Graphs of Algebraic Objects

Matthew Ashford

Linacre College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity Term 2016

This thesis is dedicated to
my Nan and Nana
You didn't quite live to see its completion
but I hope you would have been proud

Acknowledgements

I would like to thank my supervisor, Professor Oliver Riordan, for his support and advice throughout. In particular I would like to thank him for picking two very interesting topics to work on, for making countless insightful suggestions and for persevering with me through some trying times.

I would like to thank my college advisor Alison Reid for her care and support, and for being an encouraging voice when I most needed one.

I would also like to thank the Engineering and Physical Sciences Research Council for the invaluable financial support that they have given me over the last few years.

I must also thank my Mum, my Dad, and Nick for their incredible levels of encouragement and support that they have lent me, giving me the strength to keep going during some dark times. I also thank my friends and housemates (Paul, Sam, James and Adam) for their good humour and great understanding for the last three years.

And finally, I thank Shannon, for being the truly amazing person that she is.

Abstract

A rack on $[n]$ can be thought of as a set of permutations $(f_x)_{x\in[n]}$ such that $f_{(x)f_y} = f_y^{-1} f_x f_y$ for all x and y; additionally, a rack such that $(x)f_x = x$ for all x is a quandle, and a quandle such that $f_x^2 = \iota$ for all x is a kei. These objects can be represented by loopless, edge-coloured directed multigraphs on $[n]$; we put an edge of colour y between x and z if and only if $(x)f_y = z$. This thesis studies racks and kei via these graphical representations.

In 2013, Blackburn [2] showed that for any $\epsilon > 0$ and n sufficiently large, the number of isomorphism classes of kei on $[n]$ is at least $2^{(1/4-\epsilon)n^2}$, while the number of isomorphism classes of racks on $[n]$ is at most $2^{(c+\epsilon)n^2}$ for a constant c. The main result of this thesis is that the lower bound is asymptotically correct; for any $\epsilon > 0$ the number of isomorphism classes of racks on $[n]$ is at most $2^{(1/4+\epsilon)n^2}$, for n sufficiently large. We also show that in almost all of these racks almost all components of the corresponding graph have size 2. Additionally, we show several more detailed results regarding kei, including a full classification of all kei in which each component has size equal to an odd prime.

This thesis also contains several partial results on another topic. The commuting graph of a finite group is defined to be the graph on the non-trivial cosets of $G/Z(G)$ where $\bar{x}\bar{y}$ is an edge if and only if x and y commute. In 2012, Hegarty and Zhelezov [12] proposed a family of random groups and published some results on the corresponding family of random graphs; this thesis contains some generalisations and extensions of these results.

Statement of Originality

The work described in Sections 2.1, 3.1, 3.2, 8.1 and 8.2, the whole of Chapters 5 and 6, and the Appendix is entirely my own; the work described in the rest of the thesis was carried out jointly with my supervisor, Professor Oliver Riordan.

Contents

Introduction and Notation

This thesis consists of two distinct parts, although both parts concern families of algebraic objects that can be represented as graphs or multigraphs.

In the first part, we consider the commuting graph of a (finite) group, defined to be the graph on the non-trivial cosets of $G/Z(G)$ where $\bar{x}\bar{y}$ is an edge if and only if x and y commute. A 2008 conjecture of Iranmahesh and Jafarzadeh (see [17]) claimed that there was an upper bound on the diameters of all such connected graphs, but in 2012 Hegarty and Zhelezov [12] proposed a random group model that provided some evidence against this. The corresponding series of random graphs share several properties with the Erdős–Rényi model $\mathcal{G}(n, n^{-1+\epsilon})$, which has diameter k for $k-1 < 1/\epsilon < k$. This part further explores this proposed model with an aim towards finding its diameter, developing several preliminary results while also indicating some areas in which the new model demonstrates markedly different behaviour from the Erdős–Rényi model. The case of diameter 2 is proved, but the case of diameter 3 seems to be much harder; several problems and obstructions in this case are also demonstrated.

In the second part, we consider a less studied family of algebraic objects. A rack on $[n]$ can be thought of as a set of permutations $(f_x)_{x\in[n]}$ such that $f_{(x)f_y} = f_y^{-1} f_x f_y$ for all x and y; additionally, a rack such that $(x)f_x = x$ for all x is a quandle, and a quandle such that $f_x^2 = \iota$ for all x is a kei. These objects can be represented by loopless, edge-coloured directed multigraphs on $[n]$; we put an edge of colour y between x and z if and only if $(x)f_y = z$. The components of this graph coincide with the orbits of the operator group $G = \langle f_1, \ldots, f_n \rangle$ in its natural action on $[n]$.

In 2013, Blackburn [2] showed upper and lower bounds on the number of isomorphism classes of racks of order n. For any $\epsilon > 0$ and n sufficiently large, a lower bound of $2^{(1/4-\epsilon)n^2}$ was obtained by considering kei for which almost every orbit of G had size 2, while an upper bound of $2^{(c+\epsilon)n^2}$ (for a constant c) was obtained using more algebraic methods. The main result of this thesis is that the lower bound is asymptotically correct; for any $\epsilon > 0$ and n sufficiently large, the number of isomorphism classes of racks on $[n]$ is at most $2^{(1/4+\epsilon)n^2}$, and in almost all of these racks almost all orbits of G haze size 2. We show this result by considering the corresponding graphs of racks, and apply a variety of combinatorial methods. We also show several more detailed results regarding the (undirected) graphs representing kei, including a full classification of all kei in which each orbit has size equal to an odd prime.

The following notation will be used in this thesis. Let $f, g : \mathbb{N} \to [0, \infty)$ be sequences.

- $f = o(g)$ means $\frac{f(n)}{g(n)} \to 0$ as $n \to \infty$.
- $f = O(g)$ means there exists some constant $C > 0$ such that $f(n) \leqslant Cg(n)$ for n sufficiently large.
- $f = \omega(g)$ means $g = o(f)$.
- $f = \Omega(g)$ means $g = O(f)$.
- $f = \Theta(g)$ means $f = O(g)$ and $f = \Omega(g)$, i.e. there exist constants $c, C > 0$ such that $cg(n) \leqslant f(n) \leqslant Cg(n)$ for n sufficiently large.
- $f \sim g$ means $\frac{f(n)}{g(n)} \to 1$ as $n \to \infty$.

Let $E = E_n$ be an event (or formally a sequence of events, one for each $n = 1, 2, \ldots$) in a probability space. We say that E *occurs with high probability* (or *w.h.p.* for short) if $\mathbb{P}(E_n) \to 1$ as $n \to \infty$. If F_n also occurs with high probability, it is straightforward to see that $E_n \cap F_n$ occurs w.h.p.; this can be extended by induction to a finite number of events.

When defining multigraphs, we don't consider edges to have their own identity, i.e. a multigraph G is defined by a vertex set $V = V(G)$ and an edge multiset $E = E(G)$ of unordered pairs from V. For multisets A and B, $A \uplus B$ is the multiset obtained by including each element v with multiplicity m, where m is the sum of the multiplicity of v in A and the multiplicity of v in B. If F is a multiset of unordered pairs from $V(G)$, $G + F$ is the multigraph with vertex set $V(G)$ and edge multiset $E(G) \uplus F$.

