AMSI VACATION RESEARCH SCHOLARSHIPS 2019–20

EXPLORE THE MATHEMATICAL SCIENCES THIS SUMMER

Graph-Encoded Manifolds

Steven Condell

Supervised by Dr Jonathan Spreer

The University of Sydney

Vacation Research Scholarships are funded jointly by the Department of Education and Training and the Australian Mathematical Sciences Institute.

AMSI AUSTRALIAN MATHEMATICAL SCIENCES INSTITUTE

Australian Government Department of Education

Abstract

Certain manifold triangulations admit a particular kind of encoding as a coloured graph, in a way that allows us to uniquely recover the triangulation from the graph. This graph is called a graph-encoded manifold (gem). Gems have additional combinatorial structure that general triangulations lack. We focus on gems corresponding to closed, orientable surfaces, and on a particular type of gem called a crystallisation. We use the combinatorial structure of gems to reduce questions about gems to questions about permutations, and hence prove results on the number of crystallisations of a given surface and the nature of symmetries in crystallisaitions. These results on surface crystallisations lay the foundations for finding more general results on surface gems and on gems of higher dimensional manifolds.

1 Introduction

A *surface* is any space that looks locally like $\mathbb{R}^2$ at every point, and more generally, a *d-manifold* is a space that looks locally like $\mathbb{R}^n$ at every point. Surfaces that are smooth and closed can be *triangulated*, meaning (informally) that they can be cut up into triangles in a way that preserves the structure of the surface. Triangulations can also be defined for general d-manifolds, using d-simplices as the unit pieces instead of triangles. For an introduction to surfaces and 3-manifolds, including closed surfaces, orientable surfaces and the classification of surfaces, see Meyerhoff (1992). For an introduction to triangulations, see Chapter 6 of Armstrong (1983).

Triangulations are useful as they give a finite combinatorial description of manifolds, but the space of triangulations as a whole is difficult to understand. For example, the question of what the average triangulation of a given manifold looks like is one that seems to be a very difficult problem. In this project, we restrict to a subset of triangulations that satisfy certain properties (which will be made precise in the next section), such that they can be uniquely encoded as a graph with coloured arcs. This graph encoding of a triangulation is what we call a *graph-encoded manifold*, or *gem*. Gems have additional combinatorial properties that are not present in general triangulations, and by exploiting these properties we can learn about gems in levels of detail that are not possible for general triangulations.

In this report, we focus on gems corresponding to closed, orientable surfaces, and in particular look at a specific kind of gem called a crystallisation. The study of crystallisations is justfied by the fact that any closed, connected manifold M admits a crystallisation (Pezzana (1974)).

We begin by formally defining a gem, before explaining how their combinatorial structure allows us to describe them using only a small number of permutations. We then use this permutation-based description to find the number of crystallisations that a given surface has. We go on to define isomorphism of gems and symmetries of gems, and consider in terms of our permutation-based description how to map between isomorphic gems. We finally explore a highly symmetric subset of crystallisations and completely describe this subset in terms of its size and symmetries.

1.1 Statement of authorship

Section 2 and Section 3 of this report present existing theory on graph-encoded manifolds corresponding to surfaces. This theory was presented to me by my supervisor, Dr Jonathan Spreer, and is largely based on the work of Pezzana and his colleagues (see Ferri et al. (1986) for a survey of this work). In the later sections of the report, new work is presented. This work was conducted under the guidance of Dr Spreer, by myself and three other research students: Timothy Earl Lapuz, Taylor Ruber and Max Tobin. In this report I have focused primiarily on the results that I worked on myself, but the results in Section 5 were obtained as part of a group effort. This report has been written by me.

2 Definition of a graph-encoded manifold and notation

Given a triangulation T of a surface S, we can construct the *dual graph* $\Gamma(T)$ of the triangulation by placing a node on each triangle in the triangulation, and connecting two nodes via an arc if and only if the corresponding triangles are joined via an edge. We can analogously define the dual graph of an d-manifold triangulation to be the graph obtained by placing a node on each d-simplex and connecting two nodes via an arc if and only if the corresponding n-simplices are joined via an $(d-1)$-simplex. This gives us a mapping from an arbitrary d-manifold triangulation to a graph (a set of nodes and a set of arcs). It follows from the construction of the dual graph that it must be $(d+1)$-regular; that is, it must have exactly $d+1$ arcs leaving each node.

The dual graph is a good candidate for a graph encoding of a manifold; however, perhaps unsurprisingly, the dual graph does not contain enough information to recover the triangulation from which it was derived. We move towards a better graph encoding by adding some additional information to our triangulation.

Definition 2.1 (Rainbow-coloured simplex)**.** Let $d \in \mathbb{N}$, and fix a 'rainbow' of d colours: $0, 1, \ldots, n-1$. A *rainbow-coloured simplex* is a d-simplex with its vertices coloured such that it has exactly one vertex of each of the d colours.

Definition 2.2 (Rainbow colouring, balanced triangulation)**.** Let T be a triangulation of a d-manifold, and fix a rainbow of d colours $0, 1, \ldots, n-1$. A *rainbow colouring* of T, if it exists, is a colouring of all vertices in T, such that each d-simplex in T is a rainbow-coloured simplex. (We require adjacent simplices to have the same vertex colouring on their common vertices.) If T has a rainbow colouring, we say that T is a *balanced triangulation*.

It is clear that not all triangulations are balanced; for example, any triangulation with self-identification of vertices in a d-simplex cannot be balanced, since it would be impossible to assign a distinct colour to each vertex. However, if we do have a balanced triangulation, then a rainbow colouring of that triangulation induces an arc colouring on the dual graph in the following way. Firstly, fix an arc in the dual graph. This arc corresponds to a $(d-1)$-simplex in the triangulation, and this $(d-1)$-simplex must have a distinct colour on each of its $d-1$ vertices, as otherwise the d-simplices that it is a part of could not be rainbow-coloured. Thus, if we list all colours that this $(d-1)$-simplex has on its vertices, there is a unique colour from our rainbow that is missing from this list. We assign the arc this unique colour in our arc colouring. Following this procedure for every arc in our dual graph, we have an arc colouring determined by the rainbow colouring of our triangulation. The dual graph, together with this arc colouring, is a graph-encoded manifold.

Definition 2.3 (Graph-encoded manifold (gem))**.** Let T be a balanced triangulation, and let $G = \Gamma(T)$ be its dual graph. Let B be a rainbow colouring of of T, and let C be the arc colouring of G induced by B as described above. Then (G, C) is a *graph-encoded manifold* corresponding to the triangulation T.

The arc colouring adds enough information to the dual graph to allow us to recover the original triangulation. (This can be done for surface triangulations by redrawing a triangle around each node and using the colouring of the arcs to see which edges must be matched up.)

3 From graphs to permutations

Having defined gems, we now begin to explore their rich combinatorial structure. From now on, we focus on gems corresponding to closed, orientable surfaces. In this section, we show that we can completely describe such gems using only a few permutations, and then, in a restricted case, show that one permutation is sufficient to completely describe our gem. This simple permutation-based description of gems makes them much easier to study and work with.

We consider the closed, orientable surface S of genus g and a balanced triangulation T of this surface, along with a rainbow colouring of T. We know that the Euler characteristic of the surface S

is $\chi(S) = v - e + t$, where v, e and t are the number of vertices, edges and triangles in T respectively. Each triangle has exactly 3 edges, and since S is closed, each edge must be a part of exactly two triangles. Combining this information, we find that $2e = 3t$, and so $\chi(S) = v - \frac{1}{2}t$. We also know that the Euler characteristic of any closed, orientable surface can be expressed in terms of its genus g as $\chi(S) = 2 - 2g$. Combining these expressions for $\chi(S)$, we find an expression for the number of triangles in T in terms of the genus of the surface and the number of vertices in T: $t = 4g + 2v - 4$.

