(n+n'))\pi(n+n')$$

For every $n \in N$, we can substitute $n' \mapsto n' - n$, because we sum over $N \cong \mathbb{F}_p$, to get

$$P_{\psi_N} \circ P_{\psi_N} = \frac{1}{p} \sum_{n \in N} \frac{1}{p} \sum_{n' \in N} \psi_N(-n')\pi(n').$$

This can now be reduced to

$$P_{\psi_N} \circ P_{\psi_N} = \frac{1}{p} \sum_{n \in N} P_{\psi_N} = \frac{1}{p} \cdot |N| \cdot P_{\psi_N} = P_{\psi_N}.$$

□

We are now ready to prove step 1.

Lemma 4.4 (Step 1). *The dimension of $\mathcal{H}_{\psi_N}$ is 1 for every character ψ_N.*

Proof. We know that the trace of a projection is equal to the dimension of the target space. Therefore we compute

$$\begin{aligned}\operatorname{Tr}(P_{\psi_N}) &= \operatorname{Tr}\left(\frac{1}{p}\sum_{n\in N}\psi_N(-n)\pi(n)\right)\\ &= \frac{1}{p}\sum_{n\in N}\psi_N(-n)\operatorname{Tr}(\pi(n))\\ &= \frac{1}{p}\sum_{n\in N}\psi_N(-n)\chi_\pi(n).\end{aligned}$$

Now, for every non-zero $n \in N$, we get $\chi_\pi(n) = 0$ as we saw in the proof of the Stone-von Neumann theorem (Theorem 2.15). If $n = 0$, then $\chi_\pi(0) = p$ (as this is the dimension of $\mathcal{H}$) and $\psi_N(-0) = 1$. We can therefore conclude that

$$\operatorname{Tr}(P_{\psi_N}) = 1,$$

which gives us that $\dim \mathcal{H}_{\psi_N} = 1$. □

The proof of step 2 is a bit harder and will be done in two substeps. First we prove that for two different lines N and N':

$$|\langle f_{\psi_N}, f_{\psi_{N'}}\rangle|^2 = \operatorname{Tr}(P_{\psi_N}P_{\psi_{N'}}),$$

and then step 2 will follow with:

$$\operatorname{Tr}(P_{\psi_N}P_{\psi_{N'}}) = \frac{1}{p}.$$

Lemma 4.5. *Let N and N' be two different lines in V. Let ψ_N and $\psi_{N'}$ be two characters of N and N' respectively. Let f_{ψ_N} and $f_{\psi_{N'}}$ be the normal eigenfunctions of $\mathcal{H}_{\psi_N}$ and $\mathcal{H}_{\psi_{N'}}$ respectively. Then*

$$|\langle f_{\psi_N}, f_{\psi_{N'}}\rangle|^2 = \operatorname{Tr}(P_{\psi_N}P_{\psi_{N'}}).$$

Proof. Because f_{ψ_N} is a function of norm 1 in $\mathcal{H}_{\psi_N}$, we can rewrite P_{ψ_N} as

$$P_{\psi_N} f = \langle f, f_{\psi_N}\rangle \cdot f_{\psi_N}$$

where $f \in \mathcal{H}$. Now

$$\begin{aligned}P_{\psi_N}P_{\psi_{N'}} f &= P_{\psi_N}(\langle f, f_{\psi_{N'}}\rangle \cdot f_{\psi_{N'}}),\\ &= \langle f, f_{\psi_{N'}}\rangle\langle f_{\psi_{N'}}, f_{\psi_N}\rangle \cdot f_{\psi_N}.\end{aligned}$$