We may also consider a directed multigraph G, where $\overrightarrow{E}(G)$ is a multiset of *ordered* pairs from $V(G)$. A path in such a directed graph G need not respect the orientation of the edges, so for $x, y \in V(G)$, there is an xy-path in G if and only if there is an xy-path in the underlying undirected graph. A component of a directed graph G is defined to be a component of the underlying undirected graph. For $x, y \in V(G)$, a *directed* xy-path is a sequence of vertices $x = x_0, \ldots, x_r = y$ such that $\overrightarrow{x_i x_{i+1}} \in \overrightarrow{E}(G)$ for all i.

When considering a function f related to an algebraic object, we will follow the algebraic convention and write the function on the *right*; this will be used for homomorphisms of groups and racks, and for groups acting on a set. For functions defined on a subset of the real line, we will continue to write the function on the *left*.

Part I

Commuting Graphs of Random Groups

Chapter 1

Introduction to Part I

1.1 Commuting graphs

In this first section, we will introduce the concepts behind the first part of this thesis and present the model being used. The majority of this part will be on a random graph model, but it has its roots and motivation in an algebraic question, so we will need a few results and ideas from basic group theory. Recall that the *commutator* of any two group elements x, y in a group G is

$$[x, y] = x^{-1}y^{-1}xy = (yx)^{-1}xy$$

so that x and y commute if and only if $[x, y] = 1$. Taking the inverse, we see that

$$[x, y]^{-1} = y^{-1}x^{-1}yx = [y, x]$$

for any x, y. The following lemma will be important throughout.

Lemma 1.1.1. *For elements x, y and z in a group G,*

$$[zx, y] = [z, y]^x[x, y] \quad \textit{and} \quad [x, zy] = [x, y][x, z]^y,$$

where in general $g^y := y^{-1}gy$ is the conjugation map.

Proof. Using the definition of commutator we see that

$$\begin{aligned} [zx, y] &= (zx)^{-1}y^{-1}zxy \\ &= x^{-1}z^{-1}y^{-1}z(yy^{-1})xy \\ &= x^{-1}[z, y](xx^{-1})y^{-1}xy \\ &= [z, y]^x[x, y]. \end{aligned}$$

By taking inverses we see that

$$[x, zy] = [zy, x]^{-1} = ([z, x]^y [y, x])^{-1} = [x, y][x, z]^y.$$

$\square$

Now recall that the *centre* of a group G is

$$Z(G) = \{x \in G \mid [x, y] = 1 \text{ for all } y \in G\},$$

i.e. the set of all elements which commute with every other element. Our starting point is the following definition.

Definition 1.1.2. The *commuting graph* $\Gamma(G)$ of a non-abelian group G has vertex set $G \setminus Z(G)$ and edge relations given by

$$xy \in E(\Gamma(G)) \text{ if and only if } [x, y] = 1.$$

As an example, consider the dihedral groups D_{2n}, describing the symmetries of a regular n-gon in the plane. These groups have presentation

$$\langle x, y \mid x^n, y^2, yxyx\rangle$$

where x represents a rotation through $2\pi/n$ and y a reflection. We can then compute $Z(D_6) = \{1\}$ and $Z(D_8) = \{1, x^2\}$; the graphs $\Gamma(D_6)$ and $\Gamma(D_8)$ are given in Figure 1.1. The following lemma shows that commutators are invariant over cosets of central

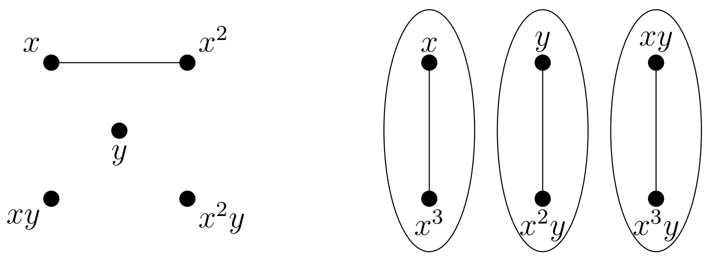

Figure 1.1: Commuting graphs of D_6 (left) and D_8 (right). The cosets of $Z(D_8)$ in D_8 are indicated; as $Z(D_6)$ is trivial, each coset of $Z(D_6)$ in D_6 is trivial.

subgroups.

Lemma 1.1.3. *Let G be a group and suppose $H \leqslant Z(G)$. Let $x, y \in G$; then for any $x' \in Hx$ and $y' \in Hy$, $[x', y'] = [x, y]$.*

Proof. By definition, there exist $z, w \in H \leqslant Z(G)$ such that $x' = zx$ and $y' = wy$. Then from Lemma 1.1.1,

$$[x', y'] = [zx, wy] = [z, wy]^x[x, wy] = [x, y][x, w]^y = [x, y],$$

giving the result. □

Note that any central subgroup of G is necessarily a normal subgroup of G, and so the quotient group G/H is always well-defined.

If we apply this lemma with $H = Z(G) =: Z$, we see that commutators are invariant over cosets of Z; thus if $[x, y] = 1$ then the induced subgraph $\Gamma(G)[Zx \cup Zy]$ is a complete bipartite graph, while if $[x, y] \neq 1$ then $\Gamma(G)[Zx \cup Zy]$ is empty. As $[x, x] = 1$ for all x, each $\Gamma(G)[Zx]$ is a clique. This allows us to redefine the commuting graph by taking $(G/Z) \setminus \{Z\}$ as the vertex set with the edge set E given by

$$ZxZy \in E \text{ if and only if } [x, y] = 1.$$

The revised commuting graphs of D_6 and D_8 under this construction are given in Figure 1.2. Informally, the second notion of commuting graph is obtained from the

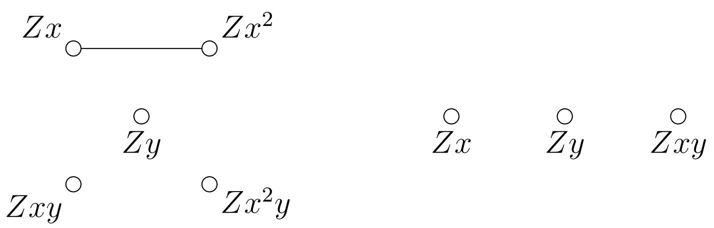

Figure 1.2: Commuting graphs of D_6 (left) and D_8 (right)

first by compressing certain cliques into single vertices and replacing certain complete bipartite graphs with single edges; it follows that the two graphs are either both disconnected or both connected, and in the latter case they have the same diameter (unless the original graph consists of just one coset of Z, in which case the diameter changes from 1 to 0). Because (in general) the second graph is smaller, we will adopt this second definition of $\Gamma(G)$.

The concept of commuting graphs was introduced by Brauer and Fowler in a 1955 paper on groups of even order [3]; they were later used in the classification of finite simple groups, as well as being studied in their own right (see [27]). Since 2000, attention has been paid most prominently to the diameters of such graphs, with various interesting results being shown for specific, classical classes of finite groups (see for example [29]). In 2008, Iranmanesh and Jafarzadeh [17] made the following strong conjecture:

Conjecture 1.1.4. *There exists an absolute constant c such that if G is a finite, non-abelian group with $\Gamma(G)$ connected, then* $\operatorname{diam}(\Gamma(G)) \leqslant c$.

