The Rubik's Cube and Minimal Representations of Split Group Extensions

Charles Daly and Justin Kingsnorth

July 30, 2025

Abstract

In this paper, we examine the groups G_2 and G_3 associated to the 2×2 and 3×3 Rubik's cubes. We express G_2 and G_3 in terms of familiar groups and exhibit a split homomorphism $\psi : G_3 \longrightarrow G_2$ to prove that G_2 embeds inside G_3 as a subgroup. In addition, we prove several results bounding the dimensions of minimal faithful representations of finite abelian groups split by some complementary subgroup. We then employ these results to determine the minimal faithful dimensions of G_2 and G_3 over both $\mathbb{C}$ and $\mathbb{R}$. We find that G_2 has minimal dimension 8 over $\mathbb{C}$ and 16 over $\mathbb{R}$, and that G_3 has minimal dimension 20 over $\mathbb{C}$ and 28 over $\mathbb{R}$.

Contents

1 Introduction

The Rubik's cube is an excellent illustration of the principles of finite group theory. It offers clear, tangible examples of concepts such as normal subgroups and group homomorphisms, as well as a compelling motivation for tools like the semi-direct product. Moreover, it is complex enough to resist brute computation, so that we are best equipped to understand it through its underlying algebraic structure.

In the same vein, the Rubik's cube provides a clear exposition of the representation theory of finite groups. Despite its apparent complexity, the cube admits natural, low-dimensional representations that reflect both the algebraic structure of the cube group and the geometry of the cube itself. The proof that these representations have minimal dimension demonstrates many foundational techniques in representation theory, and we will find that it generalizes easily to a large class of related groups.

In this paper, we will investigate the groups G_2 and G_3 associated to the 2×2 and 3×3 Rubik's cubes. We will study G_2 through the homomorphism $\phi : G_2 \longrightarrow S_8$ induced by suppressing the colors on the 2×2 cube, and we will study G_3 through the homomorphism $\psi : G_3 \longrightarrow G_2$ induced by suppressing the edge pieces of the 3×3 cube. We will then further decompose the kernels of ϕ and ψ, so that G_2 and G_3 can be expressed entirely in terms of familiar groups. In Theorem 3.10 we will prove that ψ is split, so that G_2 can be realized as a subgroup of G_3.

In addition, we will consider faithful representations of finite abelian groups and prove that the minimal dimension of these representations over both the complex and real numbers is related to the invariant factor decomposition of the group as in Theorem 4.2. We then consider split extensions by finite abelian groups whose splitting homomorphism is faithful. In Theorems 4.6 and 4.4, we relate the minimal dimension of faithful representations over both the real and complex numbers to the minimal permutation degree of the complementary subgroup.

We conclude by using our results to calculate the minimal dimension of faithful representations of both G_2 and G_3 over $\mathbb{C}$ and $\mathbb{R}$. We will show that any faithful representation of G_2 must have dimension at least 8 over $\mathbb{C}$ and 16 over $\mathbb{R}$, and we will construct representations that achieve these bounds. We will then repeat this program for G_3, and we will find that its faithful representations have minimal dimension 20 over $\mathbb{C}$ and 28 over $\mathbb{R}$.

The paper is organized as follows. In Section 2, we determine the group structure of G_2. In Section 3, we do the same for G_3. In Section 4, we establish some bounds on the minimal faithful dimensions of complex and real representations of both finite abelian groups and split extensions by them. In Section 5, we consider the faithful representations of G_2 and G_3 respectively.

Acknowledgements

Both authors extend their gratitude to Brown University's Summer/Semester Projects for Research, Internship, and Teaching (SPRINT) and Undergraduate Teaching and Research Awards (UTRA) for both their financial support and dedication to collaboration amongst junior faculty and undergraduates. The first author would like to thank Dani Kaufman for helping spark interest in this problem through her deep understanding of the group structure of the Rubik's cube. He would also like to thank both Philip Eberhart and Myles Miller for their insights into the puzzle, particularly Myles for showing him the 'bird configuration' of the Rubik's Pyramorphix. The second author thanks his algebra teachers, Isabel Vogt and Wei Zhining, for their instruction and guidance over the course of his undergraduate education. He also thanks his parents for their constant encouragement, their celebration of curiosity, and their steadfast commitment to his studies.

2 Group Structure of the 2×2 Rubik's Cube

In this section, we determine the group structure of G_2 through the map $\phi : G_2 \longrightarrow S_8$ recording permutations of the corner pieces. We prove that this map is surjective in Proposition 2.1. Next, we turn to the kernel K of ϕ, or the subgroup of $\mathbb{Z}_3^8$ describing rotations of the corners in place. In order to determine K, we introduce the concept of local orientation and prove that it induces a $\mathbb{Z}_3$ invariant preserved by G_2. We then prove in Proposition 2.5 that K is maximal under this invariant. We close by showing that the sequence $K \longrightarrow G_2 \longrightarrow S_8$ is split in

Proposition 2.6, and by enumerating the normal subgroups of G_2 in Propositions 2.7–2.9.

2.1 Setup of the Problem

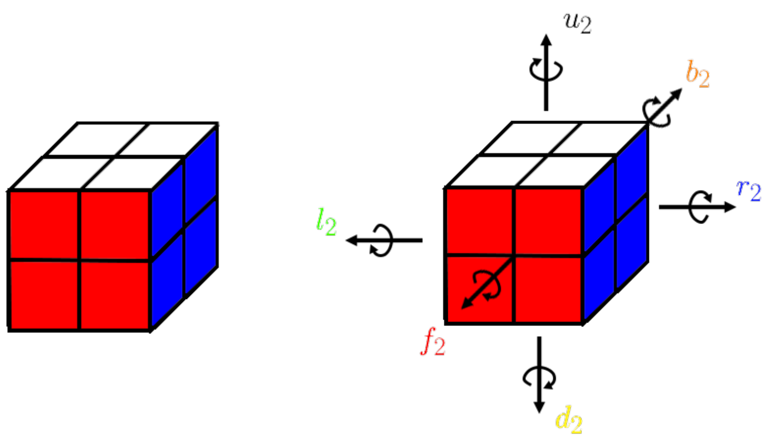

Figure 1: The $2 \times 2 \times 2$ Rubik's cube; generators of the 2×2 cube group G_2.

The $2 \times 2 \times 2$ Rubik's cube, here also called "the 2×2 cube", is a cube subdivided into eight smaller pieces called "cubelets". The six faces of the cube are each covered by four stickers of the same color, one for each cubelet on that face. (See Figure 1).

Each face can be rotated in 90-degree increments while the opposite face is held fixed, and these rotations induce permutations of the stickers on the cube. The group associated to the 2×2 cube, called G_2, can be realized through these permutations. It is generated by six elements, as shown in Figure 1:

- u_2: rotate the top (white) face 90 degrees clockwise
- d_2: rotate the bottom (yellow) face 90 degrees clockwise
- f_2: rotate the front (red) face 90 degrees clockwise
- b_2: rotate the back (orange) face 90 degrees clockwise
- l_2: rotate the left (green) face 90 degrees clockwise
- r_2: rotate the right (blue) face 90 degrees clockwise

By convention, these rotations are clockwise for an observer looking at the relevant face outside the cube. The colors associated with each face are for the cube in its solved state. Unless we are also discussing the 3×3 cube group G_3, we will suppress the subscripts on these generators.

2.2 The Quotient Map

There is a natural group homomorphism $\phi : G_2 \longrightarrow S_8$ induced by the action of G_2 on the positions of the eight cubelets. To construct this homomorphism, we read off the corner positions of the cube in the following order: top-front-left, top-front-right, top-back-left, top-back-right, bottom-front-left, bottom-front-right, bottom-back-left, bottom-back-right. We then assign a number to each position according to its place in the list, as shown in Figure 2. With these numbers, we find that

$$\begin{aligned} &\phi(u) = (1342); \phi(d) = (5687); \phi(f) = (1265); \\ &\phi(b) = (3784); \phi(l) = (1573); \phi(r) = (2486) \end{aligned} \tag{1}$$

so that, for example, u takes the cubelet at position 1 to the cubelet at position 3. We can then determine where ϕ takes an arbitrary element $g \in G_2$.

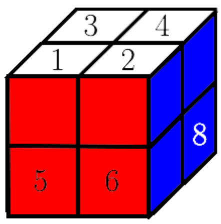

Figure 2: Labels for the corner positions of the 2×2 cube. Position 7 is hidden in the bottom-back-left.

Proposition 2.1. *The map $\phi : G_2 \longrightarrow S_8$ is surjective.*

Proof. It is enough to construct a single transposition of two adjacent cubelets; one of two diagonally opposite cubelets on the same face; and one of two cubelets connected by a long diagonal. Any transposition of cubelets falls into one of these three categories, and we can interchange transpositions within each category by rotating the cube as a whole.

We begin with two adjacent cubelets, represented by the transposition (34). We define $g_1 = rdr^{-1}f^{-1}$, $g_2 = g_1 u g_1 u^{-1}$, and $t_1 = u^{-1}g_2$. Using the listing (1) of the images of the generators under ϕ, we find that $\phi(t_1) = (34)$, so that t_1 transposes the top-back-left and top-back-right cubelets. Figure 3 gives an illustration.

It should then be clear that $t_2 = lt_1l^{-1}$ and $t_3 = l^2t_1l^{-2}$ provide the other two classes of transposition. Since these three transpositions give a way to produce an arbitrary transposition of cubelets, and since S_8 is generated by transpositions, this completes the proof that ϕ is surjective. □

Figure 3: The move t_1 described in Proposition 2.1, shown from the front and back faces of the 2×2 cube.

2.3 Kernel of the Quotient

The kernel K of ϕ is the subgroup of G_2 that changes the orientations of the cubelets while leaving their positions fixed. Since there are eight cubelets in total and three orientations for each cubelet, and since rotations of the cubelets commute, K must be a subgroup of $\mathbb{Z}_3^8$.

In this section, we will introduce the concept of local orientation and show that it produces a $\mathbb{Z}_3$ algebraic invariant preserved by G_2. We will then show that K is the maximal subgroup of $\mathbb{Z}_3^8$ that respects this invariant.

Definition 2.2. The *local orientation* of a position i is an element $s_i \in \mathbb{Z}_3$ defined by an orthonormal right-handed basis $\hat{x}_i, \hat{y}_i, \hat{z}_i$ attached to the cubelet at that position. These bases travel and rotate with the cubelets as they are acted on by G_2. s_i is the number of counterclockwise rotations by $2\pi/3$ taking the basis at position i in the solved state to the basis at i in the current state.

Figure 4 illustrates this definition. Each cubelet is arbitrarily assigned such a basis in the solved state of the cube, and the change in local orientation at each position depends on this assignment. As an example, the generator u preserves local orientation only if the top face of the cube contains four basis vectors of the same type. However, we can show that the sum of local orientations is an algebraic invariant.

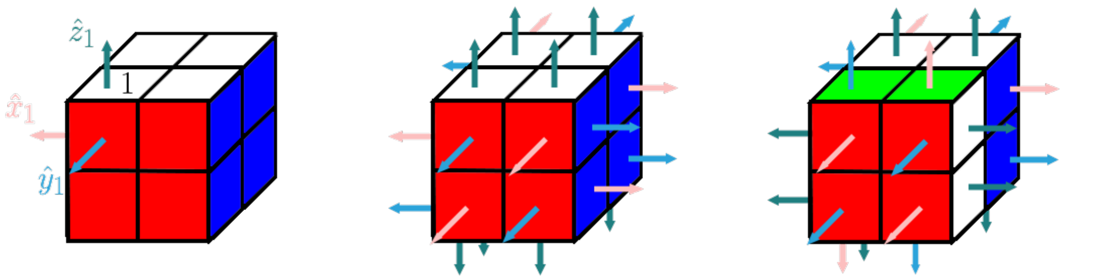

Figure 4: A basis for local orientation at position 1; bases for local orientations at each position; the change in local orientation under f. In the right-most figure, the changes in local orientations are $([1], [2], [0], [0], [2], [1], [0], [0])$.

