Graphing gamma matrices, and connecting the two Heisenberg groups in dimension 4

Niklas Kastlander

Bachelor's Thesis (15 ECTS)
Supervisor: Ingemar Bengtsson (Fil.dr.)
Department of physics
Stockholm University
Spring term 2022

Abstract

There are two Heisenberg groups in four dimensions - one generated by taking the direct product of the Heisenberg group in two dimensions with itself, and one by using the standard Weyl-Heisenberg generator definition. Although these are typically taken as distinct, and used in separate contexts, there exists an as of yet unexplored connection between them. Using a common representation, the two groups permute each other, or equivalently, lie in each other's Clifford groups. This thesis explains and explores this connection.

On the way, we take a detour to a group intimately tied to the Heisenberg group, namely the group generated by the gamma matrices. We explore a graphical representation for the commutation relations of this group in four dimensions, introduced in 1996 by David Goodmanson [1], and consider if a similar representation can be found in higher dimensions. On this front, only a partial success is reached - while a generalization is found, it lacks key features of the four-dimensional case.

Contents

1 Introduction

Although the framing might be unfamiliar, most students of quantum mechanics have in practice already worked with the Heisenberg group in two dimensions. Essentially, this is the group of the Pauli matrices, some of the most well known objects in the entire field. It is perhaps then not so surprising that the Heisenberg group is of large importance, and that it has been extended to higher dimensions as well. One can think of this as analogous to generalizing the Pauli matrices themselves to unitary matrices in higher dimensions.

When considering composite systems, where each subsystem is connected to a Pauli matrix, it is reasonable to consider the tensor product of Pauli matrices to be connected to the system as a whole. This corresponds to generating Heisenberg groups in higher dimensions by taking the direct product of smaller Heisenberg groups. For a specific example from quantum computing, where the composite system is built from qubits, see Gottesman [2]. We will cover a general approach for this in Section 3.

As part of Section 3 we will encounter the gamma matrices, which play an important physical role in the Dirac equation. Although typically introduced from that angle, i.e. strongly associated with the physical space-time, we will be mostly considering them from a different direction, namely as a practical application of the Heisenberg group. In Section 4 we dive into the possibilities of representing the commutation relations of these gamma matrices graphically, first in a four-dimensional space-time, and later in higher dimensions.

After this graphical excursion, we return to a different way of generalizing the Heisenberg group, which does not require using smaller Heisenberg groups as building blocks. The result, the so called Weyl-Heisenberg groups, are useful in several areas of quantum mechanics, as well as classical signal processing. This generalization is covered in Section 5.

Now, we could be satisfied with there simply being two distinct ways to do this generalization, and use whichever one works best for each application. But the pattern-seeking researcher might wonder if there isn't some link between them, which could potentially be used to view them in a different light. In Section 6 we will build a common matrix representation for the two different Heisenberg groups in 4 dimensions, and in Sections 7 and 8 we will use this to discuss such a link.

The link has to do with the Clifford group, the group consisting of those unitary operators which permute the elements of the Heisenberg group. As is usual in science, the name has a confusing origin. Originally the name Clifford group belonged to the group of gamma matrices, due to the fact that the gamma matrices form what is called a Clifford *algebra*. This name spread to the Heisenberg group (originally a generalization of the gamma matrices), and to its permuting group. In the end the gamma group and the Heisenberg group got different names, and the permuting group became known as the sole Clifford group.

2 H_2, and some fundamental theory

Consider the set of the Pauli matrices, including the identity element:

$$\left\{\sigma_0 = \mathbb{1},\ \sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},\ \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix},\ \sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}\right\} \tag{1}$$

These constitute a *unitary operator basis* for $\mathbb{C}^2$ - by which we mean that every linear operator acting on $\mathbb{C}^2$ can be written as a linear combination of these four operators. Equivalently, each complex 2×2-matrix can be written as such a linear combination. We can also define a scalar product between two operators A and B as $tr(A^\dagger B)$, for which this basis is orthogonal. Viewing the set of Pauli matrices as an orthogonal basis like this is one of the fundamental reasons why it is of interest. We will however mainly concern ourselves with a different perspective: viewing this set in the context of *group theory*.

A *group* is a set associated with some binary operation[1] that satisfies certain properties. The most essential one is that a group is *closed* under the operation - meaning that when acting on two elements in the group, the operation returns another group element. The remaining properties are that the group contains an identity element and inverses to all elements, and that the binary operation is associative.

In the case of the Pauli matrices, we will consider matrix multiplication (which is inherently associative) as our binary operation. It is well known that the product of two Pauli matrices yields a new Pauli matrix, which motivates the introduction of group theory. However, the product typically also includes a phase factor - for instance $\sigma_1\sigma_2 = i\sigma_3$. Technically speaking, the set of Pauli matrices isn't closed, and thus isn't a group. We can however fix this easily by extending the set to include phase factors. If we include $i\sigma_3$, $-\sigma_3$ and $-i\sigma_3$, and the same for all other Pauli matrices, we obtain a true closed group. This group is what we formally mean with the *Heisenberg group* in two dimensions, which will be denoted in this paper as H_2.

$$\begin{aligned} \{\ &\sigma_0, \ -\sigma_0, \ i\sigma_0, \ -i\sigma_0, \\ &\sigma_1, \ -\sigma_1, \ i\sigma_1, \ -i\sigma_1, \\ &\sigma_2, \ -\sigma_2, \ i\sigma_2, \ -i\sigma_2, \\ &\sigma_3, \ -\sigma_3, \ i\sigma_3, \ -i\sigma_3\ \} \end{aligned}$$

Figure 1: The Heisenberg group in two dimensions, H_2, is formally defined as the above set where we can multiply elements together by matrix multiplication.

An important question is what we mean by *a group in n dimensions*.[2] In order to explain this, we first need to discuss the concept of representations.

When we specify a set of specific matrices in order to define a group, as we did with H_2 in Figure 1, that set is called a *matrix representation* for that group. One group can have many different representations - for instance we could have written the Pauli matrices in a different basis. We call a matrix representation of a group *irreducible*, if the matrices can not be reduced to block-triangular matrices via some unitary transformation. As an example, the reader might know that the rotational group $SO(3)$ has irreducible representations for matrices of size 3,5,7, and so on.

The Heisenberg group H_2 *only* has irreducible representations for 2×2-matrices, which means the group itself is specifically connected to two-dimensional vector spaces. This is what motivates the naming convention that H_2 is a group in two dimensions. When we later discuss extending the Heisenberg group to dimension n, what we mean is to find groups which naturally extend the structure of H_2, but which have irreducible representations for $n \times n$-matrices instead.

[1] An operation taking in two objects, and returning another one, e.g. scalar multiplication.
[2] Especially since *the dimension of a group* is an entirely separate thing.

Instead of listing all of the elements in the group explicitly, one can also describe this group in terms of the *generators* $i\mathbb{1}$, σ_1 and σ_3.[3] What we mean by this is that every element in the group can be broken down to some product containing only these three factors (with possible repetition). As an example, $-\sigma_2 = i\mathbb{1} \times \sigma_3 \times \sigma_1$. Such products of generators are more commonly called *words*.

Because all Pauli-matrices either anti-commute or commute, we can always change the order of a word in H_2, as long as we are okay with a potential sign change. This means we could choose a standard form, e.g. $(i\mathbb{1})^{n_1}(\sigma_1)^{n_2}(\sigma_3)^{n_3}$, on which all elements of the group could be written. Continuing our example from above, writing $-\sigma_2$ on this form requires a sign change, which yields $-\sigma_2 = (i\mathbb{1})^3\sigma_1\sigma_3$. One advantage of having a standard form is that the group elements can now be associated with the vectors (n_1, n_2, n_3) - giving us an alternate way to think of how the group is structured.

Finally, we state a specific fact we will need later on. Consider the *Hadamard operator*:

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \tag{2}$$

Applying this operator as a change of basis (or active transformation) has the effect of swapping σ_1 and σ_3, while leaving σ_0 and σ_2 mostly unchanged:

$$\begin{cases} H\sigma_0 H^{-1} & = \sigma_0 \\ H\sigma_1 H^{-1} & = \sigma_3 \\ H\sigma_2 H^{-1} & = -\sigma_2 \\ H\sigma_3 H^{-1} & = \sigma_1 \end{cases} \tag{3}$$

3 Generalizing H_2 via the direct product

3.1 Dimension four: $H_{2\times 2}$

One simple way to generalize H_2 to dimension four is to take the direct product $H_2 \times H_2$ - that is, to consider the set of matrices $\{\sigma_i \otimes \sigma_j\}$. This yields a set of 16 4×4-matrices, which once again forms a unitary operator basis, but this time for operators on $\mathbb{C}^4$. If we extend this set in the same way as for H_2, by multiplying with phase factors, we get a true Heisenberg group containing 64 elements. One can also check that irreducible representations of this group is only possible with 4×4-matrices, meaning this is really a group in four dimensions. We will denote this group as $H_{2\times 2}$.