The authors provided some evidence in the form of the symmetric groups $\mathfrak{S}_n$ and a collection of partial results were proved over the next few years [10]. In 2012, however, Hegarty and Zhelezov [12] produced evidence against the conjecture; they constructed a sequence of random groups whose commuting graphs seem to have increasing diameter. This led them to formally conjecture the *negation* of Conjecture 1.1.4:

Conjecture 1.1.5. *There exists a sequence of groups $(G_k)_{k=1}^{\infty}$ with connected commuting graphs such that* $\operatorname{diam}(\Gamma(G_k)) \geqslant k$ *for each k.*

Hegarty and Zhelezov made progress towards this conjecture by considering the random commuting graphs resulting from the random groups that they constructed; these graphs bear a certain similarity to the classical Erdős–Rényi binomial random graph. This part of the thesis continues the study of these graphs and we are able to show several partial results, but not the intended result. To our knowledge, no results have been published concerning this model outside of [12].

However, Conjecture 1.1.5 was confirmed in all cases in 2013 by Giudici and Parker [9], using entirely deterministic methods.

1.2 A random group model

Before we can describe the random graph model, let us first construct the random group; in time we will be able to drop the group theory aspects and reduce to a different algebraic framework. The construction is that of [12] presented in a slightly different way.

Fix an abelian group H and a natural number m (we will let $m \to \infty$ later). First suppose that G is a finite group such that $H \leqslant Z(G)$ (and so $H \lhd G$) and $G/H = \langle Hx_1, \ldots, Hx_m \rangle \cong C_2^m$, where $x_1, \ldots, x_m \in G$. As G/H is a 2-group, $H = (Hx_i)^2 = Hx_i^2$ for all i, and as G/H is abelian, $Hx_ix_j = Hx_jx_i$ for all i, j; these conditions are equivalent to

$$x_i^2 \in H \text{ for each } i = 1, \ldots, m. \tag{1.2.1}$$

and

$$[x_i, x_j] \in H \text{ for each } 1 \leqslant i < j \leqslant m. \tag{1.2.2}$$

Now consider the reverse problem of *constructing* a group G given H and m; we require $H \leqslant Z(G)$ and $G/H \cong C_2^m$, where G/H is *defined* to be $\langle Hx_1, \dots, Hx_m \rangle$ for distinct elements $x_1, \dots, x_m \in G$. The statements (1.2.1) and (1.2.2) give the correct conditions for $G/H \cong C_2^m$. Any such group G will be generated by H and the set $\{x_1, \dots, x_m\}$; we will write G as an appropriate subgroup of the free group on a set of generators. Suppose that H has a presentation $\langle y_1, \dots, y_l \mid R' \rangle$ and let $h_i, h_{ij} \in H$ for $1 \leqslant i, j \leqslant m$; we consider a set of relators R consisting exactly of the following elements:

1. The set R';
2. $x_i^{-1} y_k^{-1} x_i y_k = [x_i, y_k] = 1$ for $1 \leqslant i \leqslant m$, $1 \leqslant k \leqslant l$;
3. $x_i^2 h_i^{-1}$ for $1 \leqslant i \leqslant m$;
4. $x_i^{-1} x_j^{-1} x_i x_j h_{ij}^{-1}$ for $1 \leqslant i < j \leqslant m$.

Item 1 ensures that H is reconstructed correctly, while item 2 ensures that H is central. Items 3 and 4 ensure that (1.2.1) and (1.2.2) hold, and thus that $G/H \cong C_2^m$. Hence $G = \langle x_1, \dots, x_m, y_1, \dots, y_l \mid R \rangle$ has the desired properties, regardless of the choices of the elements h_i and h_{ij} (for $i < j$).

As an example, suppose $H = \langle y \mid y^2 \rangle \cong C_2$ and $m = 2$; then

$$G = \langle x_1, x_2, y \mid y^2, [x_1, y], [x_2, y], x_1^2 h_1^{-1}, x_2^2 h_2^{-1}, [x_1, x_2] h_{12}^{-1} \rangle,$$

where each of h_1, h_2 and h_{12} is equal to 1 or y. If we choose $h_{12} = 1$ then G is abelian as all of the generators commute, and a simple calculation shows that $G \cong C_2^3$ if we take $h_1 = h_2 = 1$ and $G \cong C_4 \times C_2$ otherwise. If we take $h_{12} = y$ then G is necessarily non-abelian; a slightly more involved calculation shows that $G \cong Q_8$ (the quaternion group) if $h_1 = h_2 = y$ and $G \cong D_8$ otherwise.

In the model we will use, H is itself an elementary abelian 2-group, namely $C_2^{m+r} = \langle w_1, \dots, w_r, h_1, \dots, h_m \rangle$ where r is a natural number depending on m. Deterministically, we will take $x_i^2 = h_i$ for each i; the h_is are present more for convenience and could practically be dropped from the model. The randomness in the model comes from the choices of commutators as in (1.2.2); for for each pair $i, j \in [m]$ with $i < j$, the commutator $[x_i, x_j]$ is chosen uniformly at random from the 2^r elements of $\langle w_1, \dots, w_r \rangle \leqslant H$, independently for all pairs.

There are many facts that can be shown about this random group (see [12]), but the most important for us will be the following:

Proposition 1.2.1. *Suppose that $r = r(m) \geqslant 2$ for sufficiently large m. Then $H = Z(G)$ with high probability (w.h.p.) as $m \to \infty$.*

(This is Proposition 3.3 of [12]). The proof of this statement is deferred until we have developed a slightly different framework for the problem. In the interim we will consider the graph Γ on the vertex set $G/H \setminus \{H\}$ with edge relations given by

$$HxHy \in E(\Gamma) \text{ if and only if } [x, y] = 1.$$

As $H \leqslant Z(G)$, we have from Lemma 1.1.3 that the commutators are invariant on cosets of H, and thus Γ is well-defined.

Now let us consider the properties of the random graph Γ, detached from the inherent group theory. For each l we will identify in the obvious way the elementary abelian 2-groups C_2^l with the vector space $\mathbb{F}_2^l$; for this reason it is more natural to use additive notation. In this light, we regard the group G/H as an m-dimensional space V with basis $\{H + x_1, \ldots, H + x_m\}$ and $W := \langle w_1, \ldots, w_r\rangle \leqslant H$ as an r-dimensional space. There is a natural, random map $B : V \times V \to W$ given by

$$B : (H + x, H + x') \to [x, x']$$

and this map is well-defined from Lemma 1.1.3. The edge $\{H + x, H + x'\}$ appearing is therefore equivalent to $B(H + x, H + x') = 0$, giving us our edge distribution. We can immediately establish some key properties of this function.

Proposition 1.2.2. *The map $B : V \times V \to W$ described above satisfies, for all $v, v', v'' \in V$:*

1. $B(v, v) = 0$;
2. $B(v, v' + v'') = B(v, v') + B(v, v'')$;
3. $B(v' + v'', v) = B(v', v) + B(v'', v)$;
4. $B(v, v') = B(v', v)$.