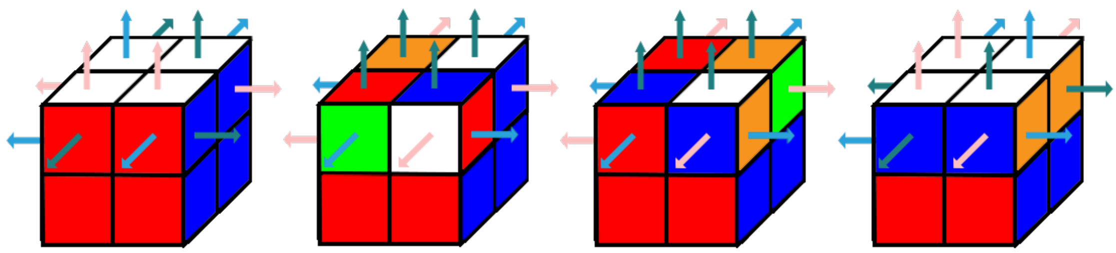

Figure 5: An illustration of Proposition 2.3. From left, the starting configuration; applying the rotation in place; applying u; inverting the rotation in place to recover u alone.

Proposition 2.3. *The sum $s \in \mathbb{Z}_3$ of local orientations $s_1, \ldots, s_8$ is independent of basis and preserved by G_2.*

Proof. It is enough to show this for an arbitrary generator of G_2, say u. We will assume that the cube is in some arbitrary state and that we have already specified bases $\hat{x}_i, \hat{y}_i, \hat{z}_i$. We can ignore the bottom face of the cube, since u fixes the positions and orientations of the bottom cubelets. We illustrate our proof in Figure 5.

First, rotate each top cubelet in place so that the basis vectors $\hat{x}_i$ lie on the top face. This changes s to some quantity s'. Next, apply u to the cube. This fixes s', as the bottom cubelets are fixed and the $\hat{x}_i$ vectors of the top cubelets remain on the top face. Finally, rotate each cubelet by the inverse of its first rotation. This changes s' back to s, so that s is preserved by this sequence of moves. This sequence is equivalent to u alone, in the sense that they induce the same permutation on the stickers, so u must fix s as well. A similar argument will show that s is preserved by each of the remaining generators, so that s is invariant under

G_2. □

Figure 6 illustrates this constraint: a clockwise twist at position 2 requires a counterclockwise twist at position 4. We will see shortly that K is maximal under the invariant s, so that the move in the figure is in fact an element of K.

We can write elements $k \in K$ as vectors $(k_1, \ldots, k_8) \in \mathbb{Z}_3^8$, where k_i denotes the change in local orientation at position i induced by k. By Proposition 2.3, $k_1 + \ldots + k_8 = [0]$. Using this notation, the conjugation action of G_2 on K can be written in a straightforward way. If $k = (k_1, \ldots, k_8)$ and $\phi(g) = \sigma$, one can show that

$$gkg^{-1} = (k_{\sigma^{-1}(1)}, \ldots, k_{\sigma^{-1}(8)}) \text{ for } \phi(g) = \sigma \tag{2}$$

We will use this fact to show that K is maximal under the $\mathbb{Z}_3$ invariant. Before we do this, we will introduce some notation that will encode this property concisely.

Definition 2.4. Let k and m be integers greater than 1 and 0 respectively. The group $\mathbb{Z}_{k,0}^m$ is the kernel of the homomorphism $\mathbb{Z}_k^m \longrightarrow \mathbb{Z}_k$ which takes the m-tuple of elements in $\mathbb{Z}_k$ to their sum.

Proposition 2.5. $K = \mathbb{Z}_{3,0}^8$.

Proof. Since K respects the algebraic invariant, we already know that $K \subseteq \mathbb{Z}_{3,0}^8$. To get equality, we first claim that the element $k = (u^2r^{-1}u^2r)(ur^{-1}ur)$ is nontrivial and contained in K. We do not consider it profitable to show this explicitly, as it would require a notation for moves as permutations of the stickers which we would not use elsewhere in the paper. The interested reader may check this statement on an online Rubik's cube simulator.

We now treat k as an arbitrary nontrivial element of K and consider its components $k_1, \ldots, k_8$. Since $k_1 + \ldots + k_8 = [0]$ and the k_i are not all $[0]$, there must be distinct i, j such that $k_i \neq k_j$.

Since ϕ is a surjective map, we can find $g \in G_2$ such that $\phi(g) = (ij)$. We then consider $k' = gkg^{-1}$. We find that $k'_i = k_j, k'_j = k_i$, and $k'_l = k_l$ for all other l, so that the difference $\tilde{k} = k - k'$ is trivial in all components except i and j. Moreover, $\tilde{k}_i + \tilde{k}_j = [0]$.

With transpositions similar to those induced by g, we can use $\tilde{k}$ to construct elements $e_1 = ([1], [2], [0], \ldots, [0])$, $e_2 = ([1], [0], [2], [0], \ldots, [0])$, and so on up to e_7. These elements form a basis for a seven-dimensional $\mathbb{Z}_3$-vector space $V \subseteq K \subseteq \mathbb{Z}_{3,0}^8$.

Since $\mathbb{Z}_{3,0}^8$ is also seven-dimensional, the inclusions become equalities, so that $K = \mathbb{Z}_{3,0}^8$ as claimed. □

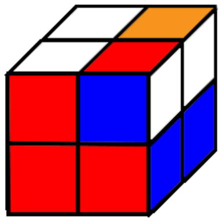

Figure 6: A typical element of K.

2.4 The Split Exact Sequence $K \longrightarrow G_2 \longrightarrow S_8$; Normal Subgroups

In this section, we will close our study of the 2×2 cube's group structure by proving two results. We will first show that the short exact sequence $K \longrightarrow G_2 \longrightarrow S_8$ induced by ϕ is split. We will then show that G_2 has only two nontrivial normal subgroups: K, the kernel of ϕ, and L, the set of elements expressible as an even number of generators.

Proposition 2.6. *The short exact sequence $K \longrightarrow G_2 \longrightarrow S_8$ is split.*

Proof. In order to construct a copy of S_8 in G_2, we will choose a basis $\hat{x}_i, \hat{y}_i, \hat{z}_i$ for a local orientation at each position. We then define $H \subseteq G_2$ as the subgroup of elements that fix local orientation pointwise for these bases. We claim that $\phi|_H$ is an isomorphism.

$\phi|_H$ is injective, since any element of H that fixes the positions of the cubelets is trivial. We can also see that $\phi|_H$ is surjective as follows. Choose any $\sigma \in S_8$. Since ϕ is surjective, we can find some element $g \in G_2$ such that $\phi(g) = \sigma$. By Proposition 2.3, g preserves the algebraic invariant s, and the same is true of all elements of H. It follows that g differs from an element of H by some rotation of the cubelets that preserves s, and by Proposition 2.5, we know that some element $k \in K$ will carry out this rotation. We then have that $kg \in H$ and $\phi(kg) = \phi(g) = \sigma$, so that $\phi|_H$ is surjective. Then $\phi|_H$ is an isomorphism, so that the sequence is split as we claimed. □

G_2 is then isomorphic to the semi-direct product $K \rtimes_{\tilde{\phi}} H$, where $\tilde{\phi} : H \longrightarrow \mathrm{Aut}(K)$ is the conjugation action given by (2). Since we have a natural isomorphism between H and S_8, we can equivalently write the semi-direct product as

$K \rtimes_{\tilde{\phi}} S_8$, and we will use this notation from here on. We illustrate a typical element of H in Figure 7.

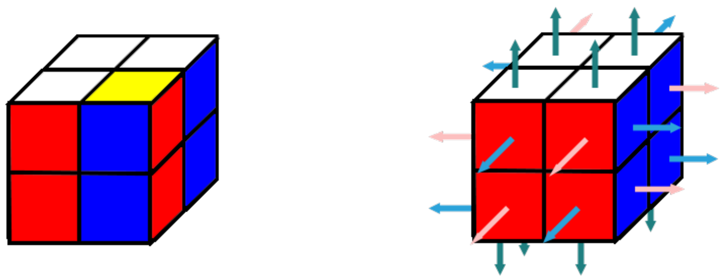

Figure 7: The element of H corresponding to the transposition (26), defined using the local orientation on the right.

Using ϕ, we will now list the non-trivial normal subgroups $N \subseteq G_2$. In addition to K, we will identify the set $L \subseteq G_2$ of elements expressible as an even number of generators. We begin by noting that for any nontrivial normal subgroup $N \subseteq G_2$, $\phi(N)$ must be isomorphic to the trivial group, A_8, or S_8.

Proposition 2.7. *If $\phi(N)$ is trivial, $N = K$.*

Proof. Since $\phi(N)$ is trivial, $N \subseteq K$. Since N is nontrivial, it contains some nontrivial element $k \in K$. Since N is closed under conjugation by G_2, we can apply the arguments of Proposition 2.5 to show that $K \subseteq N$. We conclude that if $\phi(N)$ is trivial, $N = K$. □

Proposition 2.8. *If $\phi(N) \simeq S_8$, then $N = G_2$.*

Proof. We will show that N has nontrivial kernel under ϕ, so that $K \subseteq N$ as before. To see this, take $n \in N$ such that $\phi(n) = (123)$, and take $k = ([1], [2], [0], \ldots, [0]) \in K$. Since N is normal in G_2, $kn^{-1}k^{-1} \in N$, so that $nkn^{-1}k^{-1} \in N$. By the formula (2), we know that $nkn^{-1} = ([0], [1], [2], \ldots, [0]) \in K$, so that $nkn^{-1}k^{-1} = ([2], [2], [2], \ldots, [0]) \in K$. Then $N \cap K$ is nontrivial, so that $K \subseteq N$ and $N = G_2$. □

Proposition 2.9. *If $\phi(N) \simeq A_8$, then $N = L$.*

Proof. By the same argument as in Proposition 2.8, we find that $N \cap K$ is nontrivial, so that $K \subseteq N$. Since $\phi(N) \simeq A_8$ is normal in S_8, we can define a map $\psi : S_8 \longrightarrow \mathbb{Z}_2$ by the quotient $S_8/\phi(N)$. Since $K \subseteq N$, it follows that $N = \ker \psi \circ \phi$.

To see that $N = L$, consider the image of any generator a of G_2 under $\psi \circ \phi$. From the formula (1), we can see that $\phi(a)$ has odd sign for all a, so that

$(\psi \circ \psi)(a) = [1] \in \mathbb{Z}_2$. If we then write some element $g \in G_2$ as a word $a_1 \dots a_n$, we see that n is even if and only if $g \in N$. This shows that the parity of n is well-defined and that $N = L$, as claimed. □

As a final observation, we note that G_2 is centerless, since any central element would generate a cyclic normal subgroup of G_2. We also note that the abelianization of G_2 is $\mathbb{Z}_2$, since this is the only nontrivial abelian quotient of G_2 by the propositions above. This closes our analysis of the 2×2 cube group, and we now turn to the 3×3 cube.

3 Group Structure of the 3×3 Cube

In this section, we study G_3 through the surjective map $\psi : G_3 \longrightarrow G_2$ induced by taking like generators to like generators. Since we have already studied G_2, we will concentrate on the kernel N of ψ consisting of moves that fix the corners of the cube. As in Section 2, we will study N through the map $\beta : N \longrightarrow S_{12}$ recording permutations of the edge pieces. We will find in Proposition 3.4 that $\beta(N) = A_{12}$. We then turn as before to the kernel M of N, or the subgroup of $\mathbb{Z}_2^{12}$ describing rotations of the edges in place. We review the concept of local orientation, find that it induces a $\mathbb{Z}_2$ invariant preserved by G_3, and show in Proposition 3.7 that M is maximal under this invariant. As a final group-theoretic result, we show in Theorem 3.10 that the map $\psi : G_3 \longrightarrow G_2$ is split, so that G_2 can be realized as a subgroup of G_3.