If we now pick our orthonormal basis $\mathcal{B}_{N'}$ as our basis of $\mathcal{H}$, then we find that for every $f \in \mathcal{B}_{N'}$ with $f \neq f_{\psi_{N'}}$ that

$$\langle f, f_{\psi_{N'}}\rangle = 0.$$

The trace of a transformation can be seen as the sum of the diagonal of the matrix $[P_{\psi_N} P_{\psi_{N'}}]_{\mathcal{B}_{N'}}$. But since

$$P_{\psi_N} P_{\psi_{N'}}(f) = \langle f, f_{\psi_{N'}} \rangle \langle f_{\psi_{N'}}, f_{\psi_N} \rangle \cdot f_{\psi_N} = 0$$

for all $f \in \mathcal{B}_{N'}$ with $f \neq f_{\psi_{N'}}$, and

$$P_{\psi_N} P_{\psi_{N'}}(f_{\psi_{N'}}) = \langle f_{\psi_{N'}}, f_{\psi_{N'}} \rangle \langle f_{\psi_{N'}}, f_{\psi_N} \rangle \cdot f_{\psi_N} = \langle f_{\psi_{N'}}, f_{\psi_N} \rangle \cdot f_{\psi_N},$$

we get that $\mathrm{Tr}(P_{\psi_N} P_{\psi_{N'}})$ is equal to the coefficient of $P_{\psi_{N'}}(\langle f_{\psi_{N'}}, f_{\psi_N} \rangle \cdot f_{\psi_N})$, which we can easily calculate by

$$P_{\psi_{N'}}(\langle f_{\psi_{N'}}, f_{\psi_N} \rangle \cdot f_{\psi_N}) = \langle f_{\psi_{N'}}, f_{\psi_N} \rangle \langle f_{\psi_N}, f_{\psi_{N'}} \rangle \cdot f_{\psi_{N'}}.$$

Concluding, we get

$$\mathrm{Tr}(P_{\psi_N} P_{\psi_{N'}}) = \langle f_{\psi_{N'}}, f_{\psi_N} \rangle \langle f_{\psi_N}, f_{\psi_{N'}} \rangle = |\langle f_{\psi_N}, f_{\psi_{N'}} \rangle|^2.$$

□

The second substep of step 2 is now an easy computation.

Lemma 4.6. *Let N and N' be two different lines in V and let P_{ψ_N} and $P_{\psi_{N'}}$ be the projections onto $\mathcal{H}_{\psi_N}$ and $\mathcal{H}_{\psi_{N'}}$ respectively. Then*

$$\mathrm{Tr}(P_{\psi_N} P_{\psi_{N'}}) = \frac{1}{p}.$$

Proof. We will use the first definition of P_{ψ_N} again.

$$\begin{aligned} \mathrm{Tr}(P_{\psi_N} P_{\psi_{N'}}) &= \mathrm{Tr}\left(\frac{1}{p} \sum_{n \in N} \psi_N(-n)\pi(n) \frac{1}{p} \sum_{n' \in N'} \psi_{N'}(-n')\pi(n') \right) \\ &= \frac{1}{p^2} \sum_{n \in N} \sum_{n' \in N'} \psi_N(-n)\psi_{N'}(-n') \,\mathrm{Tr}(\pi(n)\pi(n')) \end{aligned}$$

Now, as N and N' are two different lines we have that $N \cap N' = \{0\}$. So $n + n' = 0$ if and only if $n = 0 = n'$. But we know that $\mathrm{Tr}(\pi(n)\pi(n'))$ will only be non-zero if $n + n' = 0$ (as we have seen in the proof of step 1), and therefore the sum reduces to

$$\mathrm{Tr}(P_{\psi_N} P_{\psi_{N'}}) = \frac{1}{p^2} \cdot \psi_N(-0)\psi_{N'}(-0')\chi_\pi(0) = \frac{1}{p^2} \cdot p = \frac{1}{p},$$

which concludes the proof. □

Putting all these results together we see that we have found $p + 1$ MUBs, which we called $\mathcal{B}_N$ for a line N in V. However, we can get more out of this approach than just existence, we can actually get complete formulas for these eigenfunctions f_{ψ_N}, which is the goal of the following subsection.