Proof. For the first statement, write $v = H + x$, so $B(v, v) = [x, x] = 0$ by definition. For the second statement write $v = H + x$, $v' = H + x'$ and $v'' = H + x''$ for some x, x', x''. Then by Lemma 1.1.1,

$$B(v, v' + v'') = [x, x' + x''] = [x, x''] + [x, x']^{x''} = [x, x'] + [x, x''] = B(v, v') + B(v, v''),$$

as (in the group theory sense) the commutators are chosen to belong to a central subgroup. The third statement follows similarly. For the fourth, simply note

$$0 = B(v+v', v+v') = B(v,v) + B(v,v') + B(v',v) + B(v',v') = B(v,v') + B(v',v),$$

which gives the result as we are working over $\mathbb{F}_2$. $\square$

We thus see that B is a symmetric, bilinear map that vanishes when its two arguments are equal, and thus it is entirely determined by its action on the basis elements $(H + x_i, H + x_j)$ for $i < j$. But note that $B(H + x_i, H + x_j) = [x_i, x_j]$ is chosen uniformly at random from W, independently for all pairs (i, j). Γ is thus equivalent to the model in the following definition.

Definition 1.2.3. Let $V = \langle v_1, \ldots, v_m \rangle$ and W be vector spaces over $\mathbb{F}_2$ of dimensions m and r respectively, the former with a given basis. Construct a random bilinear map $[\cdot, \cdot]$ from $V \times V$ to W by choosing each $[v_i, v_j]$ $(i < j)$ independently, uniformly at random from W, and then extending to the whole of $V \times V$ subject to symmetry and the condition $[v_i, v_i] = 0$ for each i. The vertex set of the graph Γ is $V \setminus \{0\}$ and the edge vv' is present if and only if $[v, v'] = 0$.

This is the graph model we will work with throughout the rest of this part. Note that as things stand, the parameter r is independent of m; our hope is that by choosing r to depend on m in the correct way we can show that with high probability (as $m \to \infty$) the diameter of this graph is equal to a fixed, arbitrary integer k. This, along with the proof of Proposition 1.2.1, would verify the conjecture of Hegarty and Zhelezov. In order to prove this key Proposition, we first need some important results on random linear and bilinear maps.

1.3 Random linear and bilinear maps

1.3.1 Linear maps

We consider the following situation. Let V and W be vector spaces over $\mathbb{F}_2$, of dimensions m and r respectively, and fix a basis for V, say $\mathcal{B} = \{v_1, \ldots, v_m\}$. Define a random linear map $F : V \to W$ by choosing $F(v_i)$ uniformly at random from the 2^r elements of W, independently for each $i = 1, \ldots, m$. By defining a map on a basis we have defined a (unique) linear map on the whole space. We will show that we can choose *any* basis for V and retain the important property that the image of each basis vector is chosen uniformly at random from W, independently of all other basis vectors; we will need the following lemma.

Lemma 1.3.1. *Let $X_1, \dots, X_n$ be independent random variables, where each X_i is uniform on W. Then $X_1 + X_2$ is also uniform on W and the random variables $X_1 + X_2, X_2, \dots, X_n$ are independent, i.e. $(X_1 + X_2, X_2, \dots, X_n) \sim (X_1, X_2, \dots, X_n)$.*

Proof. For any $w \in W$, we have (as X_1 and X_2 are independent)

$$\begin{aligned}
\mathbb{P}(X_1 + X_2 = w) &= \sum_{u \in W} \mathbb{P}(X_1 + X_2 = w | X_1 = u)\mathbb{P}(X_1 = u) \\
&= 2^{-r} \sum_{u \in W} \mathbb{P}(X_2 = w - u | X_1 = u) \\
&= 2^{-r} \sum_{u \in W} \mathbb{P}(X_2 = w - u) \\
&= 2^{-r}.
\end{aligned}$$

As w was arbitrary we have that $X_1 + X_2$ is uniform on W.

Now let $w_1, \dots, w_m \in W$. Note that

$$\begin{aligned}
\mathbb{P}(X_1 + X_2 = w_1, X_2 = w_2) &= \mathbb{P}(X_1 = w_1 - w_2, X_2 = w_2) \\
&= \mathbb{P}(X_1 = w_1 - w_2)\mathbb{P}(X_2 = w_2) \\
&= 2^{-2r},
\end{aligned}$$

where we have used the independence of X_1 and X_2. Let A be the event $\{X_1 + X_2 = w_1, X_i = w_i \text{ for } i \geqslant 2\}$; then by the independence of $X_1, \dots, X_n$,

$$\begin{aligned}
\mathbb{P}(A) &= \mathbb{P}(X_1 + X_2 = w_1, X_2 = w_2) \cdot \prod_{i=3}^{n} \mathbb{P}(X_i = w_i) \\
&= 2^{-2r} \cdot (2^{-r})^{m-2} \\
&= 2^{-mr} \\
&= \mathbb{P}(X_1 + X_2 = w_1) \cdot \mathbb{P}(X_2 = w_2) \cdot \prod_{i=3}^{n} \mathbb{P}(X_i = w_i).
\end{aligned}$$

□

We can now show the result.

Proposition 1.3.2. *Let $\{x_1, \dots, x_l\}$ be a set of linearly independent vectors in V. Then the random variables $F(x_1), \dots, F(x_l)$ are independent, and $F(x_i)$ is uniform on W for each i.*

Proof. Let $\mathcal{B} = \{v_1, \ldots, v_n\}$ be the original basis of V. Given a linearly independent set $\{x_1, \ldots, x_l\}$ (with $l \leqslant m$ necessarily) we can extend this to a basis $\mathcal{B}' = \{x_1, \ldots, x_l, x_{l+1}, \ldots, x_m\}$. It is a basic fact from linear algebra (see for example Chapter 2 of [4]) that we can obtain $\mathcal{B}'$ from the original ordered basis $\mathcal{B}$ via a sequence of bases $\mathcal{B} = \mathcal{B}_0, \mathcal{B}_1, \ldots, \mathcal{B}_t = \mathcal{B}'$ where each $\mathcal{B}_{k+1} = \{v_1^{(k+1)}, \ldots, v_m^{(k+1)}\}$ is obtained from $\mathcal{B}_k = \{v_1^{(k)}, \ldots, v_m^{(k)}\}$ by exactly one of the following three operations:

1. Swap two elements of $\mathcal{B}_r$, i.e. put $v_i^{(k+1)} = v_j^{(k)}$ and $v_j^{(k+1)} = v_i^{(k)}$ for some $i \neq j$;

2. Replace an element by a non-zero scalar multiple, i.e. put $v_i^{(k+1)} = \lambda v_i^{(k)}$ for some i and $\lambda \neq 0$;

3. Replace an element by a scalar multiple of another, i.e. put $v_i^{(k+1)} = v_i^{(k)} + \lambda v_j^{(k)}$ for some $i \neq j$ and $\lambda \neq 0$.

Note that item 1 only corresponds to a re-ordering; for this reason we may assume without loss of generality that $i = 1$ and $j = 2$ in items 2 and 3. Further, the only non-zero constant in $\mathbb{F}_2$ is 1, so item 2 is not a change and item 3 corresponds to adding the first two elements; thus $\mathcal{B}_{k+1} = \{v_1^{(k)} + v_2^{(k)}, v_2^{(k)}, \ldots, v_m^{(k)}\}$.