3.2 The gamma matrices

One practical reason why $H_{2\times 2}$ is of interest is that its elements are closely linked to the *gamma matrices*. These are used in the Dirac equation, the generalization of the Schrödinger equation for relativistic particles in space-time. The Dirac equation can be written on the following form:

$$\left[\sum_{\mu=0}^{3}\left(i\gamma_\mu \frac{\partial}{\partial x_\mu}\right) - m\right]\Psi(\vec{x}) = 0 \tag{4}$$

Here Ψ is a wavefunction in four-dimensional space-time. We describe space-time with the four-vector $\vec{x}$, which has the components $(x_0 = t, x_1, x_2, x_3)$. In this thesis, what we are really interested in are the gamma matrices γ_μ. These matrices can be chosen in many different ways, as long as they satisfy the following two conditions: They all have to anti-commute with each other, and they have to square to either $-\mathbb{1}$ or $+\mathbb{1}$ depending on if they are associated with a time-like or space-like dimension. Sadly there is no consensus on which sign each type of dimension should be associated with - we will opt for associating the time-like dimensions with a negative sign. For a concrete example, see Figure 2 on the next page.

[3] For this group, and in general, there are many ways to choose these generators.

$$\gamma_0 = \begin{pmatrix} i & 0 & 0 & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & -i & 0 \\ 0 & 0 & 0 & -i \end{pmatrix}, \quad \gamma_1 = \begin{pmatrix} 0 & 0 & 0 & i \\ 0 & 0 & i & 0 \\ 0 & -i & 0 & 0 \\ -i & 0 & 0 & 0 \end{pmatrix}, \quad \gamma_2 = \begin{pmatrix} 0 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{pmatrix}, \quad \gamma_3 = \begin{pmatrix} 0 & 0 & i & 0 \\ 0 & 0 & 0 & -i \\ -i & 0 & 0 & 0 \\ 0 & i & 0 & 0 \end{pmatrix}$$

Figure 2: One of many possible choices of gamma matrices for a four-dimensional space-time where the first dimension is time-like, and the remaining three are space-like. Feel free to verify that they all anti-commute, and that the first one squares to $-\mathbb{1}$ while the other three square to $+\mathbb{1}$.

We can write both conditions very compactly using the anti-commutator and some tensor notation:

$$\{\gamma_\mu, \gamma_\nu\} = 2\eta_{\mu\nu}\mathbb{1} \tag{5}$$

The tensor η is called the *metric*, and is tied to the geometry of space-time. Since we will only consider η on a surface level, we can ignore its tensor nature, and see it as a symmetric 4×4-matrix where each element $\eta_{\mu\nu}$ is a number. The first condition (anti-commutativity) corresponds to all off-diagonal elements in η being zero. The second condition (the gamma matrices square to $\pm\mathbb{1}$) correspond to the diagonal elements in η being $\pm\mathbb{1}$. Thus, we can write η for our case as:

$$\eta = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{6}$$

Thus far we have been working with the 4d space-time we are familiar with, but mathematically there is no constraint on how many space- or time-like dimensions there are. In order to classify the different cases it is enough to consider the diagonal elements of the metric tensor, so we introduce the *signature* as the n-tuple of the diagonal elements, i.e. $(-\ +\ +\ +)$ in our case. The signature can also refer to the number of positive elements minus the number of negative elements, i.e. $3 - 1 = 2$.[4] Since the order of the dimensions is less important, these two definitions of the signature are more or less the same, and we will use them interchangeably.

3.3 The gamma group

While the four gamma matrices themselves are of central importance, we also care about products of gamma matrices. As it turns out, the gamma matrices are generators for the *gamma group*, whose elements are listed in Table 1 below.

Elements	Compact notation	#
$\mathbb{1}$		1
$\gamma_0,\ \gamma_1,\ \gamma_2,\ \gamma_3$		4
$\gamma_0\gamma_1,\ \gamma_0\gamma_2,\ \gamma_0\gamma_3,\ \gamma_1\gamma_2,\ \gamma_1\gamma_3,\ \gamma_2\gamma_3$	$\gamma_{01},\ \gamma_{02},\ \cdots$	6
$\gamma_0\gamma_1\gamma_2,\ \gamma_0\gamma_1\gamma_3,\ \gamma_0\gamma_2\gamma_3,\ \gamma_1\gamma_2\gamma_3$	$\gamma_3\gamma_5,\ \gamma_2\gamma_5,\ \gamma_1\gamma_5,\ \gamma_0\gamma_5$	4
$\gamma_0\gamma_1\gamma_2\gamma_3$	γ_5	1

Table 1: Elements in the gamma group. Since gamma matrices anti-commute, the ordering of the matrices is unimportant up to a sign change. Formally the group contains each element in the table twice (once with a negative sign), but we will not go into this in detail.

[4]As mentioned, some conventions flip the signs, and thus get the space-time signature $(+\ -\ -\ -)$ or -2.

The matrix γ_5 has its slightly curious name from the time when the four gamma matrices in the Dirac equation were numbered 1 through 4. The products $\gamma_\mu\gamma_5$ are compact ways of writing the triple products - effectively γ_μ cancels one of the four matrices in γ_5, leaving three left. We also note that because of the anti-commutation, $\gamma_{\mu\nu}$ is often written (or even defined) as $\gamma_{\mu\nu} = \frac{1}{2}(\gamma_\mu\gamma_\nu - \gamma_\nu\gamma_\mu)$.

Now, since $H_{2\times2}$ is a complete basis for 4×4-matrices, it is of course always possible to express any of the 16 matrices in the gamma group as a linear combination of elements in $H_{2\times2}$. But the fact that the gamma group and $H_{2\times2}$ are of the same size (ignoring phase factors) hints at a deeper connection between them. In fact, the two groups are essentially *the same group*, with some small difference in which phase factors are allowed. Formally, we can obtain all the elements in the gamma group (regardless of how we chose our four gamma matrices originally) by applying an appropriate change of basis[5] T to $H_{2\times2}$. For a proof of this, see for instance Appendix A2 of Jauch [3]. Since $\pm\mathbb{1}$ remains unchanged under a change of basis, we can easily see that the squares of the gamma matrices (and thus also the signature) is preserved during a basis change.

What this means is that in order to understand all possible ways we can choose our gamma matrices, it is enough to investigate in which ways we can choose gamma matrices directly from $H_{2\times2}$. Since $H_{2\times2}$ includes the phase factor variants of all matrices, it is also very easy to modify the signature of a gamma matrix set. By multiplying one or more gamma matrices with i, the sign of their squares flip. Thus, as soon as you have one set of gamma matrices, you can immediately generate sets with any signature you want.

There are however cases when one does not wish to work with the entirety of $H_{2\times2}$, and instead restricts oneself to only the real elements. One such example comes from modern particle physics: In the standard model, neutrinos are assumed to be mass-less, yet this has already been proven wrong experimentally. A new equation including a neutrino mass term can have one of two forms: either a modified Dirac equation, or a different equation known as the Majorana equation - for details see Ohlsson [4]. The Majorana equation, among other things, requires that the gamma matrices be chosen purely real. The difference between Dirac and Majorana can be tested experimentally - Majorana neutrinos would allow for a so called *neutrinoless double beta decay*. The search for such a reaction is ongoing, see for instance the projects [5], [6], [7] and [8].

Regardless of the physical context, a restriction to real matrices has an important impact on the underlying mathematics. We can no longer choose the phase factors freely, and thus the trick to generate all signatures no longer works. In fact, using real matrices only allows for a fraction of all possible signatures:

$$sign \equiv 0, 1, 2 \mod 8 \tag{7}$$

The derivation of this formula is outside the scope of this thesis, but the interested reader can find it in Freund [9]. Since we will only be dealing with even-dimensional space-times, all signatures we encounter will be even. In practice we can thus ignore the case where the signature is congruent with one.

We see from the formula that the signature $(- \; + \; + \; +)$ is still valid - however the alternate convention $(+ \; - \; - \; -)$ is not. To use the alternate convention, one must either use purely imaginary gamma matrices, or rewrite the Dirac equation to include a factor i in front of all gamma matrices.

3.4 Higher dimensions: $H_{2\times2\times2\cdots}$

We can generalize the above discussion to dimension 2^n by taking the direct product of n copies of H_2, and thus considering $2^n \times 2^n$-matrices built by tensor products of n Pauli matrices. Most statements made above generalize naturally, but some aspects require a little bit extra thought.

For four dimensions we did not need to make a distinction between the size of the gamma matrices, and the dimension of space-time - both had the same value. For higher dimensions these two are not equal, and we have to separate them out. The gamma matrices can either be seen as $2^n \times 2^n$ matrices, chosen from the Heisenberg group in 2^n dimensions, or (more commonly) as the matrices that allow us to express the Dirac equation in $2n$ space-time dimensions. The number of matrices you need for the Dirac equation is the same as the space-time dimension, i.e. $2n$. As an example, six-dimensional space-time requires six gamma matrices that all are of size 8×8.

[5] Note that we are here considering a basis change of the vector space C^4, and **not** a basis change of the operator space.

3.5 Commutation relations

Let us consider the group structure of H_2 in more detail. Since the general case is rather easy to obtain as a generalization of the real case,[6] we will work with the real case. We introduce $\widetilde{\sigma}_2$ as the real version of σ_2:

$$\widetilde{\sigma}_2 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \tag{8}$$

If we ever want to re-obtain the usual Pauli matrix, we simply reapply a factor i ($\sigma_2 = i\widetilde{\sigma}_2$). For the real Pauli matrices, we obtain the following multiplication table:[7]

$\times$	σ_0	σ_1	$\widetilde{\sigma}_2$	σ_3
σ_0	σ_0	σ_1	$\widetilde{\sigma}_2$	σ_3
σ_1	σ_1	σ_0	σ_3	$\widetilde{\sigma}_2$
$\widetilde{\sigma}_2$	$\widetilde{\sigma}_2$	$-\sigma_3$	$-\sigma_0$	σ_1
σ_3	σ_3	$-\widetilde{\sigma}_2$	$-\sigma_1$	σ_0

Table 2: Multiplication table for real Pauli matrices.