4.2 Finding formulas for f_{ψ_N}

To find formulas for f_{ψ_N}, we will need to find a function that satisfies

$$\pi(n)f_{\psi_N} = \psi_N(n)f_{\psi_N}$$

for every $n \in N$. The solution comes quite *ad hoc* if we consider the delta zero function $\delta_0 \in \mathcal{H}$ defined by

$$\delta_0(l) = \begin{cases} 1 & \text{if } l = 0 \\ 0 & \text{if } l \neq 0 \end{cases}$$

and look at its projection on the different subspaces $\mathcal{H}_{\psi_N}$.

Lemma 4.7. *The function $P_{\psi_N}\delta_0$ satisfies*

$$\pi(n)P_{\psi_N}\delta_0 = \psi_N(n)P_{\psi_N}\delta_0$$

for every $n \in N$.

Proof. We can calculate both sides:

$$\pi(n)P_{\psi_N}\delta_0 = \pi(n)\cdot\frac{1}{p}\sum_{n'\in N}\psi_N(-n')\pi(n')\delta_0 = \frac{1}{p}\sum_{n'\in N}\psi_N(-n')\pi(n+n')\delta_0,$$

and

$$\psi_N(n)P_{\psi_N}\delta_0 = \psi_N(n)\cdot\frac{1}{p}\sum_{n'\in N}\psi_N(-n')\pi(n')\delta_0 = \frac{1}{p}\sum_{n'\in N}\psi_N(n-n')\pi(n')\delta_0.$$

Now substitute $n' \mapsto n + n'$ to get

$$\psi_N(n)P_{\psi_N}\delta_0 = \frac{1}{p}\sum_{n'\in N}\psi_N(n-n-n')\pi(n+n')\delta_0 = \pi(n)P_{\psi_N}\delta_0.$$

□

We see that we have found the formula $P_{\psi_N}\delta_0$ for our eigenfunctions. However, we need to find *normal* eigenfunctions. The norm of $P_{\psi_N}\delta_0$ is

$$\begin{aligned}\langle P_{\psi_N}\delta_0, P_{\psi_N}\delta_0\rangle &= \langle P_{\psi_N}\delta_0, \delta_0\rangle \\ &= \sum_{l\in L} P_{\psi_N}\delta_0(l)\overline{\delta_0(l)} \\ &= P_{\psi_N}\delta_0(0) \\ &= \frac{1}{p}\sum_{n\in N}\psi_N(-n)\pi(n)\delta_0(0).\end{aligned}$$

We know that our representation π acts on δ_0 by translation with l_n. Because the only n with $l_n = 0$ is $n = 0$ itself, we get that this sum reduces to $\frac{1}{p}$, as for every other n in the sum, we will get $\delta_0(l) = 0$, because $l \neq 0$, and for $n = 0$, we get $\psi(0)\pi(0)\delta_0(0) = 1$. Therefore, the normal eigenfunctions f_{ψ_N} we are looking for are

$$f_{\psi_N} = \sqrt{p}\cdot P_{\psi_N}\delta_0.$$

4.2.1 Improving the formulas using $R^N_{M,L}(l)$

We can actually improve the formula $f_{\psi_N} = \sqrt{p} \cdot P_{\psi_N} \delta_0$. To do this, we need to introduce an operator that is best explained visually.

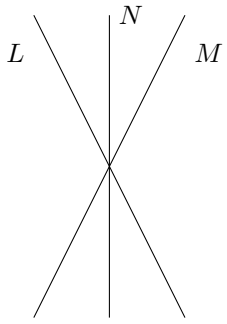

Figure 4.2: Three lines L, M and N

Let L, M and N be three lines in V as in Figure 4.2. Then, for a vector $l \in L$ there is a unique vector $n \in N$ such that $n + l \in M$, as we see in Figure 4.3.