Now in $\mathcal{B}_0$, the random variables $(F(v_1), \ldots, F(v_m)) = \left(\left((F(v_1^{(0)}\right), \ldots, F\left(v_m^{(0)}\right)\right)$ are independent and uniformly distributed on W by construction; take $k \geqslant 0$ and suppose that the random variables $F\left(v_1^{(k)}\right), \ldots, F\left(v_m^{(k)}\right)$ are also independent and uniformly distributed on W. Write $X_i = F\left(v_i^{(k)}\right)$ for each i; then $F\left(v_1^{(k+1)}\right) = F\left(v_1^{(k)} + v_2^{(k)}\right) = F\left(v_1^{(k)}\right) + F\left(v_2^{(k)}\right) = X_1 + X_2$ by the linearity of F, while $F\left(v_i^{(k+1)}\right) = F\left(v_i^{(k)}\right) = X_i$ for $i \geqslant 2$. Hence by Lemma 1.3.1, the random variables $F\left(v_1^{(k+1)}\right), \ldots, F\left(v_m^{(k+1)}\right)$ are independent and uniform on W. We have shown inductively the the variables $F\left(v_1^{(t)}\right), \ldots, F\left(v_m^{(t)}\right)$ corresponding to the basis $\mathcal{B}_t = \mathcal{B}'$ are independent and uniformly distributed on W, proving the result as (without loss of generality) $x_i = v_i^{(t)}$ for each i. $\square$

We will use this result repeatedly throughout the first part of the thesis.

1.3.2 Bilinear maps

Let V and W be the vector spaces described in the previous subsection. Now consider the situation described in Definition 1.2.3, wherein we form a random bilinear map $[\cdot, \cdot]$ from $V \times V$ to W by choosing each of the $\binom{m}{2}$ elements $[v_i, v_j]$ $(i < j)$ uniformly and independently, and extend this to the whole of $V \times V$ under the condition that

$[v_i, v_i] = 0$ for all i. We will show that this construction, as before, is invariant under a change of basis.

Proposition 1.3.3. *Let $\{x_1, \ldots, x_l\}$ be a be a set of linearly independent vectors in V. Then the list of variables $([x_i, x_j])_{i<j}$ is independent, and each $[x_i, x_j]$ is uniformly distributed on W.*

Proof. Extend $\{x_1, \ldots, x_l\}$ to a basis $\mathcal{B}'$ of V; as in Proposition 1.3.2, we can transform the original basis $\mathcal{B} = \{v_1, \ldots, v_m\}$ to $\mathcal{B}'$ via a sequence of bases $\mathcal{B} = \mathcal{B}_0, \mathcal{B}_1, \ldots, \mathcal{B}_t = \mathcal{B}'$, where we can order these bases such that for each k, $\mathcal{B}_{k+1}$ is obtained from $\mathcal{B}_k = \{v_1^{(k)}, v_2^{(k)}, \ldots, v_m^{(k)}\}$ by putting $v_1^{(k+1)} = v_1^{(k)} + v_2^{(k)}$ and $v_i^{(k+1)} = v_i^{(k)}$ for $i \geqslant 2$.

Now in $\mathcal{B}_0$, the random variables $([v_i, v_j])_{i<j} = \left(\left[v_i^{(0)}, v_j^{(0)}\right]\right)_{i<j}$ are independent and uniformly distributed on W by construction; take $k \geqslant 0$ and suppose that the random variables $\left(\left[v_i^{(k)}, v_j^{(k)}\right]\right)_{i<j}$ are also independent and uniformly distributed on W. Write $X_{ij} = \left(\left[v_i^{(k)}, v_j^{(k)}\right]\right)$ for $i < j$; then

$$\left(\left[v_1^{(k+1)}, v_2^{(k+1)}\right]\right) = \left(\left[v_1^{(k)} + v_2^{(k)}, v_2^{(k)}\right]\right) = \left(\left[v_1^{(k)}, v_2^{(k)}\right]\right) + \left(\left[v_2^{(k)}, v_2^{(k)}\right]\right) = X_{12},$$

from the properties of $[\cdot, \cdot]$, while

$$\left(\left[v_1^{(k+1)}, v_j^{(k+1)}\right]\right) = \left(\left[v_1^{(k)} + v_2^{(k)}, v_j^{(k)}\right]\right) = \left(\left[v_1^{(k)}, v_j^{(k)}\right]\right) + \left(\left[v_2^{(k)}, v_j^{(k)}\right]\right) = X_{1j} + X_{2j}$$

for $j \geqslant 3$ and $\left(\left[v_i^{(k+1)}, v_j^{(k+1)}\right]\right) = \left(\left[v_i^{(k)}, v_j^{(k)}\right]\right)$ for $i \geqslant 2$. Thus the random variables corresponding to $\left(\left[v_i^{(k+1)}, v_j^{(k+1)}\right]\right)_{i<j}$ are X_{12}, $(X_{1j} + X_{2j})_{j=3}^{m}$ and $(X_{ij})_{2\leqslant i<j}$; by applying Lemma 1.3.1 $m-2$ times we see that these variables are independent, each uniformly distributed on W. We have shown inductively the the variables $\left(\left[v_i^{(t)}, v_j^{(t)}\right]\right)_{i<j}$ corresponding to the basis $\mathcal{B}_t = \mathcal{B}'$ are independent and uniformly distributed on W, proving the result as (without loss of generality) $x_i = v_i^{(t)}$ for each i. $\square$

This is also a result that we will use repeatedly.

1.4 Equivalence of group and graph models

Recall the definition of the random graph Γ given in Definition 1.2.3. In the previous section, we showed that the key properties of the random bilinear map $[\cdot, \cdot]$ are preserved no matter what basis for v we choose, which makes many of the calculations easier. In particular, it makes the proof of Proposition 1.2.1 (restated below) much more straightforward.

Proposition 1.4.1. *Suppose that $r = r(m) \geqslant 2$ for sufficiently large m. Then $H = Z(G)$ with high probability (w.h.p.) as $m \to \infty$.*

Proof. Recall that the group G is generated by $H \leqslant Z(G)$ and a set $\{x_1, \ldots, x_m\}$, where $G/H = \langle Hx_1, \ldots, Hx_m \rangle \cong C_2^m$. H also contains an elementary abelian subgroup $W = \langle w_1, \ldots, w_r \rangle \cong C_2^r$; we can thus regard $V := G/H$ and W as vector spaces over $\mathbb{F}_2$ of dimensions m and r respectively[1].

As described in Section 1.2, the random map $B : V \times V \to W$ given by $(H + x, H + x') \mapsto [x, x']$ is well-defined, and constructed as in Section 1.3.2. Because of this, we can change basis freely.

Now in our construction, $H \leqslant Z(G)$ always; the containment is strict if and only if $g \in Z(G)$ for some $g \notin H$. As $V = G/H$ is elementary abelian, every non-trivial coset of H in G/H has a representation as Hx for some $x \in S$, where

$$S := \{x_{j_1} \cdots x_{j_l} \mid l \geqslant 1, 1 \leqslant j_1 < \cdots < j_l \leqslant m\}$$

is the set of non-trivial products of the x_is. Therefore we have strict containment if and only if there exists an $x \in S$ such that $[hx, h'x'] = [x, x'] = 1$ for every $x' \in S$.