3.1 Setup of the Problem

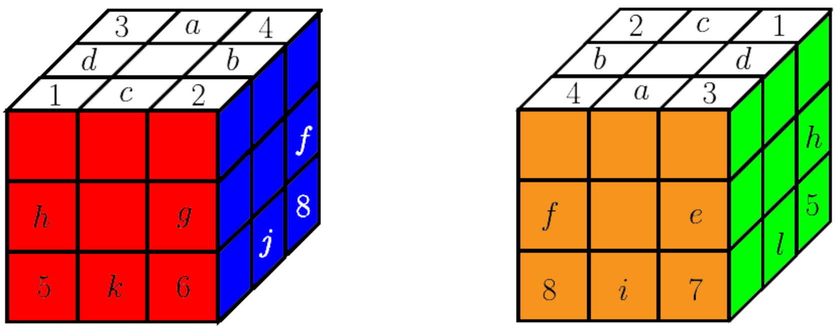

Figure 8: The $3 \times 3 \times 3$ Rubik's cube, with labels for the corner and edge positions.

The $3 \times 3 \times 3$ Rubik's cube, also called here the "3×3 cube", is divided into 26 cubelets of equal size. Each face of the 3×3 cube is three cubelets high and three wide. As before, the outer faces of each cubelet are marked with a colored sticker, so that each face of the 3×3 cube is marked with nine stickers of the same color. We will use the same color conventions as in Section 2.1. Among the cubelets, there are 8 corner pieces, each shared between three faces; 12 edge pieces, each shared between two; and 6 center pieces which belong to one face only. Figure 8 shows the cube in its solved state.

We will realize the 3×3 cube group G_3 through permutations on the stickers. As before, G_3 is generated by elements u_3, d_3, b_3, f_3, r_3, and l_3 which rotate the specified face by 90 degrees clockwise and leave the remaining cubelets fixed. These moves send corners to corners and edges to edges, and they fix the center pieces pointwise.

We will keep the listing of corner positions that we introduced in Section 2.2. As shown in Figure 8, we will also assign letters a through l to the edge positions in this order: top-back, top-right, top-front, top-left, back-left, back-right, front-right, front-left, bottom-back, bottom-right, bottom-front, bottom-left. We can then construct a homomorphism $\alpha : G_3 \longrightarrow S_{12} \times S_8$ by determining the action of each generator on the corner and edge pieces. We find that

$$\begin{aligned}&\alpha(u_3) = ((abcd), (1342)); \alpha(d_3) = ((ilkj), (5687));\\ &\alpha(b_3) = ((aeif), (3487)); \alpha(f_3) = ((cgkh), (1265));\\ &\alpha(r_3) = ((bfjg), (2486)); \alpha(l_3) = ((dhle), (1573))\end{aligned} \tag{3}$$

so that, for example, r_3 sends the corner cubelet at position 2 to position 4 and the edge cubelet at position b to position f.

3.2 The Quotient Map $\psi : G_3 \longrightarrow G_2$; Its Kernel

Given the generators $u_3, \ldots, l_3$ of G_3 and $u_2, \ldots, l_2$ of G_2, we can construct a group homomorphism $\psi : G_3 \longrightarrow G_2$ by sending u_3 to u_2, d_3 to d_2, and so on for the remaining generators. In effect, ψ suppresses the positions and orientations of the edges. We can see that ψ is a homomorphism by recognizing that for each generator g_3 of G_3 and the associated generator g_2 of G_2, g_3 and g_2 induce the same permutation on the corner stickers.

The kernel N of ψ is the normal subgroup of G_3 that acts on edge pieces while leaving the corners fixed. We will analyze N in the same way that we analyzed G_2 in Section 2, through the map $\beta : G_3 \longrightarrow S_{12}$ recording the positions of the edges. We will show that β induces the sequence $\mathbb{Z}^2_{12,0} \longrightarrow N \longrightarrow A_{12}$, recalling Definition 2.4 for the group $\mathbb{Z}^m_{k,0}$. We begin with the following proposition.

Proposition 3.1. *Let $\{g_{2,i}\} = \{{u_2}^2, {r_2}^{-1}, \ldots, r_2\}$ be the sequence of generators of G_2 that provided a nontrivial element of K in Proposition 2.5, and let $\{g_{3,i}\}$ be the associated sequence of generators of G_3. If $h \in G_3$ is the composition of the elements $g_{3,i}$, then $\alpha(h) = ((abc), 1)$.*

Proof. By Proposition 2.5 and the map $\psi : G_3 \longrightarrow G_2$, h will fix the positions of the corners, so it is enough to consider its action on the positions of the edges. We recall from Proposition 2.5 that $h = (u_3)^2(r_3)^{-1}(u_3)^2 r_3 u_3 (r_3)^{-1} u_3 r_3$. Referring to the map α defined above, it is straightforward to check that h induces the permutation (abc) on the edges, as claimed. Figure 9 provides an illustration. □

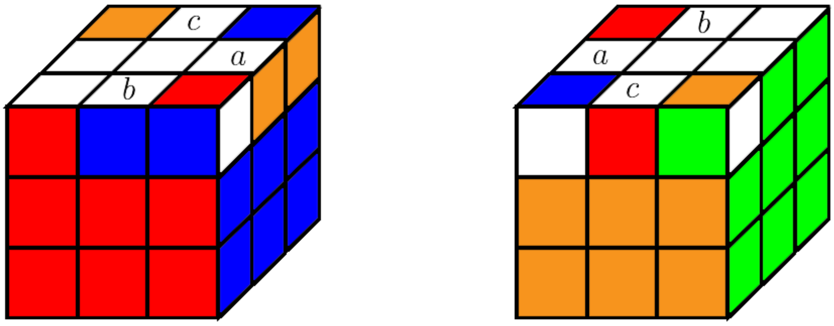

Figure 9: The move h described in Proposition 3.1, shown from the front and back of the 3×3 cube.

By applying this process to different faces of the cube, we can find moves that induce a 3-cycle on any three edges with a common face. We will call this set of moves H. H itself is not contained in N, since we know by Proposition 2.5 that its elements rotate the corners in place. However, $[h, g] \in N$ for any $h \in H$ and $g \in H \cup N$, since rotations of the corners commute and H and N permute the corners trivially. In the following propositions, we will use H to determine the group structure of N.

Proposition 3.2. $\beta(N) \supseteq A_{12}$.

Proof. We define a subset S of the edge pieces of the cube such that $\beta(N)$ contains the alternating group on the edge pieces in S. We first claim that a, b, and f are in S. By Proposition 3.1, we can find elements $h_1, h_2 \in H$ such that $\beta(h_1) = (abc)$ and $\beta(h_2) = (afe)$. As explained above, $[h_1, h_2] \in N$. $\beta([h_1, h_2]) = (abf)$, so that $\beta(N)$ contains the alternating group on these letters. We illustrate this move in Figure 10.

We then introduce a procedure for adding an edge piece e_1 to S when S contains two edge pieces e_2, e_3 on the same face of the cube as e_1. We claim that for any two edge pieces x, y in S, $(e_1xy) \in \beta(N)$. We first assume that one of these, say x, is distinct from e_2 and e_3. We can then find some $h \in H$ such that $\beta(h) = (e_1e_2e_3)$, some $n_1 \in N$ such that $\beta(n_1) = (e_2e_3x)$, and some $n_2 \in N$ such that $\beta(n_2) = (e_2e_3)(xy)$. We then have that $n = [h, n_1]n_2 \in N$ and that $\beta(n) = (e_1xy)$.

We next assume that $x = e_2$ and $y = e_3$. Since S contains at least three edges, we can find some edge $z \neq x, y$ in S. We can then find $h \in H$ and $n_1 \in N$ such that $\beta(h) = (e_1xy)$ and $\beta(n_1) = (xzy)$. As before, we have that $n = n_1[h, n_1] \in N$ and that $\beta(n) \in (e_1xy)$. We then have that $\beta(N)$ contains all 3-cycles on its prior letters and e_1, so that we can expand S to include e_1 as claimed.

Starting from $S = \{a, b, f\}$, we use this procedure to add edge pieces to S in this order: $c, d, g, h, e, i, j, k, l$. One can readily check that each edge piece shares a face with two edge pieces in S when it is added. Since S contains all of the edge pieces once we finish, $\beta(N)$ contains the alternating group on all edges of the cube, so that $\beta(N) \supseteq A_{12}$ as claimed. □

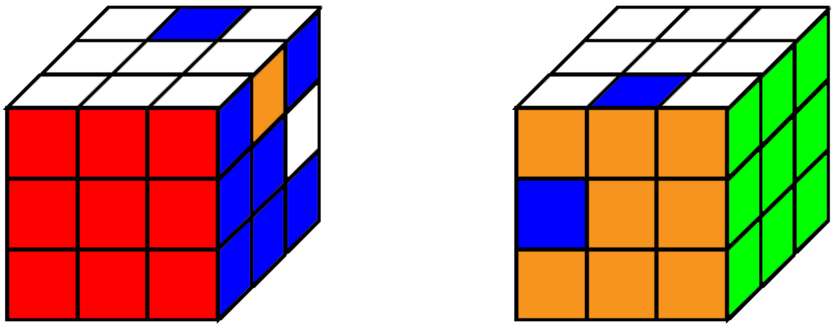

Figure 10: The move $[h_1, h_2]$ described in Proposition 3.2, shown from the front and back of the 3×3 cube.

We will now show that $\beta(n) \in A_{12}$ for any $n \in N$, so that $\beta(N)$ is in fact equal to A_{12}. For future convenience, we will split the proof into two short propositions.

Proposition 3.3. *For any $g \in G_3$, the factors of $\alpha(g) \in S_{12} \times S_8$ have the same sign.*

Proof. We can see from the formula (3) that for any generator a of G_3, the factors of $\alpha(g)$ have odd sign in S_8 and S_{12}. Since the signs of these factors agree for all generators a, they will agree for all elements $g \in G_3$ as well. □

Corollary 3.4. $\beta(N) = A_{12}$.

Proof. Since each $n \in N$ acts trivially on the corners of the cube, $\alpha(n)$ will have even sign in S_8 and so in S_{12}. Since $\beta(n)$ is simply the restriction of $\alpha(n)$ to the S_{12} factor, it will have even sign for all n, so that $\beta(N) \subseteq A_{12}$. Combining this with Proposition 3.2 proves the claim. □

We are left to determine the kernel M of β, or the group of rotations on the edges. In the next two propositions, we will show that M is isomorphic to $\mathbb{Z}_{2,0}^{12}$, using the same techniques as in Propositions 2.3 and 2.5. We will also provide a formula for the conjugation action of G_3 on M analogous to the formula (2). With this information, we can describe N through the short exact sequence $\mathbb{Z}_{2,0}^{12} \longrightarrow N \longrightarrow A_{12}$ induced by β. As a first step, we extend the definition of local orientation to edge positions and provide an illustration in Figure 11.

Definition 3.5. Following Definition 2.2, the local orientation of an edge position x is an element $t_x \in \mathbb{Z}_2$ defined by an orthonormal basis $\hat{v}_x$, $\hat{w}_x$ attached to the cubelet at that position. t_x is the number of rotations by π taking the current basis at i to the basis in the solved state.