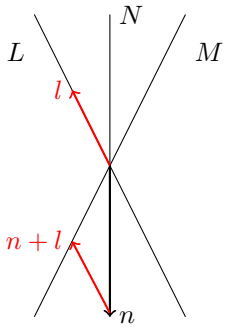

Figure 4.3: The unique vector $n \in N$ such that $n + l \in M$

It is clear that the function that associates this vector $n \in N$ to a vector $l \in L$ is linear, as it is linear with respect to addition and scalar multiplication. It is therefore a linear transformation from L to N that depends on the third line M. We shall denote this operator $R^N_{M,L}$ or, more formally:

Definition 4.8. We define the operator $R^N_{M,L} : L \to N$ as the function that gives us the unique vector $n \in N$ for an $l \in L$ such that $n + l \in M$.

Let us now look again at our formula $f_{\psi_N} = \sqrt{p} \cdot P_{\psi_N} \delta_0$ and try to improve it. Recall that ψ is the central character of π and that ω is a symplectic form on $V = \mathbb{F}_p \times \mathbb{F}_p$.

$$\begin{aligned}\sqrt{p}\cdot P_{\psi_N}\delta_0(l) &= \frac{1}{\sqrt{p}}\sum_{n\in N}\psi_N(-n)\pi(n)\delta_0(l)\\ &= \frac{1}{\sqrt{p}}\sum_{n\in N}\psi_N(-n)\pi(l_n+m_n)\delta_0(l)\\ &= \frac{1}{\sqrt{p}}\sum_{n\in N}\psi_N(-n)\psi\circ\omega\left(\begin{pmatrix}0\\m_n\end{pmatrix},\begin{pmatrix}l\\0\end{pmatrix}\right)\delta_0(l+l_n).\end{aligned}$$

Now, because $\delta_0(l)$ is only non-zero if $l=0$, we only need to check the case where $l+l_n=0$. But we know that there is a unique $n\in N$ with L-component equal to l_n, namely $R^N_{M,L}(-l)$. For the M-component of $R^N_{M,L}(-l)$, let us write $m_{R(-l)}$ instead of $m_{R^N_{M,L}(-l)}$. We can now reduce the sum and improve the formula to

$$\begin{aligned}f_{\psi_N} &= \sqrt{p}\cdot P_{\psi_N}\delta_0(l)\\ &= \frac{1}{\sqrt{p}}\psi_N\circ R^N_{M,L}(l)\cdot\psi\circ\omega\left(\begin{pmatrix}0\\m_{R(-l)}\end{pmatrix},\begin{pmatrix}l\\0\end{pmatrix}\right).\end{aligned}$$

This concludes the abstract part of this section on MUBs, as we have now found an abstract formula in terms of L, M, N, ψ, ω and ψ_N that give us the $p+1$ mutually unbiased bases $\mathcal{B}_N=\{f_{\psi_N}\mid\psi_N:N\to\mathbb{C}^*\}$. In the following subsection we see how we can construct actual bases with this formula.

4.2.2 Actual bases of $\mathbb{C}^p$

We will construct mutually unbiased bases of $\mathbb{C}^p$. To do this we can just pick our lines L and M. If we pick a base for L and M, we get ω as we saw in Example 1.33. Along the way we will choose our central character ψ of our representation π so that things look nicer.

Let $L=\{(x,0)\mid x\in\mathbb{F}_p\}$ and let $M=\{(0,x)\mid x\in\mathbb{F}_p\}$. This gives us the symplectic form

$$\omega: V\times V\to Z$$

$$\omega\left(\begin{pmatrix}a\\b\end{pmatrix},\begin{pmatrix}c\\d\end{pmatrix}\right)=ad-bc.$$