We know that the dual graph $\Gamma(T)$ of our triangulation is 3-regular, and it is clear from the way the arc colouring is constructed that we must have exactly one arc of each of the colours $0, 1, 2$ leaving each node. We also have that since S is orientable, the dual graph $\Gamma(T)$ is a bipartite graph. (This is a natural generalisation of a result from Cavicchioli et al. (1980).) Using this information, we can see that it is possible to divide the $t = 4g + 2v - 4$ nodes in $\Gamma(T)$ into two sets of nodes, N_1 and N_2, such that each node in N_1 connects only to nodes in N_2 (and vice versa), and such that each set is of size $2g + v - 2$. Considering the arcs of a single colour c, we can see that each of the nodes in N_1 connects to exactly one of the nodes in N_2 via a c-coloured arc. This gives us a way to use permutations to describe the coloured dual graph.

To do this, we can arbitrarily number the nodes in N_1 from 0 up to $2g + v - 3$. Then, we can correspondingly label the nodes in N_2 by following the 0-coloured arc from each node in N_1 and giving the connected node in N_2 the same label. Using this labelling of nodes, the set of 1-coloured arcs can be seen as describing a permutation: a bijective mapping from the set $\{0, 1, \ldots, 2g + v - 3\}$ to itself, found by taking as input each node in N_1, and as output the node in N_2 it connects to via a 1-coloured arc. The same is true of the set of 2-coloured arcs. Note that, by construction, the 0-colored arcs simply describe the identity permutation. (See Figure 1 for an example of a gem labelled in this way, and the permutations corresponding to each colour.) Knowing the two permutations marked out by the 1- and 2-coloured arcs then describes the dual graph completely, and thus given the permutations we can reconstruct the original triangulation of the closed, orientable surface S.

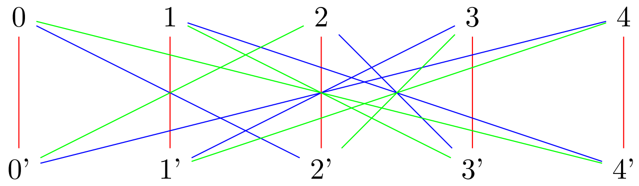

Figure 1: An example of a node-labelled gem, corresponding to a 10-triangle triangulation of the 2-holed torus. The 0-coloured (red) arcs correspond to the identity permutation, the 1-coloured (blue) arcs correspond to the permutation (0 2 3 1 4) and the 2-coloured (green) arcs correspond to the permutation (0 4 1 3 2).

We now restrict ourselves to triangulations of S with exactly three vertices ($v = 3$). This is the minimum number of vertices possible in a surface triangulation that has a gem, since we must have at least one vertex of each of the 3 colours. This notion of a minimum number of vertices generalises trivially to higher dimensions and motivates the following definition.

Definition 3.1 (Crystallisation)**.** Let M be a d-manifold, and suppose that T is a $(d+1)$-vertex balanced triangulation of M. Then, the gem corresponding to T is called a *crystallisation* of M.

The restriction that $v = 3$ gives us is that $t = 4g + 2$, and thus that the sets N_1 and N_2 each have size $2g+1$. More importantly, it gives us that the balanced triangulation T features exactly one vertex of each of the colours $0, 1, 2$. So, if $\{c_0, c_1, c_2\}$ is any arrangement of the three colours, then every edge in T that is opposite either a c_1- or c_2-coloured vertex must be adjacent to the single vertex of colour c_0. Hence, in the dual graph $\Gamma(T)$, all of the arcs of colours c_1 and c_2 are connected in a single cycle.

This restriction on the cycles present in $\Gamma(T)$ can be incorporated into the permutation description of the graph. Letting $\sigma_0, \sigma_1, \sigma_2$ be the permutations corresponding to colours $0, 1, 2$ respectively (where σ_0 is the identity), the single cycle condition gives the following restrictions on the gem:

- $\sigma_0^{-1} \circ \sigma_1 = \sigma_1$ is a single cycle permutation.
- $\sigma_0^{-1} \circ \sigma_2 = \sigma_2$ is a single cycle permutation.
- $\sigma_1^{-1} \circ \sigma_2$ is a single cycle permutation.

Now, recall that the labelling of the nodes in N_1 was chosen arbitrarily. Thus, since we know that σ_1 must be a single cycle permutation, we can choose a labelling of the nodes in N_1 (and a corresponding labelling of the nodes in N_2 such that 0-coloured arcs still map out the identity) such that $\sigma_1 = \sigma_* := (0\ 1\ \dots\ 2g)$. (We could choose σ_1 to be any single cycle permutation of length $2g+1$;

we simply choose this permutation as it is relatively easy to work with.) The permutation σ_2 is then fixed by this labelling. Using this labelling convention, the permutation σ_2 completely describes the dual graph (and thus the triangulation). (See Figure 2 for a relabelling of Figure 1 according to this labelling convention.)

The fact that we can always choose a labelling of this form allows us to make the following definition.

Definition 3.2 (Standard form)**.** Suppose (G, C) is a crystallisation of the closed, orientable surface of genus g, where $G = ((N_1, N_2), E)$ is a bipartite graph. Then the gem is said to be positioned in *standard form* if the node labelling is such that the 0-coloured arcs map out the identity permutation from N_1 to N_2 and the 1-coloured arcs map out the permutation $\sigma_* = (0\ 1\ \ldots\ 2g)$ from N_1 to N_2.

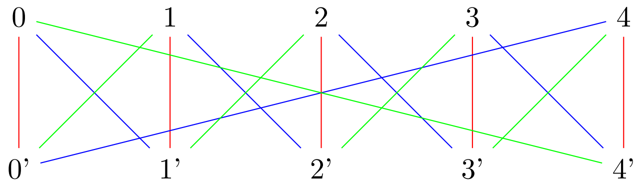

Figure 2: A relabelling of Figure 1 such that the 1-coloured (blue) permutation becomes (0 1 2 3 4). The 2-coloured (green) permutation here becomes (0 4 3 2 1).

The above tells us that if we have a valid 3-vertex triangulation T of the closed, orientable surface S and position $\Gamma(T)$ in standard form, the permutation σ_2 corresponding to colour 2 must be such that both σ_2 and $(0\ 1\ \ldots\ 2g)^{-1} \circ \sigma_2$ are single cycle permutations. Conversely, if we have a permutation σ_2 satisfying these two conditions, we can construct a gem corresponding to a balanced 3-vertex triangulation of a closed, orientable surface. The above is summarised in the following theorem.

Theorem 3.1 (Crystallisations of closed orientable surfaces)**.** *Suppose we have a surface gem* (G, C)*, where* $G = ((N_1, N_2), E)$ *is a bipartite graph with* $4g + 2$ *vertices for some* $g \in \mathbb{N}$*. Then* $|N_1| = |N_2| = 2g + 1$*. Suppose further that there exists a labelling of the nodes such that colour* 0 *maps out the identity permutation from* N_1 *to* N_2 *and colour* 1 *maps out the permutation* $\sigma_* := (0\ 1\ \ldots\ 2g)$ *from* N_1 *to* N_2*. Fix this labelling and let* σ *be the permutation from* N_1 *to* N_2 *mapped out by colour* 2*.*

Then (G, C) *is a crystallisation of the closed, orientable surface of genus* g *if and only if:*

- σ *is a single cycle permutation, and*
- $(2g\ 2g-1\ \ldots\ 1\ 0) \circ \sigma = \sigma_*^{-1} \circ \sigma$ *is a single cycle permutation.*

We finish this section with a definition motivated by the above theorem.

Definition 3.3 (Valid permutation)**.** Let $g \in \mathbb{N}$, and let $\sigma \in S_{2g+1}$. σ is called a *valid permutation* if:

- σ is a single cycle permutation, and
- $(2g\ 2g-1\ \dots\ 1\ 0) \circ \sigma = \sigma_*^{-1} \circ \sigma$ is a single cycle permutation.

4 Number of valid permutations of a given size

Upon defining a valid permutation, the first natural question to ask is how many valid permutations of a given size there are. To answer this question, we use a result from Boccara (1980).