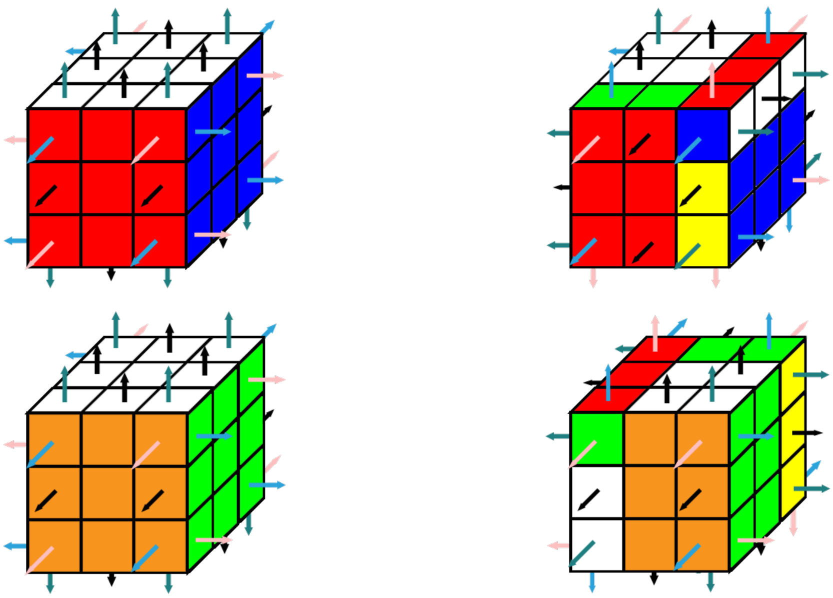

Figure 11: At left, bases for local orientations at each position. The edge orientation vectors v_x are denoted by black arrows, and the vectors w_x are left implicit. At right, the change in local orientation under rf. The changes in local orientations are $(a, b, c, d) = ([0], [1], [1], [0])$, $(e, f, g, h) = ([0], [0], [0], [1])$, $(i, j, k, l) = ([0], [0], [1], [0])$; $(1, 2, 3, 4) = ([1], [2], [0], [1])$, $(5, 6, 7, 8) = ([2], [2], [0], [1])$.

We can then extend Proposition 2.3 to produce an algebraic invariant on the edge pieces.

Proposition 3.6. *Let $t \in \mathbb{Z}_2$ denote the sum of local orientations $t_a, \ldots, t_j$ of the edges. t is preserved by G_3.*

Proof. As in Proposition 2.3, it is enough to show this for an arbitrary generator g of G_3, say u. We follow the same procedure as before, rotating the edges of the top face in place, rotating the entire top face, and then reversing the first rotation in place. This permutation of the stickers preserves t and is equivalent to u, so that u preserves t as well. Since u is arbitrary, any generator of G_3 will preserve t. Then any element of G_3 will preserve t, completing the proof. □

We then have that $M \subseteq \mathbb{Z}_{2,0}^{12}$, as each of the twelve edges has two orientations and the sum of these orientations is fixed. We can then write elements $m \in M$ in the form $(m_a, \ldots, m_l)$, where $m_x \in \mathbb{Z}_2$ denotes the rotation on edge x. As with the corners of the 2×2 cube, M is a normal subgroup of G_3, and the conjugation action of G_3 on M is given by the formula

$$gmg^{-1} = (m_{\sigma^{-1}(a)}, m_{\sigma^{-1}(b)}, \ldots, m_{\sigma^{-1}(l)}) \tag{4}$$

where g is an arbitrary element of G_3 and $\sigma = \beta(g)$. We will use this action to determine M completely.

Proposition 3.7. *M is maximal under the invariant t, so that $M \simeq \mathbb{Z}_{2,0}^{12}$.*

Proof. Returning to Proposition 3.1, let $h_1 = u^2 r^{-1} \ldots r$, so that $\beta(h_1) = (abc)$. Let $h_2 = f^2 u^{-1} \ldots u$ be the corresponding permutation on the front face, so that $\beta(h_2) = (cgh)$. Let $h_3 = r^2 f^{-1} \ldots f$ be the corresponding permutation on the right face, so that $\beta(h_3) = (bjg)$. Finally, let $m = [{h_3}^{-1}, h_1][h_2, {h_1}^{-1}]$.

We then assert that $m \in M$. In the component form given above, $m_c = m_g = [1]$, with all other components zero. As in Proposition 2.5, we leave the computation to the reader, although we illustrate m in Figure 12.

Using the component notation, we can consider M as a vector subspace of $\mathbb{Z}_2^{12}$. Since $M \subseteq \mathbb{Z}_{2,0}^{12}$, we know that M is at most 11-dimensional. We can choose a candidate basis for M of elements q_x $(x \neq a)$, where $(q_x)_a = (q_x)_x = [1]$ and all other components are zero; for example, $q_b = ([1], [1], [0], [0], \ldots, [0])$. We then recall the element $m \in M$ given above. Since the nonzero components of m in this space are sparse, we can take them to the nonzero components of any q_x through a permutation in A_{12} by transposing trivial components if necessary. By Proposition 3.2, any such permutation corresponds to conjugation by some element of H, so that each q_x is in fact included in M by normality. Then since the q_x are a set of 11 $\mathbb{Z}_2$-linearly independent elements in M, M is exactly 11-dimensional, so that $M \simeq \mathbb{Z}_{2,0}^{12}$ as claimed. □

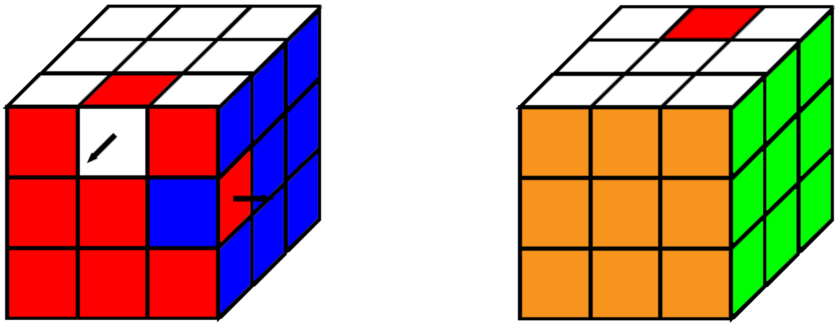

Figure 12: The move m described in Proposition 3.7, shown from the front and back of the 3×3 cube.

In sum, then, we have the following results: the kernel N of $\psi : G_3 \longrightarrow G_2$ fits into the short exact sequence $M \longrightarrow N \longrightarrow N/M$ induced by β, where M is the group of rotations on the edges and N/M records permutations of their positions. By Corollary 3.4 and Proposition 3.7, we can write this sequence as $\mathbb{Z}_{2,0}^{12} \longrightarrow N \longrightarrow A_{12}$, with the conjugation action given by the formula (4). By repeating the analysis in Section 2.4, one can show that this sequence is split, although we will omit the proof of this.

3.3 The Split Exact Sequence $N \longrightarrow G_3 \longrightarrow G_2$

In this section, we will analyze the short exact sequence $N \longrightarrow G_3 \longrightarrow G_2$ induced by ψ. We will show that this sequence is split, so that we can realize the 2×2 cube as a subgroup of the 3×3.

First, recall the map $\alpha : G_3 \longrightarrow S_{12} \times S_8$ given in Section 3.1. Let $P = \alpha(G_3)$, and let $J = \ker \alpha$. We can write $J = M \oplus L$, where L is the subgroup of G_3 rotating the corners in place and M is the subgroup rotating the edges as before. We begin by establishing some facts about L.

Proposition 3.8. *ψ induces an isomorphism $L \longrightarrow K$, where K is the group of corner rotations in G_2.*

Proof. It is clear that $\psi|_L$ is injective and that $\psi(L) \subseteq K$. To see that $\psi(L)$ is surjective, take any $k \in K$, and take some $g \in \psi^{-1}(k)$. Since g fixes the corners, we know by Proposition 3.3 that $\beta(g)$ has even sign, and we also know that g

respects the algebraic invariant t. However, we know from Section 3.2 that N is maximal under these constraints, so that $n^{-1}g$ acts trivially on the edges for some $n \in N$. We then have that $n^{-1}g \in L$ and that $\psi(n^{-1}g) = \psi(g) = k$, so that $\psi(L) = K$ as claimed. □

We then have that $J \simeq \mathbb{Z}_{2,0}^{12} \oplus \mathbb{Z}_{3,0}^{8}$, so that it is maximal under the invariants s and t. Figure 13 illustrates a typical element.

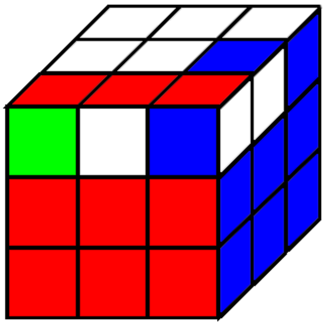

Figure 13: A typical element of J.

We will use the maximality of J to prove the next proposition.

Proposition 3.9. *The sequence* $J \longrightarrow G_3 \longrightarrow P$ *induced by* α *is split.*

Proof. Let $Q \subseteq G_3$ be the group of moves preserving local orientation on both the corners and edges. $\alpha|_Q$ is injective, since if $\alpha(q)$ is trivial, then q fixes both position and orientation. To see that $\alpha|_Q$ is surjective, take any $p \in P$ and $g \in \alpha^{-1}(p)$, and let $j \in \mathbb{Z}_{2,0}^{12} \oplus \mathbb{Z}_{3,0}^{8}$ denote the change in local orientation induced by g. Since j preserves the algebraic invariants s and t, and since J consists of all rotations of the cubelets preserving s and t, $j \in J$. Then $\alpha(j^{-1}g) = p$, and $j^{-1}g \in Q$, so that $\alpha^{-1}(p) \cap Q$ is nontrivial as claimed. This completes the proof. □

We can then write G_3 as $J \rtimes_{\tilde{\alpha}} P$, where the map $\tilde{\alpha} : P \longrightarrow \text{Aut}(J)$ is the conjugation action of P on J. This is given by the formulas (2) and (4), where the S_8 factor of P acts on L according to (2) and the S_{12} factor acts on M according to (4).

The next step is to choose a sign map $\sigma : S_8 \longrightarrow S_{12}$, where we embed $\mathbb{Z}_2$ in S_{12} as the group $\{(bc), 1\}$. We can then construct the subgroup S of $S_{12} \times S_8$ consisting of elements $(\sigma(p_2), p_2)$ for $p_2 \in S_8$.

Theorem 3.10. $L \rtimes_{\tilde{\alpha}} S$ *is contained in* G_3 *and is isomorphic to* G_2 *under* ψ.

Proof. In order to show that $S \subseteq P$, we note that the map $\phi \circ \psi : G_3 \longrightarrow S_8$ is surjective, so that G_3 contains all possible permutations of the corners. Given any $p_2 \in S_8$, we then know that P contains an element (p_1, p_2) for some $p_1 \in S_{12}$, and Proposition 3.3 requires that p_1 and p_2 have the same sign. It follows that $\sigma(p_2) = p_1 q$ for some $q \in A_{12}$, and since $A_{12} \times 1 \subseteq P$ by Corollary 3.4, $(\sigma(p_2), p_2) = (p_1, p_2)(q, 1) \in P$. We then have that $S \subseteq P$, and since $L \subseteq J$ is normal in G_3, $L \rtimes_{\tilde{\alpha}} S \subseteq J \rtimes_{\tilde{\alpha}} P = G_3$.

Using the standard component notation for elements of $L \rtimes_{\tilde{\alpha}} S$ and $G_2 = K \rtimes_{\tilde{\phi}} S_8$, we can then see that $\psi(l, (\sigma(p_2), p_2)) = (l, p_2)$. This map is clearly surjective in the second component, and we know by Proposition 3.8 that it is also surjective in the first. It is also clearly injective, so that the restriction of ψ to $L \rtimes_{\tilde{\alpha}} S$ is an isomorphism as claimed. □

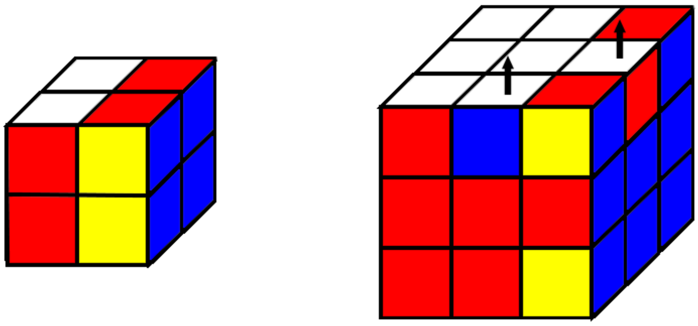

Figure 14: The element of $L \rtimes_{\tilde{\alpha}} S$ corresponding to $r_2 \in G_2$.