The important takeaways are these two properties:

(1a) $(\widetilde{\sigma}_2)^2 = -\mathbb{1}$, the remaining matrices square to $+\mathbb{1}$.

(1b) σ_0 commutes with everything; $\sigma_1, \widetilde{\sigma}_2$ and σ_3 anti-commute with each other.

When we jump to $H_{2\times 2}$ (or more generally H_{2^n}), we can apply this multiplication table and Property 1a & 1b separately for each term in the tensor product. As an example, that looks something like this:

$$(\sigma_0 \otimes \sigma_1) \cdot (\sigma_3 \otimes \widetilde{\sigma}_2) = \sigma_0\sigma_3 \otimes \sigma_1\widetilde{\sigma}_2 = \sigma_3 \otimes \sigma_3 \tag{9}$$

This also means we can generalize the two properties to a general number of Pauli matrix factors:

(2a) A matrix has a square of $-\mathbb{1}$ when its tensor product contains an odd number of $\widetilde{\sigma}_2$ factors, and otherwise has a square of $+\mathbb{1}$.

(2b) Compare the tensor products for two gamma matrices factor by factor, and count the number of times the factors from the first and the second matrix anti-commute. If this number is odd, the gamma matrices anti-commute, and if the number is even, they commute.

$$\begin{aligned} \gamma_a &= \sigma_1 \otimes \boxed{\widetilde{\sigma}_2} \otimes \boxed{\sigma_3} \\ \gamma_b &= \sigma_0 \otimes \boxed{\sigma_1} \otimes \boxed{\widetilde{\sigma}_2} \end{aligned} \qquad\qquad \begin{aligned} \gamma_a &= \sigma_1 \otimes \widetilde{\sigma}_2 \otimes \boxed{\sigma_3} \\ \gamma_c &= \sigma_1 \otimes \sigma_0 \otimes \boxed{\sigma_1} \end{aligned}$$

Figure 3: Examples for determining if two gamma matrices commute or anti-commute according to Property 2b. Anti-commuting pairs of Pauli matrices are marked with bold outlines. For γ_a and γ_b, the first pair of factors commute, and the last two pairs anti-commute. We have an *even* number of anti-commuting pairs, and thus γ_a commutes with γ_b. For γ_a and γ_c only the last pair of factors anti-commute with each other, meaning we have an *odd* number of anti-commuting pairs. Thus γ_a and γ_c anti-commute.

[6] One only needs to add phase factors at appropriate places.

[7] As a sidenote, multiplication tables like these highlight a more abstract perspective on the idea of a group. The real variant of H_2 can be seen as a set of 8 abstract objects (which we've happened to label $\pm\sigma_i$ in this case, but they might as well be named $\pm a, \pm b, \cdots$), and rules for combining these objects.

4 Graphical representations of gamma matrices

4.1 Dimension four: Goodmanson's graph

In order to better visualize the commutation relations in the four-dimensional case, we can use a very compact graph representation by Goodmanson [1]:

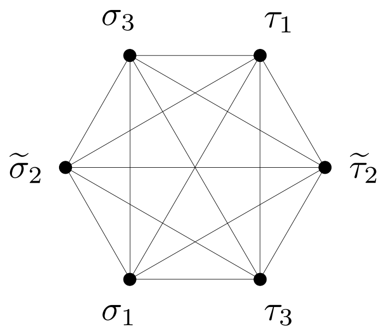

Figure 4: Goodmanson's graph. Each edge represents a 4×4 gamma matrix, and two such matrices anti-commute precisely when their corresponding edges meet at a vertex.

Each edge[8] in this graph represents a gamma matrix in $H_{2\times 2}$. In order to keep track of the order in the tensor product, we introduce a second set of Pauli matrices $\{\tau_i\}$. These are mostly equivalent to the σ:s - the only important distinction is that whenever we take a tensor product it is always in the order $\sigma \otimes \tau$ and never $\tau \otimes \sigma$. As can be expected, edges in the graph that connect σ_i and τ_j correspond to their tensor product $\sigma_i \otimes \tau_j$.

This includes all elements except for those with σ_0 or τ_0 as a factor. Luckily we also have some edges left, namely those that connect either two σ:s or two τ:s. The association made here is slightly more involved: an edge connecting two σ:s represents the tensor product $\sigma_i \otimes \tau_0$, where σ_i is the σ vertex *not* connected to the edge in question. Thus for example the edge $\sigma_1 \widetilde{\sigma}_2$ corresponds to $\sigma_3 \otimes \tau_0$. The same logic applies for edges connecting two τ:s.

Finally, the identity element $\sigma_0 \otimes \tau_0$ is not included in the graph. This is a very small problem however, since it trivially commutes with all other gamma matrices, and thus typically doesn't have to be considered.

The beauty of this graph representation is that Property 2b from Section 3.5 becomes very simple: two edges anti-commute if they meet at a vertex, and commute otherwise. Using this, we can immediately see that there can at most be five mutually anti-commuting matrices, since there are five edges sharing a single vertex. We can also see that there can at most be three mutually commuting matrices (four if you count the identity matrix), since there are six vertices in total, and each edge has to have two distinct ones. These examples are illustrated in Figure 5 below. This type of relation is of course possible to find purely algebraically (see for instance Eddington [10]), but it is clear that the speed and certainty with which one identifies them from this graph is very valuable.

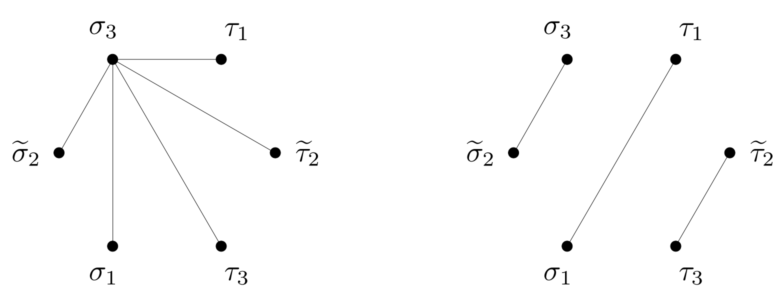

Figure 5: To the left: The largest possible set of anti-commuting gamma matrices; 5 edges meeting at a single vertex. To the right: The largest possible set of commuting gamma matrices; 3 edges sharing no vertices.

[8] If the reader is unfamiliar with graph theory, the lines are called *edges*, and the points *vertices* or *nodes*.

From one of these maximally anti-commuting sets, we can then pick four of the five matrices as our de facto gamma matrices to use in the Dirac equation.[9] Since we're working with real gamma matrices, where the question of available signatures becomes important, we can mark in the graph which matrices have positive and negative squares. This can be seen in Figure 6 below. We see that each anti-commuting 5-tuple contains three positive, and two negative edges. When choosing our four gamma matrices we can thus end up with the signatures $(++--)$ or $(+++-)$. This is precisely the signatures we expect from Formula 7.

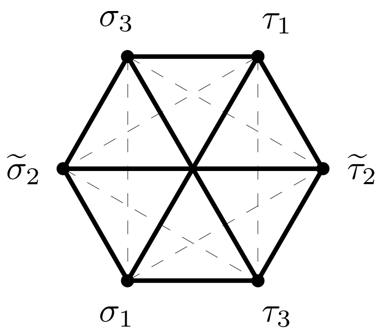

Figure 6: Goodmanson's graph, where the edges have been marked to indicate the squares of their corresponding gamma matrices. Thick continuous edges square to $+\mathbb{1}$, whereas thin dashed edges square to $-\mathbb{1}$.

4.2 Higher dimensions: the polygonal representation

Now, let us introduce a more general graphical representation, that in principle works regardless of how many Pauli matrices are included in the tensor product. To separate this from Goodmanson's representation, we'll refer to this as the *polygonal* representation. The name has been chosen since the gamma matrices will be represented with arbitrarily large polygons, and not only edges as in the Goodmanson case.

To obtain the polygonal representation, we first associate each gamma matrix $\sigma_i \otimes \tau_j \otimes \cdots$ with the set of its Pauli matrix factors $\{\sigma_i, \tau_j, ...\}$. Then we remove all identity matrices (σ_0, τ_0, etc.) from these sets. The resulting set is seen as a list of vertices, for which a polygon can be drawn. If the polygon has four or more vertices there is some ambiguity in exactly how to draw it - but the important part is really the vertices it visits, so this ambiguity is a visual problem rather than a mathematical one. See Table 3 below for some examples of how the representation works.

Note that technically the Greek letter following τ is υ (upsilon), meaning that tensor products with three factors should be $\sigma \otimes \tau \otimes \upsilon$. Given the risk of confusing υ with an italic v, we instead go backwards in the Greek alphabet, and write the tensor products on the form $\rho \otimes \sigma \otimes \tau$.