Lemma 4.1 (Number of representations of a permutation as the product of two n-cycles)**.** Let $\sigma \in S_n$ be a permutation with a total of k cycles of respective lengths $n_1, n_2, \dots, n_k$. Define the function $\psi(\sigma, x) = \Pi_{i=1}^{k}(x^{n_i} - (x-1)^{n_i})$. Let $\gamma_n(l, m, \sigma)$ be the number of ways σ can be expressed as the product of an l-cycle and an m-cycle. Then

$$\gamma_n(n, n, \sigma) = (n-1)! \int_0^1 \psi(\sigma, x)\ dx.$$

Proposition 4.1 (Number of valid permutations of a given size)**.** *Let $g \in \mathbb{N}$, and let $n = 2g+1$. Let* $\text{valid}(n)$ *be the number of valid permutations in S_n. Then*

$$\text{valid}(n) = \frac{2 \cdot (n-1)!}{n+1} = \frac{(2g)!}{g+1}.$$

Proof. We first construct a bijection between the set of valid permutations of size n and the set of pairs of n-cycles with product $(0\ 1\ \dots\ n-1)$.

Let α be a valid permutation of size n, and let $\beta = (0\ 1\ \dots\ n-1)^{-1} \circ \alpha$. Then α and β are both n-cycles. From the expression for β, we have that $\beta^{-1} = \alpha^{-1} \circ (0\ 1\ \dots\ n-1)$, and so $\alpha \circ \beta^{-1} = (0\ 1\ \dots\ n-1)$. Thus, for every valid permutation α, we have a pair of n-cycles (α, β^{-1}) with product $(0\ 1\ \dots\ n-1)$.

Conversely, let α, β^{-1} be two n-cycles satisfying $\alpha \circ \beta^{-1} = (0\ 1\ \dots\ n-1)$. Then $\beta^{-1} = \alpha^{-1} \circ (0\ 1\ \dots\ n-1)$, and so $(0\ 1\ \dots\ n-1)^{-1} \circ \alpha = \beta$. Thus, since α and β are both n-cycles, α is a valid permutation of size n.

So, there is a bijection between the two sets, and therefore the number of valid permutations of size n is equal to the number of pairs of n-cycles with product $(0\ 1\ \dots\ n-1)$. Using the notation from Lemma 4.1, this bijection gives us that $\text{valid}(n) = \gamma_n(n, n, (0\ 1\ \ldots\ n-1))$. $(0\ 1\ \ldots\ n-1)$ is an n-cycle, so $\psi((0\ 1\ \ldots\ n-1), x) = x^n - (x-1)^n$. Thus,

$$\begin{aligned}\text{valid}(n) &= (n-1)! \int_0^1 (x^n - (x-1)^n)\ dx \\ &= (n-1)! \cdot \frac{1 + (-1)^{n+1}}{n+1} \\ &= (n-1)! \cdot \frac{2}{n+1}\end{aligned}$$

where the last equality holds since $n = 2g+1$ is odd.

□

5 Isomorphism of crystallisations

The previous section gave us an expression for the number of valid permutations of a given size. This result may be interpreted as saying that there are $\frac{2 \cdot (n-1)!}{n+1}$ crystallisations of closed, orientable surface of genus g. However, this ignores the fact that, given a 3-vertex triangulation T, we have a degree of choice when determining the labelling of $\Gamma(T)$, and the choices we make can affect the permutation we end up with. So, we could have 2 crystallisations described by different permutations while being fundamentally the same. There are 3 kinds of choices we have in the labelling process:

1. Consider the stage in the labelling process where we must choose a labelling of the nodes in N_1 (and a corresponding labelling of the nodes in N_2 such that 0-coloured arcs still map out the identity) such that the 1-coloured permutation becomes $\sigma_* = (0\ 1\ \ldots\ n-1)$. Given one such labelling, we can construct a new labelling by considering the standard form depiction of the graph and cyclically shifting all of the node labels backward by k steps, so that the 0-coloured arc that began with label i now has label $i + k \pmod n$, for $0 \leq i \leq n-1$. This cyclic shift can be performed for any k satisfying $0 \leq k \leq n-1$ to produce n possible labelling choices. We call the operation of cyclically shifting all node labels backward by k steps a *rotation* of our gem.

2. In an early stage of the labelling process, we choose to regard our coloured arcs as functions from N_1 to N_2, but this choice of direction is arbitrary. To find the what the gem would look like had we chosen the other direction, we can perform a *reflection* on our graph: we "flip" the standard form depiction of our graph to interchange N_1 and N_2, and we then must relabel our our nodes to put the graph back into standard form.

3. When putting our gem into standard form, we labelled it so that the 0-coloured arcs mapped out the identity permutation (call this position 0), the 1-coloured arcs mapped out the permutation σ_* (call this position 1), and the 2-coloured arcs mapped out the valid permutation (call this position 2). However, we could have chosen any of the 3! orderings of the colours and positioned them in that order. To find out what the gem would look like had we chosen another colour ordering, we can simply recolour the graph according to this ordering and then relabel nodes to bring the graph into standard form. We call this operation a *colour swap* of our gem.

By analysing the entire labelling process, it is clear that the above 3 kinds of labelling choices are the only ones we have. Thus, the corresponding 3 kinds of operations can be composed to construct all possible valid labellings of our crystallisation. The total number of operations is $n \cdot 2 \cdot 3! = 12n$, where the three factors come from the number of rotations, reflections and colour swaps respectively. The above argument leads to the following formal definition of isomorphism in the context of crystallisations, as well as the concept of a symmetry.

Definition 5.1 (Isomorphism of crystallisations, isomorphism class)**.** Two surface crystallisations are called *isomorphic* if one can be reached from the other by some sequence of rotations, reflections and colour swaps. We also say that the valid permutations corresponding to to the two crystallisations are isomorphic. The *isomorphism class* of a crystallisation (permutation) is the set of all crystallisations (permutations) that are isomorphic to it.

Definition 5.2 (Symmetry)**.** A sequence of rotation, reflection and colour swap operations is called a *symmetry* of a crystallisation (permutation) if its image under the sequence is itself.

5.1 Examining isomorphisms using only permutations

We wish to be able to perform rotation, reflection and colour swap operations using only the permutations involved without having to consider the graph directly, so we can easily find the image of a given permutation under a given sequence of operations. This would also give us a simple method for constructing the isomorphism class of a given permutation, by performing all possible sequences of operations on it. In order to achieve this, we first make the following definitions.

Definition 5.3 (Relabelling)**.** Let $\tau = (\tau_0\ \tau_1\ \dots\ \tau_{n-1}) \in S_n$ be a single cycle permutation, and let $\rho \in S_n$. Then the *relabelling of* τ *under* ρ is the permutation $(\rho(\tau_0)\ \rho(\tau_1)\ \dots\ \rho(\tau_{n-1}))$.

Definition 5.4 (Relabelling map)**.** Let $\sigma_{old} = (0\ (\sigma_{old})_1\ \dots\ (\sigma_{old})_{n-1}) \in S_n$ and $\sigma_{new} = (0\ (\sigma_{new})_1\ \dots\ (\sigma_{new})_{n-1})$ S_n be single cycle permutations, each expressed in cycle notation starting from 0. Let ρ be the permutation mapping $(\sigma_{old})_i$ to $(\sigma_{new})_i$ for $i \in \{0, 1, \dots, n-1\}$. Then ρ is called the *relabelling map specified by the ordered pair* $(\sigma_{old}, \sigma_{new})$.

The above definitions allow us to state the following proposition on performing gem operations using only permutations.

Proposition 5.1 (Permutation-based gem operations)**.** *Let* $\sigma = (\sigma_0\ \sigma_1\ \dots\ \sigma_{n-1}) \in S_n$ *be a valid permutation. Then:*

1. *The* n *images of* σ *under rotation are given by* $((\sigma_0 + i)\ (\sigma_1 + i)\ \dots\ (\sigma_{n-1} + i))$ *for* $i \in \{0, 1, \dots, n-1\}$*, where addition is mod* n*.*
2. *The image of* σ *under reflection is given by the relabelling of* σ^{-1} *under the relabelling map specified by* $(\sigma_*^{-1}, \sigma_*)$*.*
3. *Let* $C \in S_3$ *be one of the* $3!$ *colour swaps, where for* $i = 0, 1, 2$*,* $C(i)$ *describes the position that colour* i *is in after the colour swap.*
   - *Let* $\sigma_0 = id$*,* $\sigma_1 = \sigma_*$ *and* $\sigma_2 = \sigma$*.*
   - *Let* $\alpha_i = \sigma_{C^{-1}(i)}$ *for* $i = 0, 1, 2$*.*
   - *Let* $\beta_1 = \alpha_0^{-1} \circ \alpha_1$ *and* $\beta_2 = \alpha_0^{-1} \circ \alpha_2$*.*
   - *Then the image of* σ *under the colour swap* C *is given by the relabelling of* β_2 *under the relabelling map specified by* (β_1, σ_*)*.*

This proposition follows fairly straightforwardly from the definitions of the rotation, reflection and colour swap operations: we must transform the gem and reposition it into standard form, keeping track as we go of how the permutations corresponding to each colour change. The details of the proof are given in the appendix.