The subgroup $L \rtimes_{\tilde{\alpha}} S$ is easy to describe verbally. For any $g_2 \in G_2$, the corresponding element $h \in L \rtimes_{\tilde{\alpha}} S$ acts on the corners of the 3×3 cube as g_2 acts on the 2×2. If $\phi(g_2)$ is odd, h exchanges edges b and c while fixing local orientation. If not, h fixes the edges completely. Figure 14 illustrates a typical element.

In fact, this construction will also work with any $\sigma : S_8 \longrightarrow S_{12}$ that preserves sign, so that there are many different sections for ψ. However, since σ must preserve sign, there is no section that fixes the edges of the 3×3 cube.

This completes our analysis of the 3×3 cube group. Before moving on to minimal representations, we prove a final result concerning the sequence $J \longrightarrow G_3 \longrightarrow P$ that will be useful in Section 5.

Proposition 3.11. $P = (A_{12} \times A_8) \rtimes \mathbb{Z}_2$, *the subgroup of* $S_{12} \times S_8$ *consisting of elements whose factors have the same sign.*

Proof. By Proposition 3.3, we know that $P = \alpha(G_3) \subseteq (A_{12} \times A_8) \rtimes \mathbb{Z}_2$. To get the reverse inclusion, we can recycle the argument in the first paragraph of Theorem 3.10. Given any $p_1 \in S_{12}$ and $p_2 \in S_8$ with matching signs, we find some (p_1', p_2) in P, note that $(p_1(p_1')^{-1}, 1) \in A_{12} \times 1 \subseteq P$, and conclude that $(p_1, p_2) = (p_1(p_1')^{-1}, 1)(p_1', p_2) \in P$. We then have that $P \supseteq (A_{12} \times A_8) \rtimes \mathbb{Z}_2$, so that we get equality as claimed. □

4 Minimal Representations of Split Extensions by Abelian Groups

The analyses carried out in Sections 2 and Section 3 reveal the group structures of both the 2×2 and 3×3 Rubik's cubes to be surprisingly elementary. Upon choosing a different generating set, the puzzles become relatively straightforward. This is in large part due to the simplicity of the group structures. Both G_2 and G_3 are split extensions by large abelian groups of a complementary group which acts by permutations. G_2 is isomorphic to the split extension of S_8 by $\mathbb{Z}_{3,0}^8$ as seen in Proposition 2.6, whereas G_3 is the split extension by $\mathbb{Z}_{2,0}^{12} \oplus \mathbb{Z}_{3,0}^8$ of $(A_{12} \times A_8) \rtimes \mathbb{Z}_2$ by Propositions 3.9 and 3.11.

We wish to calculate the minimal dimension of faithful representations of G_2 and G_3. By Cayley's theorem, an upper bound on the dimension is given by the order of the group. However, this is generally far from minimal. Moreto sharpened this bound over the complex numbers to the square root of the order of the group [3]. The question of how to improve these bounds with additional assumptions on the group structure, such as if the group is nilpotent or solvable, seems to remain wide open. Matters grow even more complicated when the base field is not algebraically closed.

As we are interested in minimal representations over both the complexes and reals, we begin this section with a brief refresher on real, complex, and quaternionic representations. Motivated by our ambitions to better understand representations of the groups G_2 and G_3, we proceed to study minimal faithful representations of finite abelian groups. This analysis is carried out in the first part of this section where we compute the minimal dimension of both real and complex representations of finite abelian groups A. We express these bounds entirely in terms of the invariant factor decomposition of A.

We then determine some bounds on degrees of faithful real and complex representations of split extensions by abelian groups A of complementary groups H whose splitting homomorphism $H \longrightarrow \mathrm{Aut}(A)$ is injective. We relate these bounds to the so called *minimal permutation degree* as in [4]. Simply stated, the minimal permutation degree of a finite group G is the smallest number n for which G embeds inside an S_n. This number is denoted by $\mu(G)$. Similarly, we denote by $\mathrm{mdim}_F(G)$ the smallest dimension for which G admits a faithful representation over a field F. For our purposes, we will largely be concerned with $F = \mathbb{C}$ or $F = \mathbb{R}$.

Many of these calculations are afforded by Schur's lemma which states that the only G-equivariant endomorphisms between irreducible complex vector spaces are multiples of the scalar identity. As a consequence, the irreducible complex representations of abelian groups are all one-dimensional. We leverage this fact to prove that the minimal faithful complex dimension of a split product of an abelian group A and a complementary subgroup H that acts faithfully on A is at least the minimal permutation degree of H in Theorem 4.4.

In the context of real representations, the irreducible representations of abelian groups become slightly more complicated. Every irreducible representation of an abelian group is either 1 or 2 real dimensional. This follows from the fact that if one starts with a real representation of an abelian group A and extends scalars to $\mathbb{C}$ to obtain a complex representation whose degree is equal to the original real degree, this new complex representation either remains irreducible as a complex representation or splits into two conjugate irreducible factors. In the former case, the complex dimension must equal one, as complex irreducible representations are one-dimensional. In the latter case, the complex dimension must equal two, because the representation splits into conjugate irreducible pieces. Thus, the original representation is either 1 or 2 real dimensional. This additional dimension allows for a more rich automorphism group associated to the splitting homomorphism of the semi-direct product as seen in Theorem 4.6.

4.1 Real, Complex, and Quaternionic Representations

In this subsection we provide a brief review of real, complex, and quaternionic representations. Much of what follows is taken directly from Fulton and Harris's Section 3.5 on real, complex, and quaternionic representations [1].

A complex representation (ρ, V) is called *real* if there exists a representation (σ, W) over $\mathbb{R}$ for which $\mathbb{C} \otimes_{\mathbb{R}} W \simeq V$. That is to say, a complex representation is real if and only if it is obtainable by an extension of scalars from a representation over a real vector space. Equivalently, there exists a basis of V for which the corresponding matrix representation of ρ is real-valued. Note in this instance, the character of the representation in question is real-valued.

A real-valued character does not guarantee a real representation. One can construct complex representations whose character is real-valued but which cannot be obtained as an extension of scalars from any representation (σ, W) over $\mathbb{R}$. A straight-forward example can be constructed by considering the quaternion group Q_8 sitting inside $\mathrm{SU}(2, \mathbb{C})$ acting naturally on $\mathbb{C}^2$.

Representations such as the example provided above are called *quaternionic*. Specifically, a complex representation (ρ, V) is called quaternionic if it is the restriction of scalars of a representation (ϕ, U) over $\mathbb{H}$. Such representations are also commonly called *pseudo-real*.

Complex representations which are neither real nor quaternionic are simply called complex. For our purposes we will not need to differentiate between quaternionic and complex representations, so we call either type of such representation *non-real*.

Given a non-real representation (ρ, V), one may restrict scalars to obtain a real vector space W satisfying $\dim_{\mathbb{R}} W = 2 \dim_{\mathbb{C}} V$. If we extend scalars on W, we may construct an isomorphism between $\mathbb{C} \otimes_{\mathbb{R}} W$ and $V \oplus \overline{V}$. Define the map $\mathbb{C} \times W \longrightarrow V \oplus \overline{V}$ which takes (c, v) to $(cv, \overline{c}v)$. It is routine to check that this map induces an isomorphism between $\mathbb{C} \otimes_{\mathbb{R}} W$ and $V \oplus \overline{V}$ as complex vector spaces.

Consequently, if we take a non-real representation V and add its conjugate piece to obtain $V \oplus \overline{V}$, the resulting representation is real in the sense that it is an extension of scalars from a representation over a real vector space.

In our efforts to understand the relation between the minimal complex and minimal real dimensions of a finite group, it is natural to consider the following process. Take a minimal faithful complex representation (ρ, V) of a finite group G. V splits into the sum of irreducible pieces which we separate into real and non-real pieces, and denote by V_i and W_i respectively. One may obtain a faithful real representation of G by adding the conjugate pieces $\overline{W_i}$ to V. We informally refer to this process as *realification* or *real-ifying* a complex representation.

A natural question which we will investigate in Section 5.2 is whether or not every minimal faithful real representation is the realification of some minimal faithful complex representation.

With the preliminaries established, we may move on to investigating the minimal faithful real and complex representations of finite abelian groups.

4.2 Minimal Representations of Abelian Groups

We begin our analysis by calculating the minimal real and complex dimensions of finite abelian groups. To this end, we utilize the invariant factor decomposition of a finitely generated module M over a principal ideal domain R. Given such an M there exists a sequence of non-zero elements $d_i \in R$ satisfying $d_1|d_2|\ldots|d_s$, and a non-negative integer r for which

$$M \simeq R^r \oplus R/(d_1) \oplus R/(d_2) \oplus \ldots \oplus R/(d_s) \tag{5}$$

The above decomposition in (5) is an R-module isomorphism and is called the *invariant factor decomposition* of M. The integer r is called the *rank* of M, and the elements d_i are called the *elementary divisors.* The rank is uniquely determined whereas the elementary divisors are unique up to unit in R. A module M is called a *torsion module* if the rank $r = 0$. An advantage of this decomposition is that it is 'preserved' under homomorphism. This is made precise by the following statement which can be found in Jacobson [2].

Proposition 4.1. [Jacobson Exercise 3.9.4 (pg. 194)] *Let M be a torsion module over the principal ideal domain R with invariant factors $(d_1) \supset (d_2) \supset \ldots \supset (d_s)$. The homomorphic image $\overline{M}$ of M is also a torsion module whose invariant factor ideals $(\overline{d}_1) \supset (\overline{d}_2) \supset \ldots \supset (\overline{d}_t)$ satisfy $t \leq s$ and $\overline{d}_t|d_s, \overline{d}_{t-1}|d_{s-1}, \ldots,$and $\overline{d}_1|d_{s-t+1}$.*

As we are currently concerned with *finite abelian* groups, we are interested in the case where our principal ideal domain is $R = \mathbb{Z}$ and our module is torsion. In this setting, the conclusions of Proposition 4.1 still hold for *subgroups* of our finite abelian group A. This is a simple consequence of the fact that every subgroup $B \subset A$ of a finite abelian group A is isomorphic to a quotient of A.

With these preliminaries established, we are able to prove the first theorem of this section regarding $\mathrm{mdim}_{\mathbb{C}}(A)$ and $\mathrm{mdim}_{\mathbb{R}}(A)$ for finite abelian groups A.

Theorem 4.2. *Let A be a finite abelian group whose invariant factor decomposition is given by $A \simeq \mathbb{Z}_2^a \oplus \mathbb{Z}_{d_{a+1}} \oplus \ldots \oplus \mathbb{Z}_{d_{a+b}}$ where $d_1 = \ldots = d_a = 2$ and $d_{a+1} > 2$. Then $\mathrm{mdim}_{\mathbb{C}}(A)$ is equal to the number of invariant factors in the invariant factor decomposition of A, $a + b$. The real dimension, $\mathrm{mdim}_{\mathbb{R}}(A)$, is equal to $a + 2b$.*

Proof. We first begin with the complex case. It is straightforward to construct a faithful $(a + b)$-complex dimensional representation of A. For each elementary divisor d_i, let ω_{d_i} be the d_i-th root of unity $\exp(2\pi i/d_i)$. Identify each $\mathbb{Z}_{d_i}$ with the cyclic subgroup of $\langle \omega_{d_i} \rangle \subset \mathbb{C}^\times$. Here $\mathbb{C}^\times$ acts on $\mathbb{C}$ by multiplication. Form the direct sum of these representations to obtain a faithful representation of A on $\mathbb{C}^{a+b}$, so $\mathrm{mdim}_{\mathbb{C}}(A) \leq a + b$.