With this we get that $\mathcal{H}\cong\mathbb{C}(\mathbb{F}_p)$, the space of functions from $\mathbb{F}_p\to\mathbb{C}$. Now let us rewrite our formula f_{ψ_N} for a line $N\neq M$. Such a line looks like

$$N=\{(x,ax)\mid x\in\mathbb{F}_p\}$$

for some $a\in\mathbb{F}_p$. But then we get that

$$R^N_{M,L}(-l)=R^N_{M,L}((-x,0))=(x,ax),$$

because $(-x,0)+(x,ax) \in M$. We see that $\psi_N(n) \cong \psi(cx)$ with $c \in \mathbb{F}_p$ because $N \cong \mathbb{F}_p$. This simplifies our formula to

$$f_{\psi_N}(x) = \frac{1}{\sqrt{p}}\psi(-ax^2 + cx).$$

We know that our a value is dependent only on the line N that we pick, which means that every $a \in \mathbb{F}_p$ will be used. Therefore the minus sign in the formula does not matter and can be neglected. This gives us the final formula.

$$f_{\psi_N}(x) = \frac{1}{\sqrt{p}}\psi(ax^2 + cx). \tag{4.1}$$

Recall that in this formula a depends on the line N that we picked, and c depends on the character ψ_N that we picked. Therefore, for every line, we have p functions that form an orthonormal base together. Furthermore, we know that the bases thus obtained are mutually unbiased.

Let us write out these MUBs explicitly. Pick a line N. There are two different situations. First, if $N \neq M$ we can use the formula to get the base of eigenfunctions

$$\mathcal{B}_a = \{f_{\psi_N}(x) = \frac{1}{\sqrt{p}}\psi(ax^2 + cx) \mid c \in \mathbb{F}_p\},$$

where $a \in \mathbb{F}_p$ is the steepness of our line N. This gives us the p orthonormal bases for $\mathbb{C}(\mathbb{F}_p)$. The second situation, where $N = M$, gives us the last base

$$\mathcal{B} = \{\delta_x \mid x \in \mathbb{F}_p\}.$$

For the sake of completeness we can write out the bases $\mathcal{B}_a$ even further in terms of exponential functions, where again it does not matter which central character ψ we have picked, as we vary over all $a \in \mathbb{F}_p$:

$$\mathcal{B}_a = \{\frac{1}{\sqrt{p}} \cdot e^{\frac{2\pi i}{p}(ax^2+cx)} \mid c \in \mathbb{F}_p\}, \qquad a \in \mathbb{F}_p$$

From this we get the MUBs for $\mathbb{C}^p$ by mapping $f \mapsto v$ such that $f(i) \mapsto v_i$ for $v = (v_1, ..., v_p) \in \mathbb{C}^p$. Our bases will stay mutually unbiased because for $f \mapsto v$ and $f' \mapsto v'$ we get

$$\langle v, v' \rangle = \sum_i v_i \overline{v'_i} = \sum_{i \in \mathbb{F}_p} f(i)\overline{f'(i)} = \langle f, f' \rangle.$$

Example 4.9. Let us demonstrate the process for $p = 3$. The base $\mathcal{B}$ becomes

$$\{(1,0,0),(0,1,0),(0,0,1)\}.$$

The base $\mathcal{B}_0$ becomes

$$\begin{aligned}\{&\frac{1}{\sqrt{3}}\cdot\left(e^{\frac{2\pi i}{3}(0\cdot 0^2+0\cdot 0)},e^{\frac{2\pi i}{3}(0\cdot 1^2+0\cdot 1)},e^{\frac{2\pi i}{3}(0\cdot 2^2+0\cdot 2)}\right),\\ &\frac{1}{\sqrt{3}}\cdot\left(e^{\frac{2\pi i}{3}(0\cdot 0^2+1\cdot 0)},e^{\frac{2\pi i}{3}(0\cdot 1^2+1\cdot 1)},e^{\frac{2\pi i}{3}(0\cdot 2^2+1\cdot 2)}\right),\\ &\frac{1}{\sqrt{3}}\cdot\left(e^{\frac{2\pi i}{3}(0\cdot 0^2+2\cdot 0)},e^{\frac{2\pi i}{3}(0\cdot 0^2+2\cdot 1)},e^{\frac{2\pi i}{3}(0\cdot 0^2+2\cdot 2)}\right)\},\end{aligned}$$