Matrix	Set of vertices	Graph object
$\rho_0 \otimes \sigma_0 \otimes \tau_1$	$\{\tau_1\}$	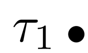
$\rho_0 \otimes \sigma_1 \otimes \tau_3$	$\{\sigma_1, \tau_3\}$	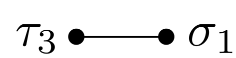
$\widetilde{\rho}_2 \otimes \sigma_3 \otimes \tau_1$	$\{\widetilde{\rho}_2, \sigma_3, \tau_1\}$	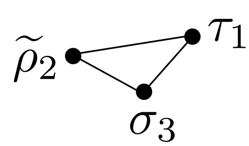

Table 3: Examples of how different 8×8-matrices are represented in the polygonal graph.

[9]This method of finding four gamma matrices will in fact yield *all* possible choices.

In Figure 7 below we've drawn the full polygonal graph for 4×4-matrices. Note the difference to Goodmanson's graph: There are no edges between vertices of the same type, i.e between two σ:s or between two τ:s. The gamma matrices those edges corresponded to, are here instead represented by the six individual vertices, as indicated by the rings around them.

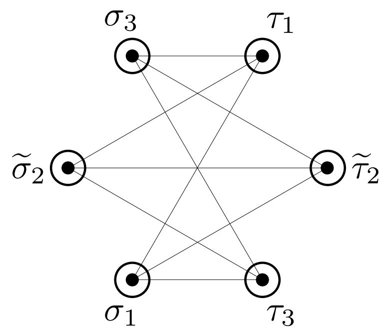

Figure 7: Polygonal graph. Here the gamma matrices are represented both by the edges and the vertices, as indicated by the rings.

For the polygonal graph, the anti-commutation rule follows directly from Property 2b in Section 3.5. Since the position in the tensor product is represented by the vertex type (σ, τ, etc.), an anti-commuting pair of factors here correspond to the two polygons visiting two anti-commuting vertices of the same type. However by not having σ_0 as a vertex, *all* vertices of the same type are anti-commuting. Thus what matters is really whether the two polygons visit *distinct* vertices of the same type. If this occurs an odd number of times, the polygons anti-commute. Examples of this rule can be found in Figure 8:

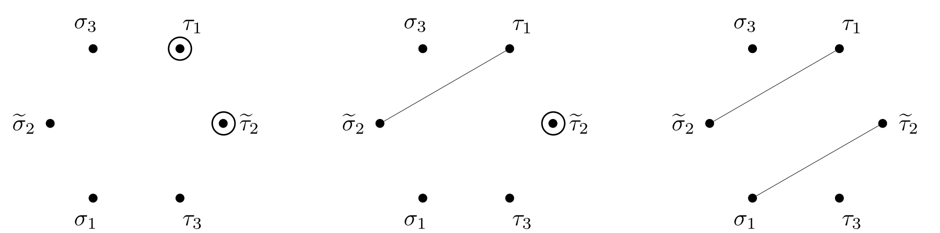

Figure 8: Examples of anti-commutation relations. To the left we have two distinct τ vertices, which form an anti-commuting pair. The gamma matrices associated with these vertices thus also anti-commute. In the middle we have a similar situation, except now one of the polygons also includes a σ vertex. Since only one of them visits a σ vertex this doesn't matter, and the two matrices still anti-commute. Finally, we have the situation where both polygons visit σ vertices. Since they visit distinct σ vertices we now have two anti-commuting pairs (one $\sigma - \sigma$ and one $\tau - \tau$). This is an even amount, and thus the two edges commute.

Specifically for 4×4-matrices we can split the rule up into three cases:

(3a) Two vertices anti-commute if they are of the same type.

(3b) An edge and a vertex anti-commute if they do not intersect.[10]

(3c) Two edges anti-commute if they intersect.

[10]We say that two edges *intersect* when they meet at the same vertex.

When the anti-commutation rule is expressed like this, the similarity to Goodmanson's representation is hopefully apparent. As one can verify, if we now change back to representing all gamma matrices with edges, the three cases above all result in "two edges anti-commute if they intersect", as we expect.

In four dimensions it is clear that Goodmanson's graph is vastly superior to the polygonal graph - one doesn't have to keep track of both edges and vertices, and the anti-commutation rule is as simple as can be. Thus the polygonal graph for 4×4-matrices is more of a prototype, used to prove and/or generate the Goodmanson representation. Maybe we could do the same in higher dimensions - first generate the relevant polygonal graph, and then try to change the representation so that all gamma matrices are associated with edges? Sadly, we encounter two important roadblocks already in six dimensions. These are illustrated in Figure 9 below.

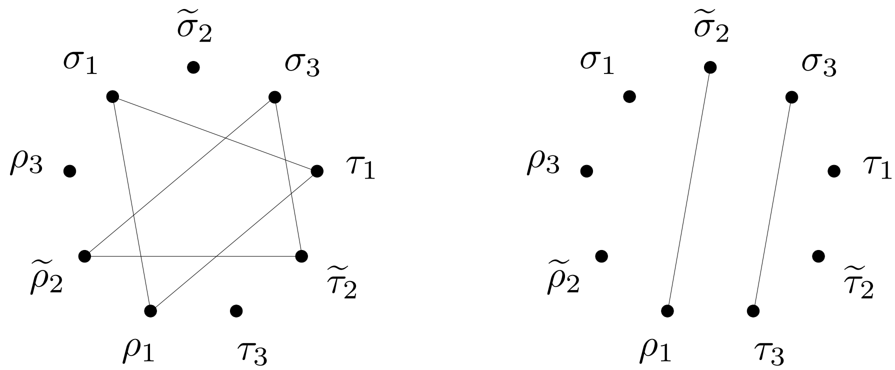

Figure 9: Two exceptions to the general rule, in a graph for 8×8-matrices. The two triangles to the left anti-commute, despite sharing no vertices. The same applies to the two edges to the right.

Firstly, the polygonal graph we start with now contains triangles in addition to the vertices and edges. Not only can the triangles not be reduced to edges, but they also anti-commute with each other when sharing *zero or two* vertices, instead of when sharing one vertex.

Secondly, the introduction of a new type of vertex (ρ) allows for more complex interplay between the edges. For edges that connect the same pair of vertex types (e.g. two $\sigma - \tau$ edges) the rule is the same as before, but for edges connecting different pairs of vertex types (as in Figure 9), they anti-commute when sharing *no* vertices.

This means that even if we perform the Goodmanson trick and turn vertices into edges, we are still stuck with at least two exceptions to the simple rule we have in four dimensions.[11] Either exception alone would already be enough to ruin the main advantage of a graphical representation - one can no longer immediately picture the structure, and no longer find maximally anti-commuting or commuting sets at a glance.

When looking at this failure from a larger perspective it is perhaps not that surprising. We know that the fundamental rule (Property 2b) is based on the number of anti-commuting pairs, or more loosely put, the number of "un-intersections" in the graph. We might be able to twist this into an intersection-based rule in four dimensions, but only because that graph is highly constrained. We can't expect it to work in higher dimensions where we have far more cases to work with.

Despite not being as good as Goodmanson's graph, the higher dimensional polygonal representation might still find some uses. It is possible to do calculations with it, albeit with an amount of work that approaches the case of doing it purely algebraically. A a specific example, one can use graphical arguments to extend a small set of anti-commuting gamma matrices into a larger one.

[11]There are other avenues, such as turning the triangles into hexagons, and the edges into quadrilaterals. However, of all the possibilities investigated, some made the commutation rules more consistent, but none made them less complicated.

A more interesting use is that the graph can serve as a visual aid for displaying sets of gamma matrices, once they have been specified by other means. As an example, Figure 10 below shows all (meaningfully distinct) maximally anti-commuting sets of 8×8-gamma matrices. In order to work with the Dirac equation in a 6-dimensional space-time, one would pick six gamma matrices from one of these sets, which also would determine the signature. For completeness, the available signatures turn out to be -6, 0 and 2, once again precisely what we expect from Formula 7.

A problem with drawing the polygonal graphs, is that in general polygons may share one or multiple sides, and it isn't always clear which polygons a specific edge belongs to. After having worked with them for a while, one develops a visual intuition - but one can also make some aesthetic changes to make them more readable. We've removed some of the shared sides - which is completely fine since what really matters is the vertices each polygon visits. We've also erased parts of some edges, to better illustrate which edges connect to form polygons, and which ones do not.