We claim this construction is minimal in dimension. Let (ρ, V) be a faithful complex representation of A. Because A is abelian, V decomposes as the direct sum of one-dimensional irreducibles, say $V_1, V_2, \ldots, V_p$. For each $a \in A$ and each i where $1 \leq i \leq p$, $\rho(a)$ acts by scalar multiplication on V_i. Because a has order at most d_{a+b}, as the order of a divides d_{a+b}, this means $\rho(a)$ acts on each V_i by some d_{a+b}-th root of unity. Hence we may identify each $\rho(a)$ with p elements in $\mathbb{Z}_{d_{a+b}}$. These elements correspond to how $\rho(a)$ acts on each V_i. By faithfulness of the representation, A itself may be identified with a subgroup of $\mathbb{Z}_{d_{a+b}}^p$. By the remarks following Proposition 4.1, we have that there are at least as many invariant factors in $\mathbb{Z}_{d_{a+b}}^p$ as there are in A. Thus $p \geq a+b$ as claimed, hence $\operatorname{mdim}_{\mathbb{C}}(A) = a+b$.

To prove the analogous statement in the real case, we begin by first noting that one can obtain a faithful $(a+2b)$-dimensional real representation by simply restricting scalars of the faithful complex representation constructed in the beginning of this proof, so $\operatorname{mdim}_{\mathbb{R}}(A) \leq a+2b$. We claim this construction is minimal in dimension.

Let (ρ, V) be a faithful real representation of A. Begin by decomposing V into its 1 and 2 real-dimensional irreducible factors. Say $V_1, \ldots, V_p$ are the 1-dimensional pieces and $W_1, \ldots, W_q$ are the 2-dimensional pieces. On each V_i, $\rho(a)$ acts by multiplication by ± 1. On each W_i, $\rho(a)$ acts by rotation. The angle of rotation is some integer multiple of $2\pi/d_{a+b}$ as $\rho(a)$ has order at most d_{a+b}. We may very well consider this angle of rotation as an integer multiple of $2\pi/(2d_{a+b})$ instead.

Via our representation ρ, we may assign to each $a \in A$ a p-tuple of elements in $\mathbb{Z}_2$ and a q-tuple of elements in $\mathbb{Z}_{2d_{a+b}}$. As ρ is faithful, this map defines an embedding of A into $\mathbb{Z}_2^p \oplus \mathbb{Z}_{2d_{a+b}}^q$. The first p factors are determined by whether $\rho(a)$ acts by ± 1 on each V_i, whereas the latter q factors are determined by how many integer multiple rotations by $2\pi/(2d_{a+b})$ were applied to W_i by $\rho(a)$.

Because A has $(a+b)$ invariant factors in its invariant decomposition and is a subgroup of $\mathbb{Z}_2^p \oplus \mathbb{Z}_{2d_{a+b}}^q$, by the remarks following Proposition 4.1, $p+q \geq a+b$. Similarly, because for each $a < i \leq a+b$, we have $d_i > 2$, there must be at least b copies of $\mathbb{Z}_{2d_{a+b}}$ in the invariant factor decomposition $\mathbb{Z}_2^p \oplus \mathbb{Z}_{2d_{a+b}}^q$. Thus, $q \geq b$, so $\dim_{\mathbb{R}} V = p + 2q \geq a+2b$. Hence the dimension of any faithful real representation of A is at least $a+2b$, thus $\operatorname{mdim}_{\mathbb{R}}(A) = a+2b$. □

4.3 A Lower Bound on Minimal Faithful Representations of Split Extensions by Abelian Groups

Having established the precise minimal faithful dimensions of real and complex finite abelian groups, we move on to calculating a lower bound of faithful representations of groups G which are split extensions by finite abelian groups A of some complementary group H where the splitting homomorphism $H \longrightarrow \mathrm{Aut}(A)$ is faithful. This class of groups contains our Rubik's cube groups G_2 and G_3.

The crux of these arguments comes from utilizing the faithfulness hypotheses of both the representation itself and the splitting homomorphism $H \longrightarrow \mathrm{Aut}(A)$. Given a representation G, we may restrict it to A to obtain several A-invariant pieces. These pieces are permuted by the complement H, and thus we obtain a permutation representation of H into some symmetric group. If we are better able to understand the kernel of this homomorphism, if any, we then obtain bounds on the number of irreducible pieces, and thus, the dimension of the representation in question. To this end, we begin first with a proposition that formalizes some of the discussion in this paragraph.

Proposition 4.3. *If (π, V) is a representation of some group G over a field F and N is a normal subgroup of G, then elements of G will permute the N-irreducible subrepresentations V_i of V.*

Proof. Let W be some N-irreducible subrepresentation of V, and consider the space $\pi(g)W = \{\pi(g)w : w \in W\}$ for some $g \in G$. We can first show that this space is N-invariant. For any $n \in N$ and any $w \in W$, $\pi(n)\pi(g)w = \pi(g)\pi(g^{-1}ng)w = \pi(g)w'$ for some $w' \in W$, since N is normal and W is N-invariant. We can then show that $\pi(g)W$ is N-irreducible. If it decomposes into N-invariant subspaces U_1 and U_2, then $\pi(g^{-1})U_1$ and $\pi(g^{-1})U_2$ are N-invariant subspaces of W, which contradicts the assumption that W is N-irreducible. Then $\pi(g)$ induces a set map on the N-irreducible subrepresentations V_i, and since $\pi(g)$ is invertible, this map must be a permutation. This completes the proof. □

For our purposes, we are largely interested in the semi-direct product G of a finite abelian group A with a complementary group H that acts faithfully on A. While the permutation map is defined on all of G, it is trivial on its normal subgroup A, so we instead restrict the permutation map to the complementary subgroup H. With this lemma and the preliminaries we established, we are ready to prove our first theorem.

Theorem 4.4. *Let G be the semi-direct product of a finite abelian group A and a finite group H, and assume the splitting map $H \longrightarrow \mathrm{Aut}(A)$ is injective. Let (π, V)*

be a faithful representation of G over $\mathbb{C}$. Let $V_1, V_2, \ldots, V_p$ denote the irreducible subrepresentations of V restricted to A. The group H permutes these pieces and induces a faithful group homomorphism $H \longrightarrow S_p$, so $mdim_{\mathbb{C}}(G) \geq \mu(H)$.

Proof. The irreducible subrepresentations V_i of A are all one-dimensional by Schur's lemma. Since $V = \bigoplus_{i=1}^{p} V_i$, we have that $p = \dim_{\mathbb{C}} V$. By Proposition 4.3, H permutes these subrepresentations, and thus induces a homomorphism $H \longrightarrow S_p$. We claim this homomorphism is injective.

If $h \in H$ and stabilizes each V_i, we guarantee the existence of $c_1, \ldots, c_p \in \mathbb{C}^{\times}$ so that $\pi_h(v_i) = c_i v_i$ for any $v_i \in V_i$. We claim that h must be equal to the identity. As every $a \in A$ also stabilizes each V_i, we have that $\pi(hah^{-1}) = \pi(a)$ as both a and h act by scalar multiplication on each V_i. By faithfulness of the representation this means that $hah^{-1} = a$ for all $a \in A$. By the hypothesis that the splitting map is injective, we have that $h = 1$. Because the homomorphism $H \longrightarrow S_p$ is injective, we have that $p \geq \mu(H)$. Hence the degree of a minimal faithful representation of G over $\mathbb{C}$ is at least $\mu(H)$. $\square$

From this theorem, we immediately obtain the following corollary.

Corollary 4.5. *Let A be an abelian group and let G be a semi-direct product of A with S_m where S_m acts faithfully on A by automorphisms. Let (π, V) be a faithful representation of G over $\mathbb{C}$. The map $S_m \longrightarrow S_p$ induced by permutations of the irreducible pieces as in Theorem 4.4 is injective, thus $mdim_{\mathbb{C}}(G) \geq m$.*

Proof. This is a simple consequence of the fact that if S_m embeds in S_p, then $p \geq m$. $\square$

We now prove an analogous result to that of Theorem 4.4 in the real setting. In this case, the irreducible representations of the normal abelian subgroup are either one or two dimensional. The two dimensional case affords instances where the complementary subgroup could stabilize irreducible pieces, but reverse orientations.

Theorem 4.6. *Let G be the semi-direct product of a finite abelian group A and a finite group H, and assume the splitting map $H \longrightarrow \mathrm{Aut}(A)$ is injective. Let (π, V) be a faithful representation of G over $\mathbb{R}$. Let $V_1, V_2, \ldots, V_p$ denote the 1-dimensional irreducible subrepresentations of V restricted to A and $W_1, W_2, \ldots, W_q$ denote the 2-dimensional irreducible subrepresentations of V restricted to A. The group H permutes these pieces and induces a faithful group homomorphism $H \longrightarrow \mathbb{Z}_2^q \rtimes (S_p \times S_q)$ where the S_p factor of $S_p \times S_q$ acts trivially on $\mathbb{Z}_2^q$ and S_q by permutations on $\mathbb{Z}_2^q$.*

Proof. We begin by noting that H must permute the 1 and 2 dimensional A-irreducible subrepresentations V_i and W_i of V amongst themselves by Proposition 4.3. This defines a homomorphism $H \longrightarrow S_p \times S_q$. We construct another homomorphism $H \longrightarrow \mathbb{Z}_2^q \rtimes (S_p \times S_q)$ that will factor through the permutation map $H \longrightarrow S_p \times S_q$.

To begin, fix an orientation o_i on each 2-dimensional A-irreducible subrepresentation W_i for $i = 1, \ldots, q$. Such a choice is arbitrary, but once chosen, we fix it once and for all. For each $h \in H$ and $i = 1, \ldots, q$, we know that $\pi(W_i)$ is equal to another one of the irreducible factors $\pi(W_i) = W_j$. For each i, assign an element of $\mathbb{Z}_2$ depending on whether $\pi(h)|_{W_i} : (W_i, o_i) \longrightarrow (W_j, o_j)$ preserves, or reverses, the orientation. Thus, to each element $h \in H$, we assign a q-tuple of elements of $\mathbb{Z}_2$ which keep record of whether $\pi(h)|_{W_i}$ is orientation preserving or reversing from W_i to $\pi(h)(W_i)$. Augmenting this with the permutation map, we have a well defined homomorphism from H to $\mathbb{Z}_2^q \rtimes (S_p \times S_q)$ where S_q acts on $\mathbb{Z}_2^q$ by permutations and S_p acts trivially on $\mathbb{Z}_2^q$. In effect, one may think of this as a 'decorated' permutation map of the representation which records where the irreducible pieces are sent via a permutation, and, whether or not orientations are preserved or reversed on each W_i via a q-tuple of elements of $\mathbb{Z}_2$.

We claim this homomorphism is injective. Fix an $h \in H$ in the kernel of the homomorphism. Such an h defines an automorphism of A for which $\pi(h)(V_i) = V_i$ and $\pi(h)(W_i) = W_i$ for all irreducible pieces. Because on each V_i, $\pi|_{V_i}$ acts by scalar multiplication by ± 1, we have that $\pi(hah^{-1})|_{V_i} = \pi(a)|_{V_i}$. On each 2-dimensional W_i, $\pi(h)|_{W_i}$ acts by a composition of a rotation and a (possibly trivial) reflection. By hypothesis, $\pi(h)$ preserves orientation, thus on each irreducible W_i, h is acting by rotation alone. However, this means that $\pi(hah^{-1})|_{W_i} = \pi(a)|_{W_i}$ as $\pi(a)$ acts by rotation as well, and therefore commutes with $\pi(h)|_{W_i}$. Hence $\pi(hah^{-1}) = \pi(a)$ for all $a \in A$, and by the faithfulness hypotheses, $h = 1$. Thus the decorated permutation homomorphism $H \longrightarrow \mathbb{Z}_2^q \rtimes (S_p \times S_q)$ is injective. □

We conclude this section with an analogue of Corollary 4.5 adapted to real representations.