the base $\mathcal{B}_1$ becomes

$$\begin{aligned}\{&\frac{1}{\sqrt{3}}\cdot\left(e^{\frac{2\pi i}{3}(1\cdot 0^2+0\cdot 0)},e^{\frac{2\pi i}{3}(1\cdot 1^2+0\cdot 1)},e^{\frac{2\pi i}{3}(1\cdot 2^2+0\cdot 2)}\right),\\ &\frac{1}{\sqrt{3}}\cdot\left(e^{\frac{2\pi i}{3}(1\cdot 0^2+1\cdot 0)},e^{\frac{2\pi i}{3}(1\cdot 1^2+1\cdot 1)},e^{\frac{2\pi i}{3}(1\cdot 2^2+1\cdot 2)}\right),\\ &\frac{1}{\sqrt{3}}\cdot\left(e^{\frac{2\pi i}{3}(1\cdot 0^2+2\cdot 0)},e^{\frac{2\pi i}{3}(1\cdot 0^2+2\cdot 1)},e^{\frac{2\pi i}{3}(1\cdot 0^2+2\cdot 2)}\right)\},\end{aligned}$$

and the base $\mathcal{B}_2$ becomes

$$\begin{aligned}\{&\frac{1}{\sqrt{3}}\cdot\left(e^{\frac{2\pi i}{3}(2\cdot 0^2+0\cdot 0)},e^{\frac{2\pi i}{3}(2\cdot 1^2+0\cdot 1)},e^{\frac{2\pi i}{3}(2\cdot 2^2+0\cdot 2)}\right),\\ &\frac{1}{\sqrt{3}}\cdot\left(e^{\frac{2\pi i}{3}(2\cdot 0^2+1\cdot 0)},e^{\frac{2\pi i}{3}(2\cdot 1^2+1\cdot 1)},e^{\frac{2\pi i}{3}(2\cdot 2^2+1\cdot 2)}\right),\\ &\frac{1}{\sqrt{3}}\cdot\left(e^{\frac{2\pi i}{3}(2\cdot 0^2+2\cdot 0)},e^{\frac{2\pi i}{3}(2\cdot 0^2+2\cdot 1)},e^{\frac{2\pi i}{3}(2\cdot 0^2+2\cdot 2)}\right)\}.\end{aligned}$$

4.3 Gauss Sums

In this section we will use the results from the previous section to prove facts about the *Gauss sum*. The sum is named after Carl Friedrich Gauss, who used it in his study on the law of quadratic reciprocity. He proved all the facts in this section using number theory and clever manipulations on the Gauss sum. For example, he proved that

$$\sum_{x\in\mathbb{F}_p} e^{\frac{2\pi i}{p}x^2} = \begin{cases}\sqrt{p} & \text{if } p\equiv 1 \mod 4,\\ i\cdot\sqrt{p} & \text{if } p\equiv 3 \mod 4.\end{cases}$$

The philosophy of this section is different. We know that the above fact is true, but not *why* it's true. The clever calculations do not *make* it true. We hope to give more insight into this problem with the use of MUBs. To do that, we will first introduce the more abstract definition of a Gauss sum. For the rest of the section, let p be prime and odd.

Definition 4.10. Let ψ be a non-trivial character of $\mathbb{F}_p$. Then we define a *Gauss sum* $G(\psi)$ by

$$G(\psi) = \sum_{x \in \mathbb{F}_p} \psi(x^2).$$

Furthermore, we define

$$G_a(\psi) := \sum_{x \in \mathbb{F}_p} \psi(ax^2)$$

for all $a \in \mathbb{F}_p^*$.

We see that if we pick the character $\psi : x \mapsto e^{\frac{2\pi i}{p}x}$, we get the Gauss sum we discussed above.