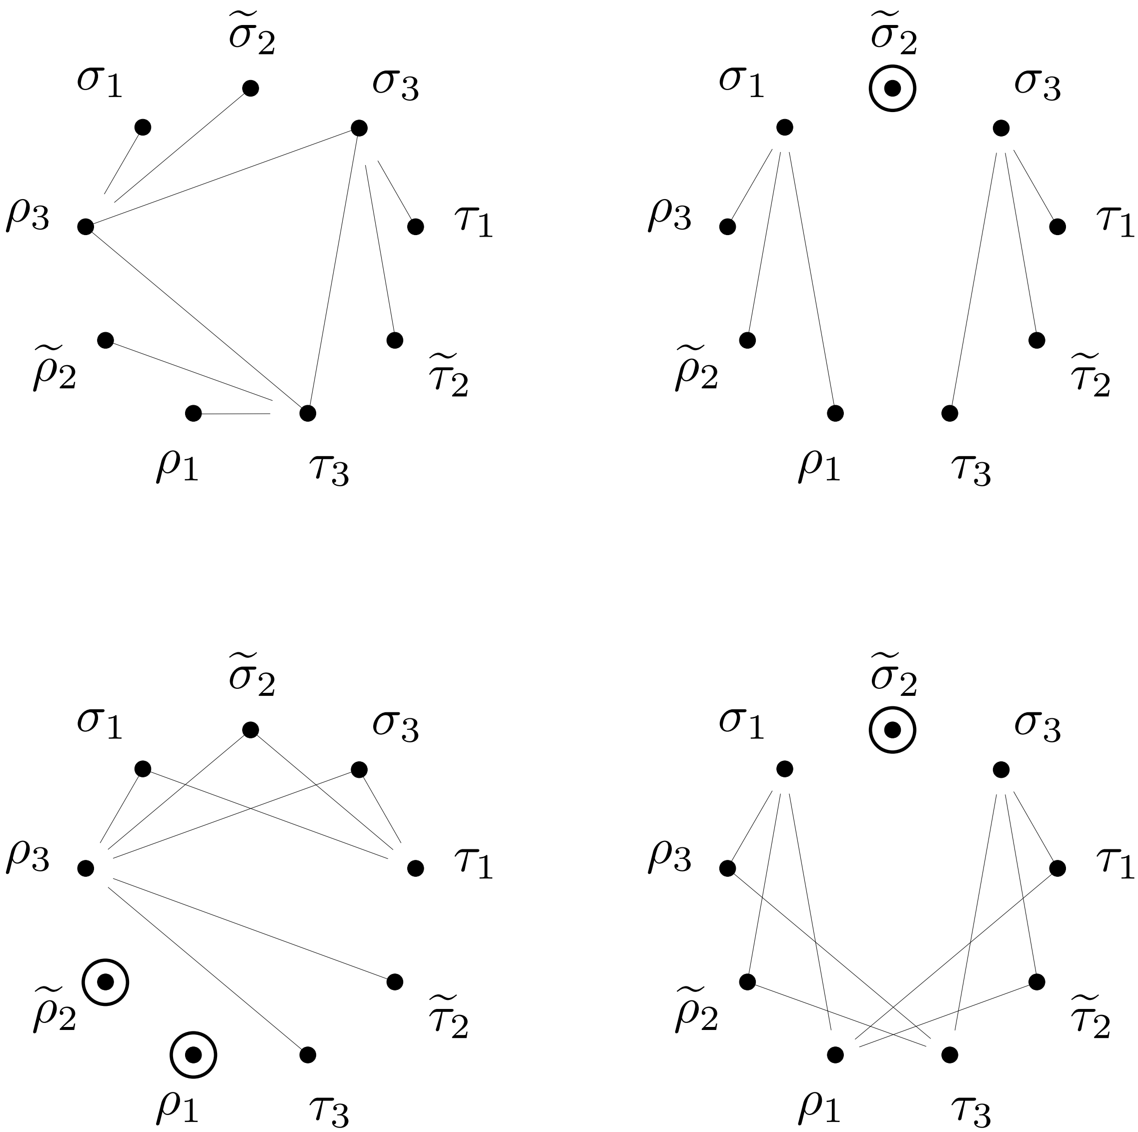

Figure 10: Polygonal graphs for all distinct maximally anti-commuting sets of 8×8 gamma matrices. In the two lower graphs some triangle sides have been removed to make the graphs more readable: in the left one we've removed $\rho_3 - \tau_1$, and in the right one $\sigma_1 - \tau_3$ and $\rho_1 - \sigma_3$.

5 Generalizing H_2 via generators (the Weyl-Heisenberg groups)

Having now analyzed one way to generalize Heisenberg groups to higher dimensions, we turn to the other. For H_2, we noted that it could be built up using three generators. One generator was essentially a phase factor (the mathematical term is that it *generated the group's center*), and then we had two generators which were "regular" group elements. When generalizing this generator definition to higher dimensions, we obtain a different type of Heisenberg groups than we studied in Section 3. These new groups are called the *Weyl-Heisenberg* groups.

The Weyl-Heisenberg groups are defined in all dimensions d. This is more general than the generalization from before, which only worked for powers of two. We will denote the Weyl-Heisenberg groups simply as H_d. The three generators we use are named ω, X and Z - analogously to $\sigma_1 = \sigma_x$ and $\sigma_3 = \sigma_z$ in two dimensions. They obey the following conditions:

$$\omega^d = X^d = Z^d = \mathbb{1}, \quad ZX = \omega XZ, \quad \omega X = X\omega, \quad \omega Z = Z\omega \tag{10}$$

Since ω commutes with every element, it is effectively a (primitive[12]) dth root of unity times the identity operator.[13] In the case of $d = 4$ we can choose $\omega = i\mathbb{1}$. Because X and Z commute up to a phase factor, we can specify a standard form $\omega^{n_1} X^{n_2} Z^{n_3}$, and as with H_2 we can always rearrange any word into this form. This also allows us to enumerate how many objects a Weyl-Heisenberg group contains. As an example, in four dimensions X^4 and Z^4 reduce to identity, which means n_2 and n_3 can each only take on the four values 0,1,2 and 3. This means that H_4 contains $4 \times 4 = 16$ objects, ignoring phase. Interestingly, this is the same amount of elements as in $H_{2\times 2}$.

5.1 The standard matrix representation

Thus far, we've discussed X and Z abstractly, without giving them an explicit matrix representation. As is usual with groups, there are many possible representations to choose from. A common choice (the standard representation) is to make the representation unitary and irreducible, and make Z be a diagonal matrix (similarly to the fact that we diagonalize σ_3 when setting up the Pauli matrices from scratch in QM).

Since $Z^4 = \mathbb{1}$, the eigenvalues, and thus the diagonal elements, must be chosen from the fourth roots of unity. The choice of eigenvalues is heavily constrained by the following relation:

$$Z|\lambda\rangle = \lambda|\lambda\rangle \quad \Rightarrow \quad Z\Big[X|\lambda\rangle\Big] = \omega XZ|\lambda\rangle = i\lambda\Big[X|\lambda\rangle\Big] \tag{11}$$

If λ is an eigenvalue to Z, then $i\lambda$ is also an eigenvalue. This forces us to include all four roots of unity as the diagonal elements. Ordering them in order of increasing phase, we obtain:

$$Z = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -i \end{pmatrix} \tag{12}$$

Diagonalizing Z greatly constrains what X can be as well. Just from equating ZX and ωXZ we find that most of the elements in X have to be zero:

$$X = \begin{pmatrix} 0 & 0 & 0 & a \\ b & 0 & 0 & 0 \\ 0 & c & 0 & 0 \\ 0 & 0 & d & 0 \end{pmatrix} \tag{13}$$

[12] The definition technically also contains that $\omega^n \neq 1$ for all $n < d$.

[13] This hopefully seems rather intuitive - but for the nit-picky reader this follows from *Schur's lemma*.

The remaining conditions (unitarity and $X^4 = \mathbb{1}$) give that these four numbers must all have an absolute value of 1, and that $abcd = 1$. All possible choices we make here will only differ up to a unitary transformation - see for instance the argument in Weyl [11] (they diagonalize X instead of Z, but the reasoning still applies in our case). We choose the simplest option which is to set all non-zero elements to 1:

$$X = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \tag{14}$$

This then completes the description of the standard matrix representation of H_4 - all other matrices can be written as words in ω, X and Z, and are thus uniquely defined.

6 Finding a common representation

Now, having specified both Heisenberg groups, we would like to seek a common ground to work with both of them simultaneously. We note that since the groups are defined independently, it is not possible to take products of abstract group elements from the two different groups. We can however work with the explicit matrix representations. However - those were also thus far chosen independently, and as might be expected this turns out to not give the best result. We have to put some thought into how we represent the groups in order to be able to work with both simultaneously in a way that connects them.

To start with, we should consider how much group structure the two groups have in common. More concretely, we seek to identify some elements from $H_{2\times 2}$ with some elements in H_4, as grounds for building a common representation. As might be expected, we are somewhat constrained in which elements we can identify. All elements in $H_{2\times 2}$ square to $\mathbb{1}$ (up to phase) - if we want a true identification, this also has to apply to the elements we choose from H_4. There are in fact only four such elements in H_4: $\mathbb{1}$, X^2, Z^2 and X^2Z^2. The best we can do is to identify these four elements with four elements from $H_{2\times 2}$.

Which elements from $H_{2\times 2}$ should we choose? It is clear that we should identify the identity element $\mathbb{1}$ with itself - it looks the same in any representation. For the remaining three elements, we note that they should have the same internal structure as X^2, Z^2 and X^2Z^2 - in particular, they must commute. There is still some freedom in which elements we choose, but fortunately all choices turn out to be functionally equivalent[14] - we are free to pick any commuting set we wish from $H_{2\times 2}$.