Fix $x \in S$; then Hx is a non-zero element of V and thus can be extended to a basis $Hx = H\tilde{x}_1, \ldots, H\tilde{x}_m$ of V. As in Proposition 1.2.2, it suffices to consider the events $\{B(H\tilde{x}_1, H\tilde{x}_j) = 0\}_{j=2}^m = \{[\tilde{x}_1, \tilde{x}_j] = 1\}_{j=2}^m$; all of these events holding is equivalent to the event $\{x \in Z(G)\}$. By Proposition 1.3.3, these events are all independent with probability 2^{-r}, so

$$\mathbb{P}(x \in Z(G)) = 2^{-r(m-1)}$$

for each x. But there are $2^m - 1$ such products x, so a union bound gives

$$\mathbb{P}(Z(G) \cap S \neq \emptyset) \leqslant \mathbb{P}\left(\bigcup_{x \in S} \{x \in Z(G)\}\right) 2^m 2^{-r(m-1)} \to 0$$

as (for sufficiently large m) $r(m-1) - m \geqslant m - 2 \to \infty$ as $m \to \infty$. □

This shows that with high probability $\Gamma = Z(G)$, and thus questions about the commuting graph can be answered by considering the model Γ, as defined in Definition 1.2.3. As the construction of Γ is independent of any group theory, we will redefine Γ as $G_{m,r}$, where we will consider r as a function of m.

The rest of the first part of the thesis is organised into two chapters. In Chapter 2, we show some results about fixed subgraphs, in particular trees and cycles, which

[1] We will, however, continue to use multiplicative notation in this proof.

allow us to give a first moment argument and make progress with a second moment argument. In Chapter 3 we will examine the neighbourhood of a vertex with an aim towards constructing paths of length k, showing the result for the case where $1/2 < \epsilon < 1$. We also explain in this chapter some of the difficulties arising when $1/3 < \epsilon < 1/2$.

Chapter 2

Paths and cycles

In this chapter we examine paths and cycles in the random graph model $G_{m,r}$; this is an extension and generalisation of the results in [12]. We are able to show that the appearance of trees and cycles in $G_{m,r}$ is approximately analogous to their appearance in $\mathcal{G}(n,p)$, which allows us to show the first moment part of the argument and make significant progress with a second moment argument.

2.1 Appearance of fixed subgraphs

In this section, we will explore a key aspect of the random graph model $\mathcal{G}_{m,r}$: the appearance of fixed subgraphs. Let us first consider the probability of a single edge appearing in this model.

Lemma 2.1.1. *For any two vertices $u, v \in V(G_{m,r})$, $\mathbb{P}(uv \in E(G_{m,r})) = 2^{-r}$.*

Proof. u and v are distinct, non-zero vectors of V; as V is defined over $\mathbb{F}_2$ this means that $\{u, v\}$ is a linearly independent set. Now extend this set to a basis of V; using this basis we see that $[u, v]$ is uniformly distributed over W, so we have $\mathbb{P}(uv \in E(G_{m,r})) = \mathbb{P}([u, v] = 0) = 2^{-r}$. □

This shows that each edge appears with the same probability 2^{-r}, which we will denote by p. However, the major difficulty with the model $G_{m,r}$ is that the edges do *not* appear independently; once the $\binom{m}{2}$ elements $([v_i, v_j])_{i<j}$ have been chosen, the remaining edges are all determined. This dependence makes calculations of certain important properties much more complicated.

As an example, consider the probability of $G_{m,r}$ containing a subgraph H on a fixed set of vertices S. In $\mathcal{G}(n,p)$, this probability is $p^{e(H)}$, regardless of the vertex set S. It is straightforward to see that this is not the case in $G_{m,r}$ by considering the simple example of the triangle K_3.

We can take any two vertices, say x and y, and they will be linearly independent. If the third vertex is $z = x + y$, then

$$\mathbb{P}(G_{m,r}[x,y,z] \cong K_3) = \mathbb{P}\big([x,y] = [x,x+y] = [y,x+y] = 0\big) = \mathbb{P}([x,y] = 0) = p,$$

from the bilinearity of $[\cdot,\cdot]$. If on the other hand the third vertex is some $z \neq x+y$ then $\{x,y,z\}$ is linearly independent and so can be extended to a basis. Adopting this basis as in Proposition 1.3.3, the elements $[x,y]$, $[x,z]$ and $[y,z]$ are chosen independently, so

$$\mathbb{P}(G_{m,r}[x,y,z] \cong K_3) = \mathbb{P}([x,y] = 0)\mathbb{P}([x,z] = 0)\mathbb{P}([y,z] = 0) = p^3.$$

In the light of this example, we will need to work with *labelled* graphs rather than just isomorphism classes. The most straightforward case is when the set of vertices $V(H)$ is linearly independent, as then we can extend it to a basis, say $\mathcal{B}'$, of V and work with this basis. Each edge is then determined by an element $[u,v]$ for some $u, v \in V(H) \subseteq \mathcal{B}'$, and so these edges appear independently and with probability p (as in the second triangle example). We can consider the *dimension* of a fixed subgraph H in the following natural way: write $\langle H \rangle := \langle V(H) \rangle$ and define $\dim(H) := \dim(\langle H \rangle) = \dim(\langle V(H) \rangle)$.

By noting that the set $\{x_1, \ldots, x_{k+1}\}$ is linearly independent if and only if $x_{k+1} \notin \langle x_1, \ldots, x_k \rangle$, we obtain the following straightforward result, which will be used repeatedly.

Proposition 2.1.2. *Let V be any vector space and $x_1, \ldots, x_k, x_{k+1} \in V$. Suppose that $\dim(\langle x_1, \ldots, x_k \rangle) = d$. Then*

$$\dim(\langle x_1, \ldots, x_k, x_{k+1} \rangle) = \begin{cases} d & \text{if } x_{k+1} \in \langle x_1, \ldots, x_k \rangle \\ d+1 & \text{otherwise.} \end{cases}$$

□

By considering the linear spans of certain sets of vertices we can sometimes partition the edges into mutually independent sets.

Lemma 2.1.3. *Let $X \subset V(G)$ and let y be a vertex with $y \notin \langle X \rangle$. Then all edges from y to $\langle X \rangle$ are independent of all edges within $\langle X \rangle$.*

Proof. Let $\{x_1, \ldots, x_d\}$ be a basis of $\langle X \rangle$; as $y \notin \langle X \rangle$, $\{x_1, \ldots, x_d, y\}$ is linearly independent and can thus be extended to a basis $\mathcal{B}$ of V; we will work with this basis. Now take any distinct $u, v \in \langle X \rangle$; then there exist $\lambda_i, \mu_j \in \mathbb{F}_2$ such that $u = \sum_{i=1}^{d} \lambda_i x_i$ and $v = \sum_{j=1}^{d} \mu_j x_j$. Then

$$[u, v] = \left[\sum_{i=1}^{d} \lambda_i x_i, \sum_{j=1}^{d} \mu_j x_j \right] = \sum_{i=1}^{d} \sum_{j=1}^{d} \lambda_i \mu_j [x_i, x_j],$$

and so the event $\{uv \in E(G_{m,r})\} = \{[u, v] = 0\}$ depends only on the values of $([x_i, x_j])_{i,j=1}^{d}$. Now take $w \in \langle X \rangle$, so there exist $\rho_i \in \mathbb{F}_2$ such that $w = \sum_{i=1}^{d} \rho_i x_i$. Then

$$[y, w] = \left[y, \sum_{i=1}^{d} \rho_i x_i \right] = \sum_{i=1}^{d} \rho_i [y, x_i],$$

and so the event $\{yw \in E(G_{m,r})\} = \{[y, w] = 0\}$ depends only on the values of $([y, x_i])_{i=1}^{d}$. This set is chosen independently of $([x_i, x_j])_{i,j=1}^{d}$, and as u, v and w were arbitrary we get the result. $\square$

We can now give a bound on the probabilities of certain types of graphs appearing in the model $G_{m,r}$; these values depend crucially on the dimension of the underlying vertex set. First consider trees, which include the paths necessary for the first moment calculation in the next section (this is a generalisation of Proposition 3.5(i) in [12], which only considers paths).