Immediately we see a connection between the standard inner product on our function space $\mathbb{C}(\mathbb{F}_p)$ and the Gauss sums. We see that this connection will prove fruitful in proving facts about the Gauss sums in the following result:

Lemma 4.11. *Let ψ be a non-trivial character. Then*

$$|G(\psi)|^2 = p.$$

Proof. We can pick ψ as the central character of the representation π that we used in the construction of the MUBs. By our formula 4.1 on page 51, we can pick $f \in \mathcal{B}_a$ with $c = 0$ and $g \in \mathcal{B}_{a-1}$ with $c = 0$. As these come from different bases, we get:

$$\begin{aligned}
\frac{1}{p} &= |\langle f, g \rangle|^2 \\
&= \left| \frac{1}{p} \sum_{x \in \mathbb{F}_p} \psi(ax^2) \overline{\psi((a-1)x^2)} \right|^2 \\
&= \left| \frac{1}{p} \sum_{x \in \mathbb{F}_p} \psi(ax^2 - ax^2 + x^2) \right|^2 \\
&= |\frac{1}{p} G(\psi)|^2,
\end{aligned}$$

which proves the lemma. □

Remark. In fact, in this way we can even prove that

$$|\sum_{x \in \mathbb{F}_p} \psi(ax^2 + bx)|^2 = p$$

for all $a, b \in \mathbb{F}_p$, in much the same way as we proved the lemma.

4.4 Discussion

We see that in a sense, the absolute value of a Gauss sum is related to the lines in the plane $\mathbb{F}_p \times \mathbb{F}_p$. Here ends the results we obtained on Gauss sums using representation. In this section we discuss whether there exists a stronger relation between Gauss sums and our representation and vector space. We therefore prove some known facts about Gauss sums, that may help in finding this relation. We will end with some interesting open questions concerning these results.

4.4.1 Manipulating Gauss sums

In this subsection we manipulate Gauss sums to prove some known facts about these sums. We do this in the hope of finding a stronger relation between Gauss sums and our representation than we did in the previous section.

An interesting question arises if we look at the case where $p \equiv 1 \mod 4$, which gives us that -1 is a square in $\mathbb{F}_p$. We can easily manipulate the Gauss sum in the following sense, where $a \in \mathbb{F}_p$ is such that $a^2 = -1$:

$$G(\psi) = \sum_{x \in \mathbb{F}_p} \psi(x^2) = \sum_{x \in \mathbb{F}_p} \psi((ax)^2) = \sum_{x \in \mathbb{F}_p} \psi(-x^2) = \overline{G(\psi)}, \tag{4.2}$$

where we used the substitution $x \mapsto ax$ in the second equality, which is allowed because $\mathbb{F}_p$ is a field. This proves that in this case the Gauss sum must be real-valued, and therefore $G(\psi) = \pm\sqrt{p}$, depending on ψ. We can ask ourselves whether it is possible to prove this using our construction of π, V and $\mathcal{H}$?

Some other known facts can also quite easily be proven by manipulation. To do this we will need to do some number theory on the structure of squares in $\mathbb{F}_p^*$. We will denote the squares in $\mathbb{F}_p^*$ with $(\mathbb{F}_p^*)^2$ (notice that this is a group under multiplication). What we want to find is what $(\mathbb{F}_p^*)/(\mathbb{F}_p^*)^2$ looks like. Because

$$(\mathbb{F}_p^*) \cong \mathbb{Z}/(p-1)\mathbb{Z},$$

we see that

$$(\mathbb{F}_p^*)^2 \cong 2\mathbb{Z}/(p-1)\mathbb{Z} \cong \mathbb{Z}/\left(\frac{p-1}{2}\right)\mathbb{Z}$$

as p is prime and odd.