We would like to be able to completely understand the space of isomorphism classes for all n, but this is a more complex task than finding the number of valid permutations of size n. At the heart of this complexity is the nature of symmetries in our crystallisations. We wish to understand when and why symmetries arise. Understanding a particular highly symmetric subset of crystallisations is the focus of the next section.

6 Constant difference cycle permutations

In this section, we look to understand a particular subset of valid permutations that turn out to have a large amount of symmetry. To define this subset, we first define the *difference cycle* of a permutation.

Definition 6.1 (Difference cycle)**.** Let $\tau = (\tau_0\ \tau_1\ \ldots\ \tau_{n-1}) \in S_n$ be a single cycle permutation, where $\tau_0 = 0$. Then the *difference cycle of* τ is the n-tuple $(d_0, d_1, \ldots, d_{n-1})$, where $d_i \equiv \tau_{i+1} - \tau_i$ for $0 \leq i \leq n-1-$.

We want to consider just those permutations that have a *constant difference cycle*; that is, the differences d_i defined above are all the same constant d. We first examine which constant difference cycle permutations are valid, and then go on to explore in detail the symmetries of constant difference cycle permutations.

6.1 Valid differences

We first define a *valid difference*, then answer the question of which constant difference cycles correspond to valid permutations in the following proposition.

Definition 6.2 (Valid difference)**.** Let $n \in \mathbb{N}$, and let $d \in \{1, 2, \ldots, n-1\}$. Then d is a *valid difference w.r.t.* n if the constant difference cycle of length n, $(d, d, \ldots, d)$, corresponds to a valid permutation of length n.

Proposition 6.1 (Set of valid differences)**.** *Let P_n be the set of prime divisors of n. Let $D_n = \{k \in \mathbb{Z} : 0 \leq k < n,\ k \not\equiv 0, 1\ (mod\ p),\ \forall p \in P_n\}$.*

Then d is a valid difference w.r.t. n if and only if $d \in D_n$.

The proof of this result relies on showing that the two cycle conditions given in the definition of a valid permutation are, for constant difference cycles of length n, equivalent to the constant difference d satisfying $d \not\equiv 0, 1 \pmod{p}$ for all primes divisors p of n. The proof is given in full in the appendix.

Having seen which differences d are valid in terms of the set D_n, we look to find the size of D_n in the following proposition. A sketch of a proof is provided in the appendix.

Proposition 6.2 (Size of D_n)**.** *Let $n = \prod_{i=1}^{k} p_i^{\alpha_i}$ be a prime factorisation of n, where $\alpha_i > 0$ for all $1 \leq i \leq k$. As in Proposition 6.1, let P_n be the set of prime divisors of n, and let $D_n = \{k \in \mathbb{Z} : 0 \leq k < n,\ k \not\equiv 0, 1\ (mod\ p),\ \forall p \in P_n\}$.*

Then $\varphi_2(n) := |D_n| = \prod_{i=1}^{k}(p_i - 2)p_i^{\alpha_i - 1}$.

Combining the previous two propositions, we get the following result on the number of valid differences w.r.t n.

Corollary 6.1 (Number of valid differences)**.** The number of valid differences w.r.t. n is exactly $\varphi_2(n) = \prod_{i=1}^{k}(p_i - 2)p_i^{\alpha_i - 1}$.

6.2 Symmetries in constant difference cycle permutations

Having found which differences d are valid, we now explore the symmetries present in permutations with constant difference cycles. We first consider rotation. Any cyclic shift of elements does not affect the differences between consecutive elements, it only shifts these differences along to some new position. But since the differences between each pair of consecutive elements are the same, we have that the difference cycle of a constant difference cycle permutation does not change under any cyclic shift, and thus we have full rotational symmetry.

We now consider reflection. Let $\sigma = (0\ d\ 2d\ \dots\ (n-1)d)$ be the valid permutation with constant difference d. By part 2 of Proposition 5.1, the reflection of this permutation is the relabelling of σ^{-1} under the relabelling map ρ specified by $(\sigma_*^{-1}, \sigma_*)$. σ^{-1} is the permutation $(0\ (n-1)d\ (n-2)d\ \dots\ d)$, and it can be shown that ρ is the map that sends i to $-i \pmod n$ for $0 \leq i < n$. Thus, the relabelling of $sigma^{-1}$ under ρ is $(0\ d\ 2d\ \dots\ (n-1)d)$, which is σ, and so the constant difference cycle permutation σ has reflectional symmetry.

This means that each permutation with a constant difference cycle is guaranteed to have at least $2n$ symmetries, and thus to be in an isomorphism class of size at least $2n$.

We now focus our attention on the colour swaps of permutations with constant difference cycles; these are the only kind of symmetries that are not automatic for constant difference cycle permutations. We begin with a result on the images of a constant difference cycle permutation σ under each of the 5 non-trivial colour swaps. This result follows from part 3 of Proposition 5.1; we can use the fact that we are dealing with constant different cycle permutations to explicitly perform the required compositions and relabellings. The full proof is provided in the appendix.

Proposition 6.3 (Images of constant difference cycles under colour swaps)**.** *Let $\sigma \in S_n$ be a valid permutation of length n that has a difference cycle with constant difference d. Then, each colour swap maps σ to another permutation with a constant difference cycle. Moreover, the valid difference corresponding to the image under each colour swap is as follows:*

	Colour swap $C \in S_3$	***Constant difference of image of*** σ ***under*** C
(i)	$(0\ 1)$	$1 - d\ (mod\ n)$
(ii)	$(1\ 2)$	$d^{-1}\ (mod\ n)$
(iii)	$(0\ 2\ 1)$	$(1-d)^{-1}\ (mod\ n)$
(iv)	$(0\ 1\ 2)$	$1 - d^{-1}\ (mod\ n)$
(v)	$(0\ 2)$	$1 - (1-d)^{-1}\ (mod\ n)$

Now that we are able to find the images of constant difference cycle permutations under each colour swap, we can easily find exactly when constant difference colour swap symmetries occur.

Proposition 6.4 (Symmetries of constant difference cycles under colour swaps)**.** *Let $\sigma \in S_n$ be a valid permutation of length n that has a difference cycle with constant difference d. Then, the conditions for symmetry under each colour swap are as follows:*

	Colour swap $C \in S_3$	σ ***has symmetry under*** C ***iff:***
(i)	$(0\ 1)$	$d \equiv 2^{-1} \equiv \frac{n+1}{2}\ (mod\ n)$
(ii)	$(1\ 2)$	$d \equiv -1\ (mod\ n)$
(iii)	$(0\ 2)$	$d \equiv 2\ (mod\ n)$
(iv)	$(0\ 1\ 2)$ *or* $(0\ 2\ 1)$	$d^2 - d + 1 \equiv 0\ (mod\ n)$

The proof of this proposition, given in the appendix, relies on simply taking the difference of the relevant image from the previous proposition, and equating this difference with the initial difference d to find a condition for equality.

Now that we have found necessary and sufficient conditions for each of the colour swap symmetries to occur for constant difference cycle permutations, we are in a position to be able to completely describe the symmetries that occur across all constant difference cycle permutations of a given size, and how often they occur.