The four elements we identify turn out to have a very important property. They constitute a *maximally* commuting set - i.e. there is no set of five or more commuting elements in either group. The reason this is of interest, is that simultaneously diagonalizing the elements in a maximally commuting set fixes a complete representation.[15] Thus by diagonalizing the four identified elements, we can create a representation for both groups at the same time.[16]

We can save on calculation time by choosing the commuting set from $H_{2\times 2}$ with some thought. If we pick $\sigma_0 \otimes \sigma_3$, $\sigma_3 \otimes \sigma_0$ and $\sigma_3 \otimes \sigma_3$ as our elements, they are already diagonalized in the standard representation for $H_{2\times 2}$. Thus the common representation will represent all elements in $H_{2\times 2}$ as we are used to. We do however have to change the representation for H_4 in order for X^2, Z^2 and X^2Z^2 to match. In the standard representation for H_4, they look like this:

$$X^2 = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} = \sigma_1 \otimes \sigma_0, \qquad Z^2 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} = \sigma_0 \otimes \sigma_3, \qquad X^2Z^2 = \sigma_1 \otimes \sigma_3 \tag{15}$$

[14] One can show that transforming one choice into another can be done via a change of basis for the entire group $H_{2\times 2}$, thus leaving the group structure itself unchanged.

[15] This is essentially the same thing as choosing a basis in QM by finding the largest number of simultaneously independent measurements.

[16] We also have to require that the elements from $H_{2\times 2}$ and H_4 have the same eigenvalues - otherwise they would look different when diagonalized. As long as we're not adding any extra phase factors (for instance, we can't use $\widetilde{\sigma}_2$ here), this requirement is already fulfilled; all commuting sets from $H_{2\times 2}$ have the same eigenvalues as the four elements we chose from H_4.

As luck would have it, they are already pretty close to what we want - all we need is to tweak the representation so that the first factor in X^2 and X^2Z^2 is σ_3 instead of σ_1. Recall from section 2 that the Hadamard operator accomplishes just that. Specifically, if we perform a basis change using $H \otimes \mathbb{1}$, we can swap the first factor, while leaving the second factor unchanged.

Performing this change of basis on the generators for H_4 fixes the shared representation for the entire group:

$$X = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \; Z = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & i \\ 1 & 0 & 0 & 0 \\ 0 & i & 0 & 0 \end{pmatrix} \tag{16}$$

And, as was the intent, it diagonalizes the intersection of $H_{2\times2}$ and H_4:

$$\begin{gathered} X^2 = \sigma_3 \otimes \sigma_0 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad Z^2 = \sigma_0 \otimes \sigma_3 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \\ X^2Z^2 = \sigma_3 \otimes \sigma_3 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \end{gathered} \tag{17}$$

7 The mutual permutation of $H_{2\times2}$ and H_4

7.1 Permutations and the Clifford group

For any set S of unitary operators, and some unitary operator U, we say that U *permutes S under conjugation* when the following holds:

$$\forall \mathcal{O} \in S: \quad U\mathcal{O}U^{-1} \in S \tag{18}$$

We can view this as a transformation parameterized by U, which takes elements from S to other elements in S. Note that U is only defined up to a phase factor - if we multiply it by $e^{i\theta}$, the inverse gets multiplied by $e^{-i\theta}$, and these factors cancel when used to permute $\mathcal{O}$.

The set of all permuting operators, $\{U\}$, is a group. It is easy to see that two successive permutations yields another permutation, and it is also trivial to verify the remaining group axioms.

In the specific case where S is a Heisenberg group, the permuting group is called the *Clifford group*, which we will denote by Cl. Each Heisenberg group has its own Clifford group, so for $H_{2\times2}$ we have $Cl_{2\times2}$, and for H_4 we have Cl_4. Don't be tricked by the notation - we are not stating here that $Cl_{2\times2} = Cl_2 \times Cl_2$, only that $Cl_{2\times2}$ is the permuting group for $H_{2\times2}$.[17]

An important observation is that when the set S is a group, it permutes itself. With this we mean that every operator in S can be used as permuting operator U, or equivalently that S is a subgroup (subset and group) to the permuting group $\{U\}$. The validity of this follows directly from S being closed: if U is chosen as an element of the group S, then the product $U\mathcal{O}U^{-1}$ is a product of three operators in S, which must lie in S as well. Applying this observation to the Heisenberg groups, we have that:

$$H_{2\times2} \subset Cl_{2\times2}, \quad H_4 \subset Cl_4 \tag{19}$$

[17] In fact, $Cl_2 \times Cl_2$ is a *subgroup* of $Cl_{2\times2}$. It is possible to split a Clifford group into a direct product of smaller Clifford groups, but only if the smaller groups' dimensions are coprime, i.e. contain no shared factors. For instance, $Cl_{12} = Cl_4 \times Cl_3$. For a more in-depth discussion of this (and how it works analogously for the Heisenberg groups), see for instance Appendix B in Appleby [12].

7.2 The connection between $H_{2\times2}$ and H_4

Now, we are finally ready to discuss the interesting connection between $H_{2\times2}$ and H_4. They permute *each other*, and thus are also subgroups of *each other's* Clifford groups:

$$H_{2\times2} \subset Cl_4, \quad H_4 \subset Cl_{2\times2} \tag{20}$$

To be precise, this statement holds true *in our shared representation.* As we discussed in section 6 we can choose the representations for both groups independently, and thus break this connection. Regardless of representation, it is true that Cl_4 contains a subgroup that is *isomorphic* to $H_{2\times2}$, and analogously that $Cl_{2\times2}$ contains a subgroup isomorphic to H_4.

The underlying mechanisms of why this holds true are not fully understood. Since the groups involved are rather small, it is possible to prove Equation 20 by brute force. We will do this in Section 7.3.

One potentially interesting consequence of this mutual permutation is the following: for all elements A from $H_{2\times2}$ and all elements B from H_4, the following expression has to lie in the intersection of the two groups:

$$ABA^{-1}B^{-1} \in H_{2\times2} \cap H_4 \tag{21}$$

The reason for this is as follows: since $H_{2\times2}$ permutes H_4, ABA^{-1} has to be an element in H_4. Thus the whole expression is a product of two H_4 elements, which in turn must also be in H_4. An analogous argument shows that $BA^{-1}B^{-1}$ is an element in $H_{2\times2}$, meaning the expression as a whole lies in $H_{2\times2}$.

As discussed when we built the common representation in Section 6, the intersection contains only diagonal elements:

$$H_{2\times2} \cap H_4 = \{\mathbb{1},\ X^2 = \sigma_3 \otimes \sigma_0,\ Z^2 = \sigma_0 \otimes \sigma_3,\ X^2Z^2 = \sigma_3 \otimes \sigma_3\} \tag{22}$$

Thus we can also state that $ABA^{-1}B^{-1}$ must be a diagonal matrix in our shared representation.

7.3 Proving the connection

The most explicit way to prove Equation 20 is to go through all elements A from $H_{2\times2}$, and all elements B in H_4, and check for each pair that $ABA^{-1} \in H_4$ and $BAB^{-1} \in H_{2\times2}$. This would require doing $16 \times 16 \times 2 = 512$ matrix calculations, and then checking that each calculation yields an element in the appropriate group. By using the power of generators, we can reduce this number by a massive amount. This approach is best introduced by means of two examples:

Assume the matrix B can be factorized as a product B_1B_2. For the cases ABA^{-1}, where B is in the center, we can rewrite the expression in the following way:

$$ABA^{-1} = AB_1B_2A^{-1} = \Big(AB_1A^{-1}\Big)\Big(AB_2A^{-1}\Big) \tag{23}$$

Here we see that if we've already done the calculations AB_1A^{-1} and AB_2A^{-1}, and verified that they yield elements in H_4, we do not need to check the calculation ABA^{-1}. More generally, any time the middle matrix is a product, we can opt to do the calculations for all factors individually instead.

For the second example, we still assume B is a product, but this time we consider the cases BAB^{-1} where B is the outer matrices. Here we can rewrite the expression like this:

$$BAB^{-1} = B_1\Big(B_2AB_2^{-1}\Big)B_1^{-1} \tag{24}$$

We see that if $B_2AB_2^{-1}$ and $B_1AB_1^{-1}$ yield elements in $H_{2\times2}$ (for all A:s), then this must also hold true for BAB^{-1} as a whole. In general, any time the outer matrices are products, we can instead choose to do the calculations for their factors.

At this point it is hopefully clear how to proceed. We only need to do the calculations where both the outer and inner matrices are generators for their respective groups. This guarantees that the remaining expressions lie in the correct groups as well - since the matrices involved have to be products of matrices which we've already checked.

What generators should we use for the two groups? Ignoring phase, H_4 has by construction a generating set containing only the two elements X and Z. For $H_{2\times 2}$ it turns out that we can find generating sets with as few as four elements, such as the following:[18]

$$\{\sigma_0 \otimes \sigma_1,\ \sigma_1 \otimes \sigma_0,\ \sigma_0 \otimes \sigma_3,\ \sigma_3 \otimes \sigma_0\} \tag{25}$$

We've now reduced the number of calculations to perform down to $2 \times 4 \times 2 = 16$, but there is room for one final optimization. We have chosen the generating set for $H_{2\times 2}$ in such a way that it contains two elements from the intersection $H_{2\times 2} \cap H_4$ - namely $\sigma_0 \otimes \sigma_3$ and $\sigma_3 \otimes \sigma_0$. The reason for this is that any expression containing these generators must lie in the correct group, as we can see in the following way:

In the cases containing these two special $H_{2\times 2}$ generators, we can express the generators as elements in H_4, and thus write the triple products entirely in terms of X and Z. Recall that we can reorder the factors in such an expression freely, at the cost of incurring a phase factor. But this means we can move the outer two matrices in the triple product together, such that they cancel out. This means the expression is equal to purely the middle matrix, potentially with some added phase factor. Regardless, it is trivially in the same group as it started in.