This gives us all the structure we need, because it gives us that $(\mathbb{F}_p^*)/(\mathbb{F}_p^*)^2$ consists of two elements: $(\mathbb{F}_p^*)^2$ and $a(\mathbb{F}_p^*)^2$, where a is any non-square in $\mathbb{F}_p^*$. With this structure we can actually manipulate the Gauss sum in a pretty nice way to get the following result:

Lemma 4.12. *Let ψ be a non-trivial character of $\mathbb{F}_p$. Then*

$$G_a(\psi) = \left(\frac{a}{1}\right) G(\psi),$$

where $(a/1)$ is the Legendre symbol.

Proof. We will prove the two cases

- a is a square in $\mathbb{F}_p$,
- a is not a square in $\mathbb{F}_p$.

For the first case we get $(a/1) = 1$ and therefore we can mimic equation 4.2 on page 54 to get the result with $b \in \mathbb{F}_p$ such that $b^2 = a$ which gives

$$G_a(\psi) = \sum_{x \in \mathbb{F}_p} \psi(ax^2) = \sum_{x \in \mathbb{F}_p} \psi((bx)^2) = \sum_{x \in \mathbb{F}_p} \psi(x^2) = G(\psi).$$

The second case however is interesting. Because $x^2 = (-x)^2$ and $x \neq -x$ for $x \in \mathbb{F}_p^*$, we get that every square except 0 will be counted twice in the Gauss sum. This leads to the following, if a is not a square:

$$\begin{aligned} 0 = \sum_{x \in \mathbb{F}_p} \psi(x) &= \psi(0) + \sum_{x \in \mathbb{F}_p^*} \psi(x) \\ &= \psi(0) + \frac{1}{2} \cdot \left(\sum_{x \in (\mathbb{F}_p^*)^2} \psi(x^2) + \sum_{x \in (\mathbb{F}_p^*)^2} \psi(ax^2) \right), \end{aligned}$$

which gives us that

$$\sum_{x \in \mathbb{F}_p^*} \psi(x^2) = - \sum_{x \in \mathbb{F}_p^*} \psi(ax^2)$$

if a is not a square. □

Remark. This also gives us that, if -1 is not a square, e.g. $p \equiv 3 \mod 4$, then

$$G(\psi) = -\overline{G(\psi)},$$

which means that $G(\psi) = \pm i \cdot \sqrt{p}$, because it must be completely imaginary.

This concludes the number-theoretical manipulations of Gauss sums. The interesting question is whether or not these results can be obtained from the structure of π, V and $\mathcal{H}$, to see if that improves our insight into *why* these results are true. We could reformulate this into the open question:

Does there exist a proof of the sign of a Gauss sum that relies only on representation theory?

Bibliography

[1] L. Auslander and R. Tolimiere. Is computing with the finite fourier transform pure or applied mathematics? *Bulletin of the American Mathematical Society*, 1(6):849–897, 1979.

[2] Yakir Aharonov Berthold-Georg Englert. The mean kings problem: Prime degrees of freedom. *Physical Review Letters*, 284(1):1–5, 2001.

[3] Ian D. Brown. Representation of finitely generated nilpotent groups. *Pacific Journal of Mathematics*, 45(1):13–26, 1973.

[4] Michèle Vergne Gerard Lion. *The Weil representation, Maslov index and Theta series*. Springer, 1980.

[5] Haret C. Rosu. Michel Planat and Serge Perrine. A survey of finite algebraic geometrical structures underlying mutually unbiased quantum measurements. *Foundations of Physics*, 36(11):1662–1680, 2006.

[6] A. Karlsson N.J. Cerf, M. Bourennane and N. Gisin. Security of quantum key distribution using d-level systems. *Physical Review Letters*, 88(12):127902, 2002.

[7] Jean-Pierre Serre. *Linear Representations of Finite Groups*. Springer, 1977.

[8] Vwani Roychowdhury Somshubhro Bandyopadhyay, P. Oscar Boykin and Farrokh Vatan. A new proof for the existence of mutually unbiased bases, 2001.

Index