Theorem 6.1 (Distribution of symmetries among constant difference cycle permutations)**.** *Let $n = 2g+1$, where $g \in \mathbb{N}$. Let $n = \left(\prod_{i=1}^{k} p_i^{\alpha_i}\right)\left(\prod_{i=1}^{l} q_i^{\beta_i}\right)(3^{\gamma})$ be the prime factorisation of n, where the p_i's are distinct and are the prime divisors of n that are $1\ (mod\ 3)$, the q_i's are distinct and are the prime divisors of n that are $2\ (mod\ 3)$, $\alpha_i > 0$ for $1 \leq i \leq k$ and $\beta_i > 0$ for $1 \leq i \leq l$.*

If $n \leq 3$, then there is only one constant difference cycle permutation, with the full $12n$ symmetries. Otherwise, if $n > 3$, we have the following:

- *The constant difference cycle permuations with differences of 2, $\frac{n+1}{2}$ and $n-1$ are always valid.*

These 3 *permutations each have exactly* $4n$ *symmetries, and together they form a single isomorphism class of size* 3.

- *If* $l > 0$ *or* $\gamma > 1$*, there are no constant difference cycle permutations with* $6n$ *symmetries. Otherwise, there are exactly* 2^k *constant difference cycle permutations with exactly* $6n$ *symmetries, and these permutations form* 2^{k-1} *isomorphism classes, each of size* 2.
- *All remaining constant difference cycle permutations have exactly* $2n$ *symmetries, and are members of isomorphism classes of size* 6.

The proof of this theorem is non-trivial, but relies mostly on a careful examination of the solutions to the four congruences listed in Proposition 6.4, as well as using Proposition 6.3 to find operations that transform these solutions to one another. The most involved step is solving the congruence $d^2 - d + 1 \equiv 0 \pmod{n}$; this congruence is dealt with in a lemma which is presented in the appendix along with the full proof of the theorem.

We finish this section with a brief corollary that follows directly from the above theorem.

Corollary 6.2 (Crystallisations with full symmetry)**.** For $g > 1$, there are no perfectly symmetric crystallisations of the closed, orientable surface S of genus g. That is, no crystallisations of S have the full $12n$ possible symmetries.

7 Conclusion

In this report, we explored the crystallisations of closed, orientable surfaces. We showed how a crystallisation of the closed, orientable surface of genus g can be described using a single permutation, and used this permutation-based description to find the number of crystallisations of a given size and analyse the isomorphism types and symmetries of crystallisations.

Using the concept of a difference cycle, we analysed a highly symmetric subset of crystallisations, those corresponding to permutations with constant difference cycles, and completely described their symmetries. Another, slightly more general highly symmetric subset of crystallisations are those corresponding to permutations with difference cycles $(d_0, d_1, \ldots, d_{n-1})$ such that $d_i \equiv d_{i+3} \pmod{n}$ for $0 \leq i < n$. These crystallisations have been analysed by myself and others using methods similar to those used in the constant difference case, to give a partial characterisation of the symmetries of gems that have at least $\frac{n}{3}$ rotational symmetries. It is much more difficult to analyse the symmetries of general permutations, but future work could explore different ways of exploiting the combinatorial structure of crystallisations to characterise the symmetries of more general classes of these objects.

More generally, future work could explore ways of generalising the results in Section 4 and Section 5 to surface gems with more vertices. Eventually, we would like to be able to use this understanding of surface gems to build an understanding of gems corresponding to 3-manifolds.

8 Acknowledgements

I would like to thank my supervisor, Dr Jonathan Spreer, for his support and guidance throughout this project. I would also like to thank my partners in this research project, Timothy Earl Lapuz, Taylor Ruber and Max Tobin, all of whom made this research experience a more positive one. Lastly, I would like to thank AMSI for funding and supporting this project.

References

Armstrong, M. A. (1983), *Basic Topology*, Springer.

Boccara, G. (1980), 'Nombre de représentations d'une permutation comme produit de deux cycles de longueurs données', *Discrete Math.* **29**(2), 105–134.
URL: *https://doi.org/10.1016/0012-365X(80)90001-1*

Cavicchioli, A., Grasselli, L. & Pezzana, M. (1980), 'A normal decomposition for closed n-manifolds', *Boll. Un. Mat. Ital. B (5)* **17**(3), 1146–1165.

Ferri, M., Gagliardi, C. & Grasselli, L. (1986), 'A graph-theoretical representation of pl-manifolds — a survey on crystallizations', *aequationes mathematicae* **31**(1), 121–141.
URL: *https://doi.org/10.1007/BF02188181*

Meyerhoff, R. (1992), 'Geometric invariants for 3-manifolds', *The Mathematical Intelligencer* **14**(1), 37–53.
URL: *https://doi.org/10.1007/BF03024140*

Pezzana, M. (1974), 'Sulla struttura topologica delle varietà compatte', *Ati Sem. Mat. Fis. Univ. Modena* **23**(1), 269–277 (1975).

A Additional proofs

Proof of Proposition 5.1. 1. The images of σ under rotation are clear from the definition of rotation.

2. Let τ_i be the permutation corresponding to the i-coloured arcs at the current stage of the rearrangement, for $i = 0, 1, 2$.

   We begin with a gem labelling such that the arcs of each colour are regarded as a function from N_1 to N_2, where $\tau_0 = id$, $\tau_1 = \sigma_*$ and $\tau_2 = \sigma$. To perform a reflection, we must flip the graph, meaning we swap the roles of N_1 and N_2. This inverts all of the permutations, so in the reflected graph, we have $\tau_0 = id$, $\tau_1 = \sigma_*^{-1}$ and $\tau_2 = \sigma^{-1}$.

   After performing the flipping, we must move the crystallisation into standard form. In order to preserve τ_0 as the identity, we can fix N_2's labels as completely dependent on N_1's labels: a node in N_2 is given the label j if and only if it is connected via a 0-coloured arc to the node in N_1 labelled j. This allows us to think of the 0-coloured arcs themselves as having a label $j \in \{0, 1, \ldots, n-1\}$, corresponding to the label of both its adjacent nodes. Then, to move into standard form, we must relabel the 0-coloured arcs (and hence relabel the corresponding pairs of nodes) such that τ_1, which is currently σ_*^{-1}, becomes σ_*. We define ρ to be the relabelling map specified by $(\sigma_*^{-1}, \sigma_*)$, and apply this relabelling to all 3 permutations. This gives us a crystallisation in standard form, since after this relabelling, τ_0 remains the identity and τ_1 becomes σ_* by construction. Thus, the new value of τ_2 (the relabelling of σ^{-1} under ρ) is the image of σ under reflection.

3. We fix a colour swap C as specified above. To perform the recolouring, we first simultaneously change the $C^{-1}(i)$-coloured arcs to i-coloured arcs for $i = 0, 1, 2$. Then, defining the τ_i's in the same way as in the proof of part 2, we have $\tau_i = \alpha_i$ for $i = 0, 1, 2$, where the α_i's are defined as in the statement of the proposition.

   As a first step towards bringing this graph into standard form, we relabel the nodes in N_2 so that τ_0 becomes the identity. This can be achieved by, for $j = 0, 1, \ldots, n-1$, taking the node in N_2 originally labelled j and relabelling it as $\alpha_0^{-1}(j)$. So, for all $j \in \{0, 1, \ldots, n-1\}$, the node in N_1 with label j, which is originally mapped to the node with label $\alpha_i(j)$ in N_2, is now mapped to the node with label $\alpha_0^{-1}(\alpha_i(j))$, for $i = 0, 1, 2$. This means that, after this relabelling, $\tau_0 = id$, $\tau_1 = \beta_1 := \alpha_0^{-1} \circ \alpha_1$ and $\tau_2 = \beta_2 := \alpha_0^{-1} \circ \alpha_2$.

   We again wish to preserve τ_0 as the identity, so we fix N_2's labels as completely dependent on N_1's labels as in the proof of part 2. Then, to move into standard form, we must relabel the 0-coloured arcs (and hence relabel the corresponding pairs of nodes) such that τ_1, which is currently β_1, becomes σ_*. We define ρ to be the relabelling map specified by (β_1, σ_*), and apply this relabelling to all 3 permutations. This gives us a crystallisation in standard form, since after this relabelling, τ_0 remains the identity and τ_1 becomes σ_* by construction. Thus, the new value of τ_2 (the relabelling of β_2 under ρ) is the image of σ under the colour swap C.

□

The following is a lemma used in the proof of Proposition 6.1.