Proposition 2.1.4. *Let T be a possible subtree of $G_{m,r}$ of dimension d. Then $\mathbb{P}(T \subset G_{m,r}) \leqslant p^{d-1}$.*

Note that this result holds with equality if $V(T)$ is linearly independent, as then $d = v(T)$ and $\mathbb{P}(T \subset G_{m,r}) = p^{e(T)} = p^{d-1}$ from the discussion above.

Proof. We prove this by induction on t, the number of vertices of T. If $t = 1$ then we must have $d = 1$ and the result follows trivially. So take $t > 1$ and suppose the result is true for trees with fewer vertices.

Now let u be a leaf of T and u' its neighbour, and write $T' = T \setminus u$. From Proposition 2.1.2 it follows that T' is of dimension d if $u \in \langle T' \rangle$ and $d - 1$ otherwise. In the first case we have by induction that

$$\mathbb{P}(T \subset G_{m,r}) \leqslant \mathbb{P}(T' \subset G_{m,r}) \leqslant p^{d-1}$$

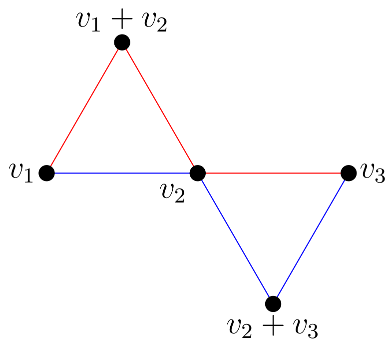

Figure 2.1: Two different paths between v_1 and v_3 of length 3 are shown, T in red and T' in blue. The dimension of $T \cup T'$ is 3, but $\mathbb{P}(T \cup T' \subset G_{m,r}) = \mathbb{P}([v_1, v_2] = [v_2, v_3] = 0) = p^2$.

and we are finished. In the second case $\dim(T') = d - 1$ and $u \notin \langle T' \rangle$, so we can apply Lemma 2.1.3 and the inductive hypothesis to see that

$$\begin{aligned}
\mathbb{P}(T \subset G_{m,r}) &= \mathbb{P}(uu' \in E(G_{m,r}),\ T' \subset G_{m,r}) \\
&= \mathbb{P}(uu' \in E(G_{m,r}))\mathbb{P}(T' \subset G_{m,r}) \\
&\leqslant p \cdot p^{d-2} \\
&= p^{d-1}.
\end{aligned}$$

The result follows by induction. □

Now let us consider some more complicated subgraphs. Proposition 3.5(iv) of [12] states (using our notation) that if T and T' are two different paths of the same length between two vertices a and b, then $\mathbb{P}(T \cup T' \subset G_{m,r}) \leqslant p^{\dim(T \cup T')}$. This statement is actually *false*; a counterexample is given in Figure 2.1. The reason for the discrepancy is that the two paths form two degenerate cycles, as mentioned earlier; by adding more such triangles, we can extend the example in Figure 2.1 to longer paths. A more careful argument than Proposition 2.1.4 shows that for all cycles C of dimension at least 3, $\mathbb{P}(C \subset G_{m,r}) \leqslant p^{\dim(C)}$; we show in Section 2.3 how this relates to counting pairs of paths.

Proposition 2.1.5. *Let C be a possible cycle in G of dimension $d \geqslant 3$. Then $\mathbb{P}(C \subset G_{m,r}) \leqslant p^d$.*

As before, note that this result holds with equality if $V(C)$ is linearly independent.

Proof. Write $t = |V(C)|$ and order $V(C)$ as $u_1, \ldots, u_t$ so that the edges of C are given by $u_1u_2, \ldots, u_{t-1}u_t, u_tu_1$. Consider the sets of vertices obtained by starting with $V(C)$

and iteratively removing $u_t, u_{t-1}, \ldots$; from Proposition 2.1.2 there exists a t' such that the segment $u_1, \ldots u_{t'}$ will have dimension $d-1$. Now choose the *longest* segment of C of dimension $d-1$ and re-order the vertices such that $u_2, \ldots, u_{s+1}$ is this segment S (of length s). Note that as $s \geqslant d-1 \geqslant 2$, the vertices u_2 and u_{s+1} are distinct. There are now two cases to consider.

First suppose that $s = t-1$, i.e. that we can remove a single vertex u_1 from C to get a set of dimension $d-1$. Then $u_1 \notin \langle S \rangle$ so all edges from u_1 to $\langle S \rangle$ are independent of those within $\langle S \rangle$ by Lemma 2.1.3. Note also that $\{u_1, u_2, u_t\}$ is linearly independent so the edges u_1u_2 and u_1u_t appear independently. Let T be the path spanned by S and note that it is of dimension $d-1$. Then

$$\begin{aligned}
\mathbb{P}(C \subset G_{m,r}) &= \mathbb{P}(u_1u_2 \in E(G_{m,r}),\ u_tu_1 \in E(G_{m,r}),\ T \subset G_{m,r}) \\
&= \mathbb{P}(u_1u_2 \in E(G_{m,r}))\mathbb{P}(u_tu_1 \in E(G_{m,r}))\mathbb{P}(T \subset G_{m,r}) \\
&\leqslant p \cdot p \cdot p^{d-2} \\
&= p^d,
\end{aligned}$$

where we have used the result for trees. Otherwise $s < t-1$, so there are distinct

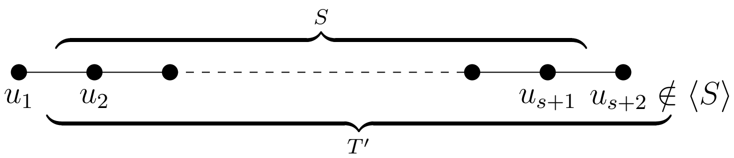

Figure 2.2: A segment of the cycle C; S is a longest segment of dimension $d-1$ while we apply the argument for trees to the d-dimensional segment T'.