What remains in the end then is to check only the 8 matrix calculations containing the other two generators for $H_{2\times 2}$, which gives a very compact way of verifying the connection between the two groups. We have not been able to spot any interesting patterns in these equations - they appear almost random - but the interested reader can find the calculations in Table 6 in the appendix.

8 The $SL(2, \mathbb{Z}_8)$ perspective

This section will dive deeper into one half of Equation 20, specifically the fact that $H_{2\times 2} \subset Cl_4$. In Section 8.1 we will explicitly express the elements in Cl_4 on a standard form that involves a different important group denoted as $SL(2, \mathbb{Z}_8)$. Since $H_{2\times 2} \subset Cl_4$ we will also be able to express all elements in $H_{2\times 2}$ on this standard form, which will be done in Section 8.2. This way we will also explicitly express the subgroup of Cl_4 corresponding to $H_{2\times 2}$.

8.1 Expressing Cl_4 on the standard form

First, let us define the group $SL(2, \mathbb{Z}_8)$. *SL* stands for *special linear*, denoting that the group consists of matrices (linear) with a determinant of 1 (special). The 2 indicates the size, i.e. 2×2-matrices. And finally, the $\mathbb{Z}_8$ tells us the nature of the elements of this matrix; they are restricted to integers modulo 8, or equivalently, the integers from 0 to 7. As a quick comment, this also means the determinant is to be calculated modulo 8.

Now, we turn to establishing the connection between Cl_4 and $SL(2, \mathbb{Z}_8)$, and in the process define the standard form of Cl_4. The underlying math is fairly complex and largely outside the scope of this paper - we will summarize the important parts, while leaving out thorough justifications for some of the statements. For a more in-depth discussion see for instance the earlier sections of Appleby [13].

The first step is to define the so called *displacement operators*. These are based on elements from the Weyl-Heisenberg groups, but we've added a particular phase factor in front:

$$D_{i,j} := \tau^{ij} X^i Z^j \tag{26}$$

[18] This is analogous to four gamma matrices generating the gamma group, as we discussed in Section 3.2.

The factor τ is, like ω, a phase factor times identity. We will fix precisely which phase factor we mean shortly, but for now we leave it unspecified. Note that the superscript ij is not an indexation - we mean to take τ to the power of the product ij.

The product of two displacement operators returns another, up to a phase factor:

$$D_{i,j}D_{k,l} = \tau^{ij+kl}X^iZ^jX^kZ^l = \omega^{jk}\tau^{ij+kl}X^{i+k}Z^{j+l} = \omega^{jk}\tau^{-jk-il}D_{i+k,j+l} \tag{27}$$

Now we specify that $\tau = \sqrt{\omega}$, which gives a determinant-like behaviour in the exponent:

$$D_{i,j}D_{k,l} = \tau^{jk-il}D_{i+k,j+l} \tag{28}$$

At this point, a word of caution on the nature of τ is necessary. By the original definition of H_4, it does not contain an element which squares to ω, meaning τ doesn't lie in H_4. The displacement operators thus live in a slightly extended version of H_4, where we've extended which phase factors are allowed. In general, this problem will occur for all even dimensions. (In odd dimensions one can use $\omega^{\frac{d+1}{2}}$ as a square root of ω, and thus avoid having to extend the group.)

The next step is to collect the indices of the displacement operators into a vector $\mathbf{p} = (i, j)$. The reason for this is that we are going to introduce linear transformations F in the index space. In practice, we consider a 2×2 matrix F which operates on $\mathbf{p}$, such that it takes a displacement operator $D_\mathbf{p}$ to a different displacement operator $D_{F\mathbf{p}}$.

Now, we state an important property without proof. Each matrix F can be associated with a unitary operator u_F such that:

$$D_{F\mathbf{p}} = u_F D_\mathbf{p} u_F^{-1} \tag{29}$$

Effectively, there exists an u_F which allows us to work with transformations in the space of the $D_\mathbf{p}$ themselves, instead of the index space.

In order for all this to work, we get some constraints on the matrix F. Firstly, the indices in $\mathbf{p}$ are only defined modulo 8,[19] and the same applies to F. Secondly, the determinant-like exponent we obtain when we take products of displacement operators (Equation 28) has consequences for the determinant of F itself - it has to be congruent with 1 modulo 8. This means that the set of matrices F is precisely the group $SL(2, \mathbb{Z}_8)$.

A quick comment on the nature of the unitary operators u_F: just as for the permuting operators we discussed in Section 7.1, they are defined only up to an arbitrary phase factor $e^{i\theta}$. The group structure of the u_F:s is equivalent to the group structure of the F:s themselves, in the sense that $F = F_1F_2$ corresponds to $u_F = e^{i\phi}u_{F_1}u_{F_2}$. Note however the added phase factor $e^{i\phi}$ - it is in general not necessary for the phase factors between different u_F:s to match up nicely. For odd (prime power) dimensions it is possible to choose the individual $e^{i\theta}$ in such a way that $e^{i\phi}$ vanishes (see for instance pp. 9–15 in Appleby [13]) - this is called *finding a faithful representation* of the u_F:s. For even dimensions this problem is more difficult, and as of yet there seems to be no answer published in the literature.

Returning to the main discussion, we note that since the displacement operators are elements in (the extended version of) H_4, the set $\{u_F\}$ permutes elements in H_4, and forms a subgroup to Cl_4. This forms the backbone of the standard form for Cl_4, however the u_F:s alone do not cover the entire group. If we multiply by elements from H_4, and a general phase, we do however obtain a standard form for all elements in Cl_4:

$$\tau^n D_{ij} u_F \tag{30}$$

[19] Since $\tau^8 = X^4 = Z^4 = \mathbb{1}$, increasing any index by 8 will leave the displacement operators unchanged.

One final thing we have to go over is the connection $F \leftrightarrow u_F$. The theory is involved, but in practice we can work with an explicit formula from Appleby [14] in order to translate F matrices into u_F operators:

$$F = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} \quad \Rightarrow \quad [u_F]_{r,s} = \frac{e^{i\theta}}{\sqrt{d}} \tau^{\beta^{-1}(\delta r^2 - 2rs + \alpha s^2)} \tag{31}$$

The matrix indices r and s count from 0 to 3. The factor β^{-1} is the multiplicative inverse of β (mod 8). In modular arithmetic, some integers will not have such an inverse, but we can always factor F into matrices which have invertible β:s, and then use the above formula on each factor. Note that this formula gives u_F for the **standard** representation of H_4, and that we thus have to perform a basis change in order to get u_F in the shared representation. Also note that we will for convenience always choose to set $e^{i\theta} = 1$.[20]

8.2 Expressing $H_{2\times 2}$ on the standard form

Now, since we've established that all elements in Cl_4 can be written on the form $\tau^n D_{ij} u_F$, and that $H_{2\times 2} \subset Cl_4$, it is natural to examine how the elements of $H_{2\times 2}$ look on this form. The method for finding this out is straightforward, if a bit cumbersome.

First we seek to find which matrix F each element A in $H_{2\times 2}$ corresponds to. In order to do that, we let A permute D_{01} and D_{10}, i.e. we calculate AXA^{-1} and AZA^{-1}. This corresponds to letting F act on the vectors $(0, 1)$ and $(1, 0)$ in the index space, which gives enough information to determine F.

Then we apply Equation 31 to transform F into u_F, and perform a basis change using $H \otimes \mathbb{1}$ to represent u_F in our shared representation. Finally, we find which D_{ij} we have to multiply with to get back the original element A, up to a phase factor τ^n. Since the matrices involved are monomial (i.e. have only one non-zero element per column and row), one can do this step reasonably easy by hand. The result is in Table 4 below:

$\otimes$	σ_0	σ_1	$\widetilde{\sigma}_2$	σ_3
σ_0	$\tau^0 D_{00}\ u_{\left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)}$	$\tau^0 D_{10}\ u_{\left(\begin{smallmatrix}3&0\\0&3\end{smallmatrix}\right)}$	$\tau^6 D_{12}\ u_{\left(\begin{smallmatrix}3&0\\0&3\end{smallmatrix}\right)}$	$\tau^0 D_{02}\ u_{\left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)}$
σ_1	$\tau^0 D_{03}\ u_{\left(\begin{smallmatrix}1&0\\2&1\end{smallmatrix}\right)}$	$\tau^5 D_{13}\ u_{\left(\begin{smallmatrix}3&0\\2&3\end{smallmatrix}\right)}$	$\tau^7 D_{11}\ u_{\left(\begin{smallmatrix}3&0\\2&3\end{smallmatrix}\right)}$	$\tau^0 D_{01}\ u_{\left(\begin{smallmatrix}1&0\\2&1\end{smallmatrix}\right)}$
$\widetilde{\sigma}_2$	$\tau^6 D_{23}\ u_{\left(\begin{smallmatrix}1&0\\2&1\end{smallmatrix}\right)}$	$\tau^3 D_{33}\ u_{\left(\begin{smallmatrix}3&0\\2&3\end{smallmatrix}\right)}$	$\tau^1 D_{31}\ u_{\left(\begin{smallmatrix}3&0\\2&3\end{smallmatrix}\right)}$	$\tau^2 D_{21}\ u_{\left(\begin{smallmatrix}1&0\\2&1\end{smallmatrix}\right)}$
σ_3	$\tau^0 D_{20}\ u_{\left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)}$	$\tau^0 D_{30}\ u_{\left(\begin{smallmatrix}3&0\\0&3\end{smallmatrix}\right)}$	$\tau^2 D_{32}\ u_{\left(\begin{smallmatrix}3&0\\0&3\end{smallmatrix}\right)}$	$\tau^4 D_{22}\ u_{\left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)}$

Table 4: Expressing elements from $H_{2\times 2}$ as products of displacement operators and unitary operators representing $SL(2, \mathbb{Z}_8)$. The arbitrary phase factor $e^{i\theta}$ within u_F has been chosen as 1.