Lemma A.1 (Composition of permutations with constant difference cycles)**.** Let σ_1 and σ_2 be two permutations with constant difference cycles, with differences d_1 and d_2, respectively. Then, in the cycle representation of $\sigma_2 \circ \sigma_1$, the difference between any two consecutive elements is $d_1 + d_2 \pmod n$.

Proof. Define σ_1 and σ_2 as above. Then, $\sigma_k(i) \equiv i + d_k \; (\text{mod} n)$, for $k = 1, 2$ and for $i = 0, 1, \ldots, n-1$. Thus, $\sigma_2(\sigma_1(i)) \equiv i + (d_1 + d_2) \pmod n$, for $i = 0, 1, \ldots, n-1$. This means that $\sigma_2 \circ \sigma_1$ maps each element to itself plus $d_1 + d_2$, and so the difference between any pair of consecutive elements in the cycle representation of $\sigma_2 \circ \sigma_1$ is $d_1 + d_2$. □

Proof of Proposition 6.1. We first prove the forward implication, and then the reverse implication.

- Suppose d is a valid difference w.r.t. n. Then, letting σ be the permutation in S_n with difference cycle $(d, d, \ldots, d)$, we have that σ is a valid permutation. This means by definition that both σ and $\sigma_*^{-1} \circ \sigma = (n-1 \; n-2 \; \ldots \; 1 \; 0) \circ \sigma$ are single n-cycles. So, σ can be expressed in cycle notation as $(0 \; d \; 2d \; \ldots \; (n-1)d)$, where multiplication is mod n. By Lemma A.1, the difference between consecutive elements in the cycle representation of $\sigma_*^{-1} \circ \sigma$ is $d - 1$. So, since $\sigma_*^{-1} \circ \sigma$ is a single n-cycle, it has a constant difference cycle with difference $d - 1$ and can be expressed in cycle notation as $(0 \; (d-1) \; 2(d-1) \; \ldots \; (n-1)(d-1))$.

  Fix a prime divisor p of n. Then $n = kp$ for some integer $k < n$. Now, suppose $d \equiv 0 \pmod p$; then $d = lp$ for some $l \in \mathbb{Z}$. Multiplying both sides of this expression by k, we find that $kd = k(lp) = l(kp) = ln$, which implies that $kd \equiv 0 \pmod n$. But kd is the kth element in the cycle expression of σ above, and this cannot be 0 since each number can only occur once in the n-cycle. Thus, it cannot be that $d \equiv 0 \pmod p$.

  Next, suppose $d \equiv 1 \pmod p$; then $(d-1) = lp$ for some (new) $l \in \mathbb{Z}$. Multiplying both sides by k, we find that $k(d-1) = k(lp) = l(kp) = ln$, which implies that $k(d-1) \equiv 0 \pmod n$. But

$k(d-1)$ is the kth element in the cycle expression of $\sigma_*^{-1} \circ \sigma$ above, and this cannot be 0 since each number can only occur once in the n-cycle. Thus, it cannot be that $d \equiv 1 \pmod p$.

The above argument holds for all prime divisors p of n. Thus, by the definition of D_n, $d \in D_n$.

- Conversely, suppose that $d \in D_n$. Then, for all prime divisors p of n, $d \not\equiv 0 \pmod p$. This means that $\gcd(d, n) = 1$, and so each multiple kd of d, for $0 \leq k < n$, is distinct mod n. Using this distinctness, we can construct the n-cycle permutation $\sigma := (0\ d\ 2d\ \ldots\ (n-1)d)$. By construction, σ is the permutation that has a constant difference cycle with difference d.

  Note that, by Lemma A.1, the difference between consecutive elements in the cycle representation of $\sigma_*^{-1} \circ \sigma$ is $d-1$. Now, since $d \in D_n$, we also have that $d \not\equiv 1 \pmod p$ for all prime divisors p of n, so $\gcd(d-1, n) = 1$. Thus, each multiple $k(d-1)$ of $d-1$, for $0 \leq k < n$, is distinct mod n. This tells us that, if we have a permutation in cycle representation where all consecutive elements differ by $d-1$, that permutation must be a single n-cycle. Thus, $\sigma_*^{-1} \circ \sigma$ is a single n-cycle.

  As the two permutations σ and $\sigma_*^{-1} \circ \sigma$ are n-cycles, we have that σ is a valid permutation by definition. Thus, d, the constant difference corresponding to σ, is a valid difference.

□

Proof (sketch) of Proposition 6.2. As the name $\varphi_2(n)$ suggests, this function is very similar to Euler's totient function $\varphi(n)$. In a similar way to how one can prove that $\varphi(n) = \prod_{i=1}^{k}(p_i - 1)p_i^{\alpha_i - 1}$, we can prove that our formula for $\varphi_2(n)$ is correct using the following two steps:

- Prove that the formula holds for n of the form p^α where p is prime. Clearly is not possible for any integer to be congruent to both 0 and 1 mod p for any prime p, so we can simply find the number of integers from 0 up to $n = p^\alpha$ that are congruent to 0 or 1 mod p ($2p^{\alpha-1}$) and subtract this from n to find the size of D_n. This gives us $\varphi_2(p^\alpha) = p^\alpha - 2p^{\alpha-1} = (p-2)p^{\alpha-1}$, so the formula holds for $n = p^\alpha$.

- Prove that the function is multiplicative; that is, that $\varphi_2(mn) = \varphi_2(m)\varphi_2(n)$ for any pair of coprime integers m and n. To do this, we can consider two arbitrary coprime integers m and n, and consider the 3 sets D_m, D_n and D_{mn}. Then, using the fact that m and n are coprime, we can construct a bijection between the two sets D_{mn} and $D_m \times D_n$. This gives us that the sizes of the two sets are equal, thus proving multiplicativity: $|D_{mn}| = |D_m||D_n|$.

After proving the above two results, it directly follows that $\varphi_2(n) = \prod_{i=1}^{k}(p_i - 2)p_i^{\alpha_i - 1}$.

□

Proof of Proposition 6.3. (i) We begin by proving that the image of σ under the colour swap $C = (0\ 1)$ is the permutation with a constant difference of $1 - d \pmod n$. We use part 3 of Proposition 5.1 to find the image of σ under the C. Using the notation from this proposition, we have $\alpha_0 = \sigma_*$, $\alpha_1 = id$ and $\alpha_2 = \sigma$. Then, $\beta_1 = \sigma_*^{-1} \circ id = \sigma_*^{-1}$, and $\beta_2 = \sigma_*^{-1} \circ \sigma$. Since σ is a valid permutation with constant difference d, we know that β_2 must be a single n-cycle, and so by Lemma A.1 we have that β_2 is the constant difference cycle permutation with difference $d - 1$.

To get the image of σ under C, we must relabel β_2 according to the relabelling map ρ specified by $(\beta_1, \sigma_*) = (\sigma_*^{-1}, \sigma_*) = ((0\ n-1\ n-2\ \dots\ 2\ 1), (0\ 1\ 2\ \dots\ n-2\ n-1))$. The map ρ sends i to $-i \pmod n$ for $0 \le i < n$.

Now, suppose that we have a single cycle permutation $\tau = (\tau_0\ \tau_1\ \dots \tau_{n-1})$ that has a difference cycle with constant difference e. Then $\tau_{i+1} - \tau_i = e$ for $0 \le i < n$, where addition and subtraction is mod n. Using the description of ρ above, we see that the relabelling of τ under ρ is $(-\tau_0\ -\tau_1\ \dots - \tau_{n-1})$. The difference between any two consecutive elements in this relabelling is $(-\tau_{i+1}) - (-\tau_i) = -(\tau_{i+1} - \tau_i) = -e$, so the relabelling has a difference cycle with constant difference $-e \pmod n$. Applying this to β_2, we find that the relabelling of β_2 under ρ is the constant difference cycle permutation with constant difference $-(d-1) \equiv 1 - d \pmod n$.