Corollary 4.7. *Let A be an abelian group and let G be a semi-direct product of A with S_m where S_m acts faithfully on A by automorphisms for some m not equal to 2 or 4. Compose the decorated permutation map $S_m \longrightarrow \mathbb{Z}_2^q \rtimes (S_p \times S_q)$ as constructed in Theorem 4.6 with the quotient $\mathbb{Z}_2^q \rtimes (S_p \times S_q) \longrightarrow S_p \times S_q$. This map $S_m \longrightarrow S_p \times S_q$ is injective onto at least one of the factors S_p or S_q. In particular, this means that either p or q is at least as large as m so $mdim_{\mathbb{R}}(G) \geq m$.*

Proof. As S_m acts faithfully on A, by Theorem 4.6, we have the map $S_m \longrightarrow \mathbb{Z}_2^q \rtimes$

$(S_p \times S_q)$ is injective. Assume there is some non-trivial kernel to $S_m \longrightarrow S_p \times S_q$. By injectivity of the map $S_m \longrightarrow \mathbb{Z}_2^q \rtimes (S_p \times S_q)$, this means that the image of S_m non-trivially intersects $\mathbb{Z}_2^q$, and thus $q \neq 0$. By injectivity, this means some subgroup of $\mathbb{Z}_2^q$ sits normally inside S_m. The only values of m for which there is a non-trivial normal 2-group in S_m are $m = 2$ and $m = 4$.

For $m = 2$, S_2 is itself a 2-group and hence the kernel is all of S_2, therefore the induced permutation map $S_2 \longrightarrow S_p \times S_q$ is trivial. For $m = 4$, S_4 has a normal Klein 4-group.

Thus, if $m \neq 2, 4$, S_m does not possess a non-trivial normal 2-group, and therefore the map $S_m \longrightarrow S_p \times S_q$ is injective. In this instance we claim that at least one factor composition, $S_m \longrightarrow S_p$ or $S_m \longrightarrow S_q$, is injective.

If both $S_m \longrightarrow S_p$ and $S_m \longrightarrow S_q$ have non-trivial kernels, we may denote them by K_1 and K_2 separately. Note that $K_1 \cap K_2$ must be trivial, because if $k \in K_1 \cap K_2$, then the image of k under $S_m \longrightarrow S_p \times S_q$ is trivial. By injectivity of this map, $k = 1$. Thus S_m possesses two distinct disjoint non-trivial normal subgroups. In particular, neither K_1 nor K_2 may equal S_m, thus these kernels are proper.

The only instance in which S_m has two non-trivial proper normal subgroups is when $m = 4$, but in this case, those normal subgroups are nested as the Klein 4-group sitting inside A_4. However, by hypothesis $m \neq 4$. Therefore, either $S_m \longrightarrow S_p$ or $S_m \longrightarrow S_q$ must have a trivial kernel, thus S_m embeds in either S_p or S_q. □

5 Examples and Applications

5.1 Minimal Faithful Representations of G_2 and G_3

In this section, we employ Theorems 3.10, 4.2, 4.4, and 4.6 to calculate the minimal real and complex dimensions of faithful representations of G_2 and G_3. We use Propositions 2.6, 3.9, and 3.11 which express G_2 and G_3 as split extensions by abelian groups of complementary permutation groups to our advantage in these calculations. More specifically, given a faithful representation of either G_2 or G_3, we are able to say that the permutation maps on irreducible pieces of their normal abelian subgroups as defined in Theorems 4.4 and 4.6 are typically injective. The fact that the complementary subgroups are permutation groups with few to no non-trivial normal subgroups allows for tractable case-by-case analysis. These statements are made precise in Theorem 5.1 and Theorem 5.2.

We close this section with a peculiar example that illustrates the subtlety between the relations amongst minimal real and complex dimensions of finite groups. One may wager that because the complex representations of a finite group determine the real ones, a minimal faithful real representation is determined by a minimal faithful complex one. In Theorem 5.3, we construct an example where a minimal faithful real representation is not induced by a minimal faithful complex one.

To begin, we focus our attention on the less complicated Rubik's cube group G_2 and determine its minimal faithful dimensions.

Theorem 5.1. *We have that* $mdim_{\mathbb{C}}(G_2) = 8$ *and* $mdim_{\mathbb{R}}(G_2) = 16$ *respectively. A minimal faithful real representation of* G_2 *may be obtained by a restriction of scalars of a complex faithful representation.*

Proof. We begin by constructing a faithful 8-dimensional complex representation of G_2. Let us first introduce some notation. Let $\omega_3 := \exp(2\pi i/3)$. We obtain a faithful representation of $\mathbb{Z}_{3,0}^8$ by taking the 8-tuple $([m_1], [m_2], \dots, [m_8]) \in \mathbb{Z}_{3,0}^8$ to the diagonal 8×8-matrix whose entries are $(\omega_3^{m_1}, \omega_3^{m_2} \dots, \omega_3^{m_8})$. Denote this representation of $\mathbb{Z}_{3,0}^8$ by $(\pi, \mathbb{C}^8)$. Because the standard permutation representation $(\rho, \mathbb{C}^8)$ of S_8 on $\mathbb{C}^8$ normalizes the representation of $\mathbb{Z}_{3,0}^8$, ρ is compatible with π in the following sense. For each $h \in S_8$ and any $a \in \mathbb{Z}_{3,0}^8$, we have that $\pi\left(hah^{-1}\right) = \rho(h)\pi(a)\rho\left(h^{-1}\right)$. Thus we may extend π to a representation of G_2 by declaring $\pi(ah) := \pi(a)\rho(h)$. This is a well-defined function by the fact that $S_8 \cap \mathbb{Z}_{3,0}^8 = 1$ in G_2. It is a homomorphism due to the compatibility relation as seen in the following equations.

$$\begin{aligned}\pi\left((ah)(a'h')\right) &:= \pi\left(a(ha'h^{-1})\right)\rho(hh') = \pi(a)\pi(ha'h^{-1})\rho(hh') \\ &= \pi(a)\left[\rho(h)\pi(a')\rho(h^{-1})\right]\rho(hh') = \pi(a)\rho(h)\pi(a')\rho(h') = \pi(ah)\pi(a'h')\end{aligned}$$

That this representation is faithful is readily seen by the faithfulness of both π and ρ which take images in complementary subgroups of the general linear group.

By Theorem 4.4, the degree of any faithful representation of G_2 is at least 8 as S_8 acts on $\mathbb{Z}_{3,0}^8$ faithfully. Thus $\mathrm{mdim}_{\mathbb{C}}(G_2) = 8$.

We now prove $\mathrm{mdim}_{\mathbb{R}}(G_2) = 16$. One can readily obtain a faithful real 16-dimensional representation of G_2 by restricting scalars of the complex representation constructed above.

We prove this dimension is minimal. G_2 is isomorphic to $\mathbb{Z}_{3,0}^8 \rtimes S_8$ where S_8 acts on $\mathbb{Z}_{3,0}^8$ by permutations. Let (π, V) be a real faithful representation of G_2.

Let $S_8 \longrightarrow S_p$ and $S_8 \longrightarrow S_q$ be the permutation representations as constructed in Theorem 4.6. By Corollary 4.7, we know that S_8 injects into either S_p or S_q.

If S_8 injects into S_q, then there are at least eight 2-dimensional $\mathbb{Z}_{3,0}^8$-irreducible pieces, thus $\dim_{\mathbb{R}} V \geq 16$. If S_8 fails to inject into S_q, then S_8 injects into S_p, and thus we have at least eight 1-dimensional $\mathbb{Z}_{3,0}^8$-irreducible pieces. Additionally, by Theorem 4.2, we have at least seven 2-dimensional $\mathbb{Z}_{3,0}^8$-irreducible pieces due to the invariant factor decomposition of $\mathbb{Z}_{3,0}^8 \simeq \mathbb{Z}_3^7$. Thus $p \geq 8$ and $q \geq 7$, so our representation is at least $8 + 2 \cdot 7 = 22$-dimensional.

The lower of these two bounds is 16. Therefore, any faithful real representation of G_2 is at least 16-dimensional. □

We now state and prove an analogous theorem concerning the 3×3 cube group G_3.

Theorem 5.2. *We have that $mdim_{\mathbb{C}}(G_3) = 20$ and $mdim_{\mathbb{R}}(G_3) = 28$. A minimal faithful real representation may be obtained from a minimal complex one by adding an appropriate conjugate piece.*

Proof. We may construct a faithful complex representation of G_3 in a manner similar to that in the proof of Theorem 5.1. We begin by constructing a faithful representation of $\mathbb{Z}_{2,0}^{12}$ by taking the 12-tuple $([m_1], [m_2], \ldots, [m_{12}]) \in \mathbb{Z}_{2,0}^{12}$ to the diagonal 12×12 matrix whose entries are $((-1)^{m_1}, (-1)^{m_2} \ldots, (-1)^{m_{12}})$. Similarly, we take an element $([m_1], [m_2], \ldots, [m_8]) \in \mathbb{Z}_{3,0}^8$ to the diagonal 8×8 matrix whose entries are $(\omega_3^{m_1}, \omega_3^{m_2} \ldots, \omega_3^{m_8})$. We obtain a faithful representation of $\mathbb{Z}_{2,0}^{12} \oplus \mathbb{Z}_{3,0}^8$ by taking the sum of these two representations. Denote this representation by $(\pi, \mathbb{C}^{20})$.

Consider the standard permutation representation $(\rho, \mathbb{C}^{20})$ of $S_{12} \times S_8$ and restrict it to the subgroup P of all pairs of permutations $(\sigma, \rho) \in S_{12} \times S_8$ with the same sign, as defined in Section 3.3. In a similar fashion to Theorem 5.1, we note that because this permutation representation normalizes the representation π of $\mathbb{Z}_{2,0}^{12} \oplus \mathbb{Z}_{3,0}^8$, we may extend π to a representation of G_3 by declaring $\pi(ah) := \pi(a)\rho(h)$ to obtain a faithful 20-dimensional complex representation of G_3.

By Theorem 4.4, the degree of any faithful representation of G_3 is at least the minimal permutation degree of P. As P contains a copy of $A_8 \times A_{12}$, this permutation degree is at least as large as $\mu(A_8 \times A_{12}) = \mu(A_8) + \mu(A_{12}) = 20$, as the minimal permutation degree of the product is additive for the alternating groups [4]. Thus, we have the minimal dimension of a faithful complex representation of

G_3 is 20.

We now prove the minimal faithful real dimension of G_3 is 28. One can obtain such a real 28-dimensional representation from a minimal complex one. Let $(\pi, \mathbb{C}^{20})$ denote the complex minimal faithful representation constructed above. Our representation splits into two pieces $\mathbb{C}^{20} = \mathbb{C}^{12} \oplus \mathbb{C}^8$, the first piece corresponding to the edges and the second piece corresponding to the corners. The first piece $\mathbb{C}^{12}$ is a reducible real representation whereas the second piece $\mathbb{C}^8$ is irreducible and non-real. The first piece splits into pieces of dimension 1 and 11, the former of which is generated by all vectors in $\mathbb{C}^{20}$ whose first 12 entries are equal and the remaining are zero. This reducibility comes from the center of the group G_3 which is isomorphic to $\mathbb{Z}_2$.