$\otimes$	σ_0	σ_1	$\widetilde{\sigma}_2$	σ_3
σ_0	$X^0 Z^0\ u_{\left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)}$	$X^1 Z^0\ u_{\left(\begin{smallmatrix}3&0\\0&3\end{smallmatrix}\right)}$	$X^1 Z^2\ u_{\left(\begin{smallmatrix}3&0\\0&3\end{smallmatrix}\right)}$	$X^0 Z^2\ u_{\left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)}$
σ_1	$X^0 Z^3\ u_{\left(\begin{smallmatrix}1&0\\2&1\end{smallmatrix}\right)}$	$X^1 Z^3\ u_{\left(\begin{smallmatrix}3&0\\2&3\end{smallmatrix}\right)}$	$X^1 Z^1\ u_{\left(\begin{smallmatrix}3&0\\2&3\end{smallmatrix}\right)}$	$X^0 Z^1\ u_{\left(\begin{smallmatrix}1&0\\2&1\end{smallmatrix}\right)}$
$\widetilde{\sigma}_2$	$-X^2 Z^3\ u_{\left(\begin{smallmatrix}1&0\\2&1\end{smallmatrix}\right)}$	$-X^3 Z^3\ u_{\left(\begin{smallmatrix}3&0\\2&3\end{smallmatrix}\right)}$	$-X^3 Z^1\ u_{\left(\begin{smallmatrix}3&0\\2&3\end{smallmatrix}\right)}$	$-X^2 Z^1\ u_{\left(\begin{smallmatrix}1&0\\2&1\end{smallmatrix}\right)}$
σ_3	$X^2 Z^0\ u_{\left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)}$	$X^3 Z^0\ u_{\left(\begin{smallmatrix}3&0\\0&3\end{smallmatrix}\right)}$	$X^3 Z^2\ u_{\left(\begin{smallmatrix}3&0\\0&3\end{smallmatrix}\right)}$	$X^2 Z^2\ u_{\left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)}$

Table 5: Expressing elements from $H_{2\times 2}$ as products of elements in H_4 and unitary operators representing $SL(2, \mathbb{Z}_8)$. The arbitrary phase factor $e^{i\theta}$ within u_F has been chosen as 1.

[20] For the interests of this paper, a non-faithful representation is sufficient. Since we are in an even dimension ($d = 4$) it is not even known if a faithful representation is possible, as discussed earlier.

Note that the phase factors are somewhat erratic when expressed like this. If we instead of D_{ij} write out the explicit powers of X and Z we get the much nicer relationship in Table 5 on the previous page.

The position of the minus signs in Table 5 is dependent on our choice of $\widetilde{\sigma}_2$. If we work with $-\widetilde{\sigma}_2$ instead, the signs would appear in the second column instead of the second row. Working with the original σ_2 would instead have given anti-symmetrically distributed i:s and $-i$:s.

To end this section, we reiterate an important point from Section 7 - even if the identification between elements from $H_{2\times 2}$ and the elements from Cl_4 is dependent on our common choice of representation, we can still generally state that $H_{2\times 2}$ is isomorphic to the specific subgroup of Cl_4 which is illustrated in Table 4 & 5.

9 Conclusion

When working in a very reasonable shared representation, $H_{2\times 2}$ and H_4 permute each other's elements, or equivalently, lie in each other's Clifford groups. Regardless of representation, the Clifford groups contain subgroups isomorphic to each other's Heisenberg groups. The subgroup of Cl_4 has been explicitly defined in terms of the shared representation.

Since this was essentially shown using a brute force method, much remains to be done in terms of understanding this connection. Is this relationship all there is - a mathematical fluke - or does it arise from some even deeper connection between the two Heisenberg groups? How sensitive is the statement to the choice of the common representation? Is it specifically a fact in four dimensions, or does it hold in higher dimensions as well?

A specific question of interest, for which there wasn't time in this thesis, is the existence of an analogue to the $SL(2,\mathbb{Z}_8)$ perspective. In the same way that the elements of Cl_4 can be expressed on a standard form $\tau^n D_{ij} u_F$, the elements of $Cl_{2\times 2}$ can also be expressed on a standard form. Even if there is some difference in the details,[21] we should still be able to express the elements of H_4 on the standard form for $Cl_{2\times 2}$, and thus specify the precise subgroup of $Cl_{2\times 2}$ which is isomorphic to H_4.

On the topic of graphical representation of gamma matrices: There does not appear to be a fully satisfactory way of extending Goodmanson's graph to higher dimensions. The partial results obtained might be of help in some applications, but we've been unable to replicate the key feature of immediately spotting anti-commuting and commuting sets. Of course, the possibility of a radically different graph representation which retains this property remains.

Acknowledgements

I wish to express my thanks to Ingemar Bengtsson for his excellent supervision: his thorough explanations of the theory, practical tips on the mathematical methods and writing, as well as proof-reading all the drafts. I also wish to thank Marcus Appleby for pointing out the important fact that building the shared representation requires that the eigenvalues match up between the two groups (Footnote 16). Finally, I wish to thank my family and friends for their support, especially Elle for early motivational support as well as pointing out some lacking explanations, and Sara for helping me stay motivated during the final edits.

[21] For instance, $Cl_{2\times 2}$ is commonly built from generators, in contrast to the more explicit definition of $SL(2,\mathbb{Z}_8)$.

10 Appendix

$$
\begin{array}{rcl}
\hline
X\ (\sigma_1 \otimes \sigma_0)\ X^{-1} & = & -\sigma_1 \otimes \sigma_3 \\
X\ (\sigma_0 \otimes \sigma_1)\ X^{-1} & = & \sigma_3 \otimes \sigma_1 \\
Z\ (\sigma_1 \otimes \sigma_0)\ Z^{-1} & = & \sigma_1 \otimes \sigma_0 \\
Z\ (\sigma_0 \otimes \sigma_1)\ Z^{-1} & = & i\sigma_0 \otimes \widetilde{\sigma}_2 \\
\hline
(\sigma_1 \otimes \sigma_0)\ X\ (\sigma_1 \otimes \sigma_0)^{-1} & = & XZ^2 \\
(\sigma_0 \otimes \sigma_1)\ X\ (\sigma_0 \otimes \sigma_1)^{-1} & = & X^3 \\
(\sigma_1 \otimes \sigma_0)\ Z\ (\sigma_1 \otimes \sigma_0)^{-1} & = & Z \\
(\sigma_0 \otimes \sigma_1)\ Z\ (\sigma_0 \otimes \sigma_1)^{-1} & = & iZ^3 \\
\hline
\end{array}
$$

Table 6: The matrix calculations proving Equation 20. Performed in the common representation obtained at the end of Section 6.

References

[1] David M. Goodmanson (1996). "A graphical representation of the Dirac algebra". *Am. J. Phys.* 64, pp. 870–880.

[2] Daniel Gottesman (1998). "The Heisenberg Representation of Quantum Computers". arXiv: quant-ph/9807006.

[3] Josef Maria Jauch and Fritz Rohrlich (1976). *The Theory of Photons and Electrons : The Relativistic Quantum Field Theory of Charged Particles with Spin One-half.* New York: Springer-Verlag.

[4] Tommy Ohlsson (2000). "Dynamics of quarks and leptons : Theoretical Studies of Baryons and Neutrinos". PhD thesis. KTH, Physics, pp. 61–75.

[5] *CUORE experiment.* URL: https://cuore.lngs.infn.it/en.

[6] *LEGEND experiment.* URL: https://legend-exp.org/.

[7] *nEXO experiment.* URL: https://nexo.llnl.gov/.

[8] *SNO+ experiment.* URL: https://snoplus.phy.queensu.ca/about/neutrinoless-double-beta-decay.html.

[9] Peter G. O. Freund (1986). *Introduction to Supersymmetry.* Cambridge University Press, pp. 15–19.

[10] Arthur Stanley Eddington (1935). *New Pathways in Science.* Cambridge University Press.

[11] Hermann Weyl (1950). *The Theory of Groups and Quantum Mechanics.* New York: Dover, pp. 277–279.

[12] D. Marcus Appleby et al. (2012). "The monomial representations of the Clifford group". *Quantum Information and Computation* 12, pp. 404–431. arXiv: 1102.1268 [quant-ph].

[13] D. Marcus Appleby (2009). "Properties of the extended Clifford group with applications to SIC-POVMs and MUBs". arXiv: 0909.5233 [quant-ph].

[14] D. Marcus Appleby (2005). "Symmetric informationally complete–positive operator valued measures and the extended Clifford group". *Journal of Mathematical Physics* 46, pp. 7–10.