(ii) Using a similar approach to the previous part, we consider performing the colour swap $C = (1\ 2)$ on σ using the notation of Proposition 5.1. We have $\alpha_0 = id$, $\alpha_1 = \sigma$ and $\alpha_2 = \sigma_*$. Then, $\beta_1 = id^{-1} \circ \sigma = \sigma$ and $\beta_2 = id^{-1} \circ \sigma_* = \sigma_*$. So, to get the image of σ under C, we must relabel σ_* according to the relabelling map ρ specified by (σ, σ_*). Since $\sigma_* = (0\ 1\ \dots\ n-1)$, ρ simply maps each number i in the cycle expression of σ to its position in the cycle expression, where 0 is placed at position 0. Thus, we can understand ρ by finding the position of each number in σ.

We first consider two consecutive numbers $i, i+1 \pmod n$, and find their relative positions in σ. Since σ has a difference cycle with constant difference d, it can be expressed as $\sigma = (0\ d\ 2d\ \dots\ (n-1)d)$. We must then have that $i = kd$ and $i + 1 = ld$ for some integers k and l. Then, the difference between $(i+1)$'s position and i's position in the cycle representation of σ is $l - k$. We also have that $(i+1) - i \equiv ld - kd \pmod n$; this implies that $1 \equiv (l-k)d$ and so $l - k \equiv d^{-1} \pmod n$. Thus, the difference between $(i+1)$'s position and i's position in the the cycle representation of σ is the constant $d^{-1} \pmod n$.

As the difference between the position of $(i+1)$ and i in σ is the same for all i, we can for all i express the difference between the position of i and the position of 0 as i times this difference: $i \cdot d^{-1} \pmod n$. Since 0 is known to be at position 0, this tells us that $i \cdot d^{-1} \pmod n$ is the position of i in σ for all i. So, $\rho(i) = i \cdot d^{-1} \pmod n$, for all i. Thus, the relabelling of $\sigma_* = (0\ 1\ 2\ \ldots\ n-1)$ under ρ is $(0\ d^{-1}\ 2d^{-1}\ \ldots\ (n-1)d^{-1})$, which is the constant difference cycle permutation with difference $d^{-1} \pmod n$.

(iii) The colour swap $C = (0\,2\,1)$ can be expressed as the product of cycles we have already considered: $(0\,2\,1) = (1\,2) \circ (0\,1)$. Thus, to find the image of σ under C, we can first find the image of σ under $(0\ 1)$, and then find the image of the result under $(1\ 2)$. From part (i), we know that the image of the constant difference permutation with difference d under $(0\ 1)$ is the constant difference permutation with difference $(1-d) \pmod n$. From part (ii), we know that the image of this permutation under $(1\ 2)$ is the constant difference permutation with difference $(1-d)^{-1} \pmod n$, since the $(1\ 2)$ swap simply inverts the difference we start with. This means that the image of σ under C is the constant difference permutation with difference $(1-d)^{-1} \pmod n$.

(iv) In a similar way to the previous part, we can write $C = (0\ 1\ 2) = (0\ 1) \circ (1\ 2)$. This tells us that the image of σ under C is the image of the constant difference cycle permutation with difference d^{-1} under $(0\ 1)$, which is the constant difference cycle permutation with difference $1 - d^{-1} \pmod n$.

(v) Using the same method as the previous parts, we write $C = (0\ 2) = (0\ 1) \circ (1\ 2) \circ (0\ 1)$, which tells us that the image of σ under C is the constant difference cycle permutation with difference $1 - (1-d)^{-1} \pmod n$.

□

Proof of Proposition 6.4. By Proposition 6.3, we know that constant difference cycle permutations map to constant difference cycle permutations under colour swaps. So, a symmetry under a given colour swap C exists if and only if the difference before the colour swap (d) is congruent mod n to the difference after the colour swap (e.g. $1-d$ for $C = (0\ 1)$). We solve this congruence in each case to find the conditions for symmetry:

(i) Symmetry under the $(0\ 1)$ colour swap exists iff $d \equiv 1-d \pmod n$. This congruence is equivalent to $2d \equiv 1 \pmod n$, which is equivalent to $d \equiv 2^{-1} \pmod n$. (2 must have an inverse mod n

since n must be odd for σ to exist, and this inverse can be expressed as $\frac{n+1}{2}$.)

(ii) Symmetry under the (1 2) colour swap exists iff $d \equiv d^{-1} \pmod{n}$. This congruence is equivalent to $d^2 \equiv 1 \pmod{n}$, so $(d-1)(d+1) \equiv 0 \pmod{n}$. By Proposition 6.1, if d is a valid difference, $d \not\equiv 1 \pmod{p} \implies (d-1) \not\equiv 0 \pmod{p}$ for any prime divisor p of n. Thus, $d-1$ and n are coprime and so $d-1$ has an inverse mod n. This means we can simplify $(d-1)(d+1) \equiv 0 \pmod{n}$ to $d+1 \equiv 0 \pmod{n}$, giving us the desired result.

(iii) Symmetry under the (0 2) colour swap exists iff $d \equiv 1-(1-d)^{-1} \pmod{n}$. This congruence is equivalent to $(d-1)^2 \equiv 1 \pmod{n}$, so $d(d-2) \equiv 0 \pmod{n}$. By Proposition 6.1, if d is a valid difference, $d \not\equiv 0 \pmod{p}$ for any prime divisor p of n. Thus, d and n are coprime and so d has an inverse mod n. This means we can simplify $d(d-2) \equiv 0 \pmod{n}$ to $d-2 \equiv 0 \pmod{n}$, giving us the desired result.

(iv) Symmetry under the (0 2 1) colour swap exists iff $d \equiv (1-d)^{-1} \pmod{n}$. This congruence is equivalent to $d(1-d) \equiv 1 \pmod{n}$, so $d^2-d+1 \equiv 0 \pmod{n}$. Similarly, symmetry under the (0 1 2) colour swap exists iff $d \equiv 1-d^{-1} \pmod{n}$. This congruence is equivalent to $d^2 \equiv d-1 \pmod{n}$, so $d^2-d+1 \equiv 0 \pmod{n}$.

□

The following is a lemma that aids in the proof of Theorem 6.1.

Lemma A.2 (Number of solutions to $d^2-d+1 \equiv 0 \pmod{n}$)**.** Let $n = \left(\prod_{i=1}^{k} p_i^{\alpha_i}\right)\left(\prod_{i=1}^{l} q_i^{\beta_i}\right)(3^{\gamma})$ be the prime factorisation of n, where the p_i's are distinct and are the prime divisors of n that are $1 \pmod{3}$, the q_i's are distinct and are the prime divisors of n that are $2 \pmod{3}$, $\alpha_i > 0$ for $1 \le i \le k$ and $\beta_i > 0$ for $1 \le i \le l$.

Then the number of solutions $d \in \{0, 1, \ldots, n-1\}$ to the congruence $d^2-d+1 \equiv 0 \pmod{n}$ is 0 if $l > 0$ or $\gamma > 1$, and is 2^k otherwise.

Proof. We want to find the solutions to the congruence $d^2-d+1 \equiv 0 \pmod{n}$. We begin by completing the square to simplify our congruence. $d^2-d+1 \equiv d^2-(n+1)d+1 \equiv d^2 - 2 \cdot \frac{n+1}{2} \cdot d + 1 \pmod{n}$, so completing the square we find that $(d-\frac{n+1}{2})^2 \equiv (\frac{n+1}{2})^2 - 1$. This can be simplified to become $(2d-n-1)^2 \equiv n^2+2n-3$, and terms in n can be removed to reduce to $(2d-1)^2 \equiv -3 \pmod{n}$. So, we must solve the quadratic congruence $y^2 \equiv -3 \pmod{n}$, with each solution y giving a unique solution $d \equiv 2^{-1}(y+1) \pmod{n}$.

To solve $y^2 \equiv -3 \pmod{n}$, we can first solve $y^2 \equiv -3 \pmod{p^\alpha}$ for all prime divisors p of n, where α is the number of times p occurs in the prime factorisation of n. Then, since the numbers p^α are pairwise coprime, we can combine the solutions via the Chinese Remainder Theorem to find the solution set mod n. The number of solutions to the congruence mod n is equal to the product of the number of solutions mod p^α for all prime factors p.