We may obtain a real representation from $\mathbb{C}^{20}$ by real-ifying and adding a complex conjugate piece to yield $\mathbb{C}^{12} \oplus (\mathbb{C}^8 \oplus \overline{\mathbb{C}^8})$. That this representation is faithful follows immediately from faithfulness of the original complex representation. Hence $\mathrm{mdim}_{\mathbb{R}}(G_3) \leq 12 + 2 \cdot 8 = 28$.

We now prove this dimension is minimal. G_3 is isomorphic to $(\mathbb{Z}_{2,0}^{12} \oplus \mathbb{Z}_{3,0}^{8}) \rtimes P$ where P acts on $(\mathbb{Z}_{2,0}^{12} \oplus \mathbb{Z}_{3,0}^{8})$ by permutations. Let (π, V) be a real faithful representation of G_3. Let $P \longrightarrow \mathbb{Z}_2^q \rtimes (S_p \times S_q)$ be the decorated permutation homomorphism afforded by Theorem 4.6. As P contains no non-trivial normal 2-groups, the decorated permutation homomorphism is injective. We claim the homomorphism onto the quotient $S_p \times S_q$ is injective. As P is isomorphic to $(A_{12} \times A_8) \rtimes \mathbb{Z}_2$ where $\mathbb{Z}_2$ acts on each factor by conjugation by an odd permutation, the only normal subgroups P contains are isomorphic to $1, A_8, A_{12}, A_8 \times A_{12}$, and P itself. None of these groups are isomorphic to a subgroup of $\mathbb{Z}_2^q$ for $q \neq 0$. Hence the kernel of the map $P \longrightarrow S_p \times S_q$ must be trivial.

We now consider the projection maps $P \longrightarrow S_p$ and $P \longrightarrow S_q$. By injectivity of $P \longrightarrow S_p \times S_q$, the kernels of these maps must intersect trivially. Denote these kernels by $K_p, K_q \subset P$ respectively. In the table in (6) below we list the possible pairs of kernels (K_p, K_q), the corresponding minimum number of 1 and 2 dimensional irreducible pieces afforded by $P/K_p \longrightarrow S_p$ and $P/K_q \longrightarrow S_q$, and a lower bound of the real dimension of V, $p + 2q$.

K_p	K_q	p	q	$p+2q$
1	1	20	20	60
$1 \times A_8$	$A_{12} \times 1$	12	8	28
$A_{12} \times 1$	$1 \times A_8$	8	12	32
1	$A_8 \times A_{12}$	20	2	24
$A_8 \times A_{12}$	1	2	20	42
1	P	20	0	20
P	1	0	20	40

(6)

For example, in line 3 of (6), the kernel of the map $P \longrightarrow S_p$ is $A_{12} \times 1$ and the kernel of the map $P \longrightarrow S_q$ is $1 \times A_8$. Thus, we have $P/(A_{12} \times 1) \simeq S_8$ embedding in S_p and $P/(1 \times A_8) \simeq S_{12}$ embedding in S_q. Therefore $p \geq 8$ and $q \geq 12$, so $\dim_{\mathbb{R}} V \geq 8 + 2 \cdot 12 = 32$.

The only cases we must further inspect are when $K_p = 1$ and $K_q = A_8 \times A_{12}$ and when $K_p = 1$ and $K_q = P$. In these cases, we would have at least twenty 1-dimensional $\mathbb{Z}_{2,0}^{12} \oplus \mathbb{Z}_{3,0}^{8}$-irreducible pieces. The invariant factor decomposition of this group is $\mathbb{Z}_2^4 \otimes \mathbb{Z}_6^7$. By Theorem 4.2, this means there are at least seven 2-dimensional $\mathbb{Z}_{2,0}^{12} \oplus \mathbb{Z}_{3,0}^{8}$-irreducible pieces. Thus in these cases, $p + 2q$ is at least $20 + 2 \cdot 14 = 48$. Consequently, every faithful representation of G_3 is at least 28 real-dimensional, so $\mathrm{mdim}_{\mathbb{R}}(G_3) = 28$. □

5.2 An Exceptional Example

The relationship between the minimal real and complex dimensions of a finite group remains unclear. Certainly for any finite group G, $\mathrm{mdim}_{\mathbb{C}}(G) \leq \mathrm{mdim}_{\mathbb{R}}(G)$. Furthermore, for any minimal faithful complex representation (ρ, V), we may take its realification as defined in Section 4.1 to obtain a faithful real representation. There are only finitely many minimal faithful complex representations up to isomorphism, so we take the one whose realification is of smallest real dimension to obtain a crude upper bound on $\mathrm{mdim}_{\mathbb{R}}(G)$. That is to say, $\mathrm{mdim}_{\mathbb{R}}(G)$ is bounded above by the smallest dimension over all realifications of minimal complex faithful representations of G.

Given the results of Theorem 5.1 and Theorem 5.2, one might conjecture that a minimal real representation is always obtainable from a minimal complex representation in the fashion described in the above paragraph and in the proof of Theorem 5.2. However, this does not always hold true as the following example illustrates. This example is motivated by Corollary 4.7 which states that so long as m is neither 2 or 4, then the decorated permutation homomorphism

$S_m \longrightarrow \mathbb{Z}_2^q \rtimes (S_p \times S_q)$ descends to an injective permutation homomorphism on the quotient $S_m \longrightarrow S_p \times S_q$.

Theorem 5.3. *Let $G = \mathbb{Z}_{3,0}^4 \rtimes S_4$ where S_4 acts by permutations on $\mathbb{Z}_{3,0}^4$. We have that $mdim_{\mathbb{C}}(G) = 4$ and $mdim_{\mathbb{R}}(G) = 6$. Every such minimal complex representation of G is irreducible and non-real, thus no minimal real representation is from a minimal complex representation.*

Proof. Let $G = \mathbb{Z}_{3,0}^4 \rtimes S_4$. That G possesses a 4-dimensional complex faithful representation is straightforward and follows from the arguments in the beginning of Theorem 5.1. If we choose any faithful 4-dimensional complex representation (π, V) of G, we claim it must be irreducible. Note that G has a minimal non-trivial normal subgroup $\mathbb{Z}_{3,0}^4$. That $\mathbb{Z}_{3,0}^4$ is minimal amongst non-trivial normal subgroups follows from an analysis similar to Propositions 2.7, 2.8 and 2.9 and is omitted for brevity.

If V is reducible and breaks into two smaller pieces, say $V = V_1 \oplus V_2$, then we have new representations (π, V_1) and (π, V_2) which by minimality of V, must both have non-trivial kernel. However, by minimality of $\mathbb{Z}_{3,0}^4$ amongst normal subgroups in G, this means $\mathbb{Z}_{3,0}^4$ is in both kernels, hence π was not faithful to begin with. Thus V is irreducible. That π is non-real may be readily seen by restricting our representation π to $\mathbb{Z}_{3,0}^4$ which has many characters with non-trivial imaginary parts.

We now prove there is a 6-dimensional real faithful representation of G. Because $\mathbb{Z}_{3,0}^4 \simeq \mathbb{Z}_3^3$, by Theorem 4.2, the real dimension of a minimal faithful representation $\mathbb{Z}_{3,0}^4$ is 6. Let us choose such a representation $(\pi, \mathbb{R}^6)$. $\mathbb{R}^6$ splits into three 2-dimensional irreducible pieces which we denote by V_1, V_2, and V_3. Without loss of generality, for each $1 \leq i \leq 3$, we may choose V_i to be generated by the $(2i-1)$-th and $2i$-th basis vectors, so for example V_2 is generated by e_3 and e_4. We also choose orientations o_i for each V_i by declaring $\{e_{2i-1}, e_{2i}\}$ to be an oriented basis.

We now construct a faithful 6-dimensional representation of S_4 using the exceptional isomorphism between S_4 and $\mathbb{Z}_{2,0}^3 \rtimes S_3$ where S_3 acts on $\mathbb{Z}_{2,0}^3$ by permutations. In what follows, we identify S_4 with $\mathbb{Z}_{2,0}^3 \rtimes S_3$.

For $1 \leq i \leq 3$, let us define orientation-reversing involutions ϵ_i of $\mathbb{R}^6$ that interchange the $(2i-1)$-th and $2i$-th basis vectors, and leave all other basis vectors fixed. For example $\epsilon_2(e_3) = e_4$, $\epsilon_2(e_4) = e_3$, and ϵ_2 fixes all vectors in V_1 and V_3. We obtain a faithful representation of $\mathbb{Z}_{2,0}^3$ into $\mathbb{R}^6$ by sending each element $([m_1, m_2, m_3]) \in \mathbb{Z}_{2,0}^3$ to $\epsilon_1^{m_1} \oplus \epsilon_2^{m_2} \oplus \epsilon_3^{m_3}$. This defines a faithful representation of $\mathbb{Z}_{2,0}^3$ into $\mathbb{R}^6$ which we denote by $(\rho, \mathbb{R}^6)$.

Similarly, for each permutation $\sigma \in S_3$, we associate an orientation preserving automorphism of $\mathbb{R}^6$ that takes e_{2i-1} to $e_{2\sigma(i)-1}$ and e_{2i} to $e_{2\sigma(i)}$ for each $1 \leq i \leq 3$. For example, $(132) \in S_3$ is the orientation preserving automorphism whose action is defined below and extended to all of $\mathbb{R}^6$ by linearity.

$$\begin{array}{ll} e_1 \to e_5 & e_2 \to e_6 \\ e_3 \to e_1 & e_4 \to e_2 \\ e_5 \to e_3 & e_6 \to e_4 \end{array}$$

The above map sends V_1 to V_3, V_2 to V_1, and V_3 to V_2. This defines a representation of S_3 into $\mathbb{R}^6$ which we denote by $(\tau, \mathbb{R}^6)$.

The representations τ and ρ are compatible in the sense that for each $\sigma \in S_3$, $\tau(\sigma)$ normalizes the image of ρ. Specifically, for each $\sigma \in S_3$ and $a \in \mathbb{Z}_{2,0}^3$, we have $\rho(\sigma a \sigma^{-1}) = \tau(\sigma)\rho(a)\tau(\sigma^{-1})$ and ρ and σ have complementary images. Hence, we may extend $(\rho, \mathbb{R}^6)$ to a representation of $\mathbb{Z}_{2,0}^3 \rtimes S_3$ on $\mathbb{R}^6$ which we denote by ρ'.

To complete the construction, we choose the typical faithful 6-dimensional representation $(\pi, \mathbb{R}^6)$ of $\mathbb{Z}_3^3$ into $\mathbb{R}^6$ where each i-th factor of $\mathbb{Z}_3^3$ acts on V_i by a rotation of $2\pi/3$-radians. The representation $(\pi, \mathbb{R}^6)$ is normalized by ρ' and has image complementary to ρ', and thus extends to a faithful representation of G of real dimension 6. This is minimal as $\mathbb{Z}_3^3$ requires at least three 2-dimensional irreducible pieces. □

What is remarkable about this example is that the minimal real representation does not come from 'real-ifying' any minimal complex representation as in the proof of Theorem 5.2. Because every minimal faithful representation of G is non-real, adding its conjugate piece would produce a $4 + 4 = 8$-dimensional real faithful representation. However, we managed to construct a 6-dimensional one. From Theorems 5.1 and 5.2, where the minimal real dimension is equal to the dimension of the realification of a minimal complex representation, we see that the upper bound of $\mathrm{mdim}_{\mathbb{R}}(G)$ by the smallest dimension of the realification over all minimal faithful complex representations cannot be improved in general. Regardless, Theorem 5.3 illustrates that this bound is not optimal in all cases.

References

[1] W. Fulton and J. Harris. *Representation Theory: A First Course.* Graduate Texts in Mathematics. Springer New York, 1991.

[2] N. Jacobson. *Basic Algebra I.* Freeman and Company, 2nd edition, 1985.

[3] Alexander Moretó. On the minimal dimension of a faithful linear representation of a finite group. 2021.

[4] D. Wright. Degrees of minimal embeddings for some direct products. *American Journal of Mathematics*, 97(4):897–903, 1975.