We first consider the case $p = 3$. The congruence $y^2 \equiv -3 \pmod{3}$ clearly has solutions exactly when $y \equiv 0 \pmod{3}$. However, by considering all squares mod 9, we find that $y^2 \equiv -3 \pmod{3^2}$ has not solutions, and so $y^2 \equiv -3 \pmod{3^\gamma}$ has no solutions for $\gamma > 1$.

To help solve the congruences in the remaining case, we use a known result on quadratic residues: -3 is a quadratic residue mod a prime p if and only if $p \equiv 1 \pmod{3}$. This tells us that, in the case $p \equiv 2 \pmod{3}$, the congruence $y^2 \equiv -3 \pmod{p}$ has no solutions. This means that $y^2 \equiv -3 \pmod{p^k}$ has no solutions for any k if $p \equiv 2 \pmod{3}$.

The QR result also tells us that, in the case $p \equiv 1 \pmod{3}$, the congruence $y^2 \equiv -3 \pmod{p}$ has exactly 2 solutions. It follows from Hensel's lifting lemma that $y^2 \equiv -3 \pmod{p^k}$ has exactly 2 solutions for all positive integers k.

Combining the above information, we have that the number of solutions to our congruence mod n is 0 if either 9 is a factor of n (i.e. $\gamma > 1$) or any primes that are 2 (mod 3) appear in the prime factorisation of n (i.e. $l > 0$). In the case where neither of these conditions hold, we have that either $n = \prod_{i=1}^k p_i^{\alpha_i}$ or $n = 3\prod_{i=1}^k p_i^{\alpha_i}$, where the p_i's and α_i's are as in the statement of the lemma. We have exactly 2 solutions mod $p_i^{\alpha_i}$ for $1 \le i \le k$, and so by the Chinese Remainder Theorem, we have a total of 2^k solutions mod $\prod_{i=1}^k p_i^{\alpha_i}$. Since we have exactly 1 solution mod 3, the factor of 3 does not affect the number of solutions. So, regardless of whether 3 is a factor of n, we have that the number of solutions to our congruence mod n is 2^k. □

Proof. For each of the cases $n = 1$ and $n = 3$, we have only 1 valid permutation of size n. In each case, it is easy to check that this valid permutation is a constant difference cycle permutation with difference $n - 1$, so we have exactly 1 constant difference cycle permutation. Rotations, reflections and colour swaps all map valid permutations of length n to valid permutations of length n, so in these cases where there is only one valid permutation of length n, each operation must map that valid permutation to itself. Thus, the permutation has every possible symmetry.

We now fix $n = 2g + 1 > 3$. Any constant difference cycle permutation has $2n$ symmetries from rotation and reflection, and has additional symmetries from colour swaps if and only if its constant difference d is a solution to at least one of the four congruences listed in Proposition 6.4. We examine solutions to the first three of these congruences by considering the three constant difference cycle permutations with differences 2, $\frac{n+1}{2}$ and $n-1$ respectively. We use the fact that n is odd, and hence that any prime divisor p of n must be greater than 2, to show that these three permutations are always valid.

- By Proposition 6.1, 2 is always a valid difference since $2 \not\equiv 0, 1 \pmod{p}$ for any prime $p > 2$.
- As noted in the proof of Proposition 6.4, 2 must have an inverse mod n since n is odd, and this inverse can be expressed as $2^{-1} \equiv \frac{n+1}{2} \pmod{n}$. Suppose $2^{-1} \equiv 0, 1 \pmod{p}$ for some prime divisor p of n; then $1 \equiv 0, 2 \pmod{p}$, which cannot be true since $p > 1$. Thus, $\frac{n+1}{2}$ is a valid difference by Proposition 6.1.
- Suppose $n - 1 \equiv 0, 1 \pmod{p}$ for some prime divisor p of n; then $-1 \equiv 0, 1 \pmod{p}$ and so $0 \equiv 1, 2 \pmod{p}$, which cannot be true since $p > 2$. Thus, $n-1$ is a valid difference by Proposition 6.1.

So, 2, $\frac{n+1}{2}$ and $n-1$ are always valid differences. Note that since $n > 3$, these three numbers are distinct mod n.

Since 2, $\frac{n+1}{2}$ and $n-1$ are valid differences, we can apply Proposition 6.4 to see that the corresponding constant difference cycle permutations have symmetry under the (0 2), (0 1) and (1 2) colour swaps respectively. As the three numbers are distinct mod n, each of them are solutions to only one of the 3 congruences $d \equiv 2,\ d \equiv 2^{-1},\ d \equiv -1 \pmod{n}$. Additionally, it is easy to check that since $n > 3$, the three numbers are not solutions to $d^2 - d + 1 \equiv 0 \pmod{n}$. This gives each of these permutations exactly 2 colour swap symmetries, giving a total of $4n$ symmetries when combined with the $2n$ symmetries from rotation and reflection.

It remains to be shown that the three permutations form a single isomorphism class of size 3. This fact follows from Proposition 6.3, since by choosing one of the 3 values we can reach the others via compositions of the $d \mapsto 1 - d$ and $d \mapsto d^{-1}$ operations. For example, starting from $d = 2$, we can perform a (0 1) colour swap to reach $1 - d \equiv 1 - 2 \equiv n - 1 \pmod{n}$, and we can perform a (1 2) colour swap to reach $d^{-1} \equiv 2^{-1} \equiv \frac{n+1}{2} \pmod{n}$. This shows that the three permutations must be in the same isomorphism class. Since each of the three permutations has $4n$ symmetries, we have covered all of the $3 \cdot 4n = 12n$ possible operations by which permutations can be isomorphic, and so there can be no other elements in the isomorphism class.

Having covered all possible solutions to the 3 congruences $d \equiv 2,\ d \equiv 2^{-1},\ d \equiv -1 \pmod{n}$, we now only have the congruence $d^2 - d + 1 \equiv 0 \pmod{n}$ left to consider. Using the factorisation of n specified in the statement of the theorem, we have by Lemma A.2 that if $l > 0$ or $\gamma > 1$, the congruence has no solutions. In this case, there are no constant difference cycle permutations with (0 1 2) or (0 2 1) colour swap symmetry, and so all valid constant difference cycle permutations apart from the 3 specified above have no non-trivial colour swap symmetries. This means they have exactly $2n$ symmetries and are thus in isomorphism classes of size 6 (each permutation has $2n$ symmetries, so exactly 6 must group together in a class to cover all $6 \cdot 2n = 12n$ operations).

Again by Lemma A.2, if $l = 0$ and $\gamma \leq 1$, we have 2^k solutions to $d^2 - d + 1 \equiv 0 \pmod n$. By Proposition 6.4, each of these solutions satisfies the conditions required for both (0 1 2) and (0 2 1) colour swap symmetry. As we have shown that $2, \frac{n+1}{2}, n-1$ are not solutions to $d^2 - d + 1 \equiv 0 \pmod n$ for $n > 3$, we have that each of the 2^k solutions to $d^2 - d + 1 \equiv 0 \pmod n$ have exactly 3 colour swap symmetries, giving a total of $6n$ symmetries for each solution.

Given a solution d' to $d^2 - d + 1 \equiv 0 \pmod n$, we can consider the image of the corresponding permutation under one of the non-symmetric colour swaps. We choose the colour swap (1 2); in this case the image has constant difference $(d')^{-1}$. This number must also be a solution to $d^2 - d + 1 \equiv 0 \pmod n$, since $(d')^2 - d' + 1 \equiv 0 \pmod n$ implies $1 - (d')^{-1} + ((d')^{-1})^2 \equiv 0 \pmod n$. Thus, the permutations corresponding to the differences d' and $(d')^{-1}$, which are distinct but isomorphic, each have exactly $6n$ symmetries. This means that these 2 permutations must form an isomorphism class of size exactly 2, since together they cover all $2 \cdot 6n = 12n$ operations.

The above argument applies for any of the 2^k solutions d'. Thus, we group each of the 2^k corresponding permutations into isomorphism classes of size 2, resulting in 2^{k-1} isomorphism classes of this size. And finally, as in the $l > 0$ or $\gamma > 1$ case, all constant difference cycle permutations with differences d that do not satisfy any of the four congruences from Proposition 6.4 do not have any non-trivial colour swap symmetries, and so have exactly $2n$ symmetries and are in isomorphism classes of size 6. □