vertices u_1 and u_{s+2} not contained in S. Now neither $S \cup \{u_1\}$ or $S \cup \{u_{s+2}\}$ can have dimension $d-1$ or there would be a longer segment of dimension $d-1$, so we must have $u_1, u_{s+2} \notin \langle S \rangle$. Hence the path T' on vertices $u_2, \ldots, u_{s+2}$ has dimension d (see Figure 2.2). Now extend the pair $\{u_2, u_{s+1}\}$ to a basis $\mathcal{B}_S$ of $\langle S \rangle$; from the above observation $\mathcal{B}_S \cup \{u_1\}$ is of dimension d and thus a basis of $\langle C \rangle$; we will work with this basis. It follows that we can write

$$u_{s+2} = u_1 + \sum_{x \in \mathcal{B}_S} \lambda_x \cdot x \tag{2.1.1}$$

for some $\lambda_x \in \mathbb{F}_2$ (note that the coefficient of u_1 is 1 as $u_{s+2} \notin \langle S \rangle = \langle \mathcal{B}_S \rangle$). Then we have

$$[u_{s+1}, u_{s+2}] = [u_{s+1}, u_1] + \sum_{x \in \mathcal{B}_S} \lambda_x \cdot [u_{s+1}, x],$$

so the event $\{u_{s+1}u_{s+2} \in E(G)\}$ depends on some of the independently chosen elements $([v,w])_{v,w\in\mathcal{B}_S}$ and the additional element $[u_1, u_{s+1}]$. Every other edge of T' has both endpoints in S, and so is determined by the set $([v,w])_{v,w\in\mathcal{B}_S}$. This means the event $\{T' \subset G\}$ depends only on the set

$$A = \{[v,w] \mid v,w \in \mathcal{B}_S\} \cup \{[u_1, u_{s+1}]\}.$$

Now note that as $u_1 \notin \langle S\rangle$ and $u_2 \neq u_{s+1}$ we have $[u_1,u_2] \notin A$. But then

$$\mathbb{P}(u_1u_2 \in E(G_{m,r}) \mid T' \subset G_{m,r}) = \mathbb{P}(u_1u_2 \in E(G_{m,r})) = p,$$

and hence

$$\begin{aligned}
\mathbb{P}(C \subset G_{m,r}) &\leqslant \mathbb{P}(u_1u_2 \in E(G_{m,r}),\ T' \subset G_{m,r}) \\
&= \mathbb{P}(u_1u_2 \in E(G_{m,r}) \mid T' \subset G_{m,r})\mathbb{P}(T' \subset G_{m,r}) \\
&\leqslant p \cdot p^{d-1} \\
&= p^d.
\end{aligned}$$

This proves the result. $\square$

The only cycles that Proposition 2.1.5 does not apply to are those of dimension 2. But such a cycle can contain at most $2^2-1=3$ vertices (as 0 is not a vertex), showing that the case of a 2-dimensional triangle is the only counterexample. Avoiding this case, we can extend the above result to unicyclic connected graphs. As before, the result holds with equality if the underlying vertex set is linearly independent.

Proposition 2.1.6. *Let H be a possible connected unicyclic subgraph of $G_{m,r}$ of dimension d, where the dimension of the sole cycle in H is at least 3. Then $\mathbb{P}(H \subset G_{m,r}) \leqslant p^d$.*

Proof. Note that H consists of a cycle $C = u_1,\dots,u_c$ of length c and a set of pairwise vertex-disjoint trees $T_1,\dots,T_c$ such that $T_i \cap C = \{u_i\}$ for each i. Let $T_i' = T_i \setminus u_i$ for each i; we will prove the result by induction on $t = |T_1'| + \cdots + |T_c'|$, the number of vertices outside of the cycle. If $t=0$ then H is a cycle of dimension $d \geqslant 3$, so from Proposition 2.1.5 we have $\mathbb{P}(H \subset G_{m,r}) \leqslant p^d$.

So take $t>0$ and suppose the result is true for smaller t. In this case we can find a tree T_i with a leaf distinct from u_i, and by relabelling we can assume this tree is T_1. Let u be the leaf in question and u' its neighbour, so that $H' = H \setminus u$ is a connected unicyclic graph with fewer vertices (but the same cycle), and hence we can apply the inductive hypothesis. From Proposition 2.1.2 we have that H' is of dimension d if $u \in \langle H' \rangle$ and $d-1$ otherwise. In the first case,

$$\mathbb{P}(H \subset G_{m,r}) \leqslant \mathbb{P}(H' \subset G_{m,r}) \leqslant p^d$$

and we are finished. Otherwise H' has dimension $d-1$ and $u \notin \langle H' \rangle$. As in the argument for trees we use the inductive hypothesis and Lemma 2.1.3 to obtain

$$\begin{aligned}
\mathbb{P}(H \subset G_{m,r}) &= \mathbb{P}(uu' \in E(G_{m,r}),\ H' \subset G_{m,r}) \\
&= \mathbb{P}(uu' \in E(G_{m,r}))\mathbb{P}(H' \subset G_{m,r}) \\
&\leqslant p \cdot p^{d-1} \\
&= p^d.
\end{aligned}$$

The result follows by induction. □

It would seem to difficult to extend these results to arbitrary graphs because of the more complicated structures involved, but these results will go some way towards solving the original problem, which we return to in the next section.

2.2 First Moment Method

2.2.1 Choosing parameters

Recall our original problem; for a fixed positive integer k, we wish to find a function $r = r(m,k)$ such that $\mathbb{P}(\mathrm{diam}(G_{m,r}) = k) \to 1$ as $m \to \infty$. In the previous section, we showed that in the specific context of fixed subgraphs $G_{m,r}$ behaves in a similar way to the classic Erdős–Rényi model $\mathcal{G}(n,p)$ with the number of vertices $n = 2^m - 1$ and the edge probability $p = 2^{-r}$. Because of this, we may suppose that a function p of n giving diameter k in $\mathcal{G}(n,p)$ can be translated to a function $r(m)$ giving diameter k in $G_{m,r}$. Such a function in $\mathcal{G}(n,p)$ is known; the following result first appeared in [22].

Theorem 2.2.1. *Fix an integer $k > 1$ and choose ϵ such that $k - 1 < 1/\epsilon < k$. Let $p = n^{-1+\epsilon}$. Then with high probability, diam$(\mathcal{G}(n,p)) = k$ as $n \to \infty$.* □

A straightforward proof of this result can be achieved in two steps. First, one applies the first moment method to X, (the number of 'short' paths of length strictly less than k between any two vertices) to show that $X = 0$ with high probability. This shows that the graph is either disconnected or has diameter at least k. Then one applies Janson's Inequality [18] to Y_{ij}, the number of paths of length k between vertices i and j to show that $\mathbb{P}(Y_{ij} = 0) \leqslant e^{-n^\delta}$ for some $\delta > 0$; this exponential bound is needed as the result must support a union bound over all $\Theta(n^2)$ pairs of vertices.

To mimic the value of this function in $G_{m,r}$, we will need to have $2^{-r} = p = n^{-1+\epsilon} = (2^m - 1)^{-1+\epsilon} \approx 2^{m(-1+\epsilon)}$, or $r \approx (1-\epsilon)m$. As r must be an integer, we will formally define

$$r := \lfloor (1-\epsilon)m \rfloor.$$

Now fix a k and choose an ϵ such that $k - 1 < 1/\epsilon < k$. The corresponding value of r will now be used exclusively.

2.2.2 Short paths

In this subsection we provide the first moment part of the argument, showing that with high probability the number of short paths between given vertices is small. This argument essentially mimics that of the $\mathcal{G}(n,p)$ result, but we will need to count paths by dimension as well as length. To facilitate this, we will first show some results about the expected number of short paths, and a useful working lemma. Recall that ϵ is chosen such that $k - 1 < 1/\epsilon < k$.

Lemma 2.2.2. *Let l be a fixed positive integer. Then*

$$n^{l-1}p^l \to \begin{cases} 0 & \text{if } l < k \\ \infty & \text{if } l \geqslant k. \end{cases}$$

Proof. Recall that $n = 2^m - 1$, so as l is fixed, $n^{l-1} \sim 2^{m(l-1)}$ as $m \to \infty$, and thus it suffices to consider the function

$$2^{m(l-1)}p^l = 2^{m(l-1)}2^{-rl} = 2^{f(m,l)}, \tag{2.2.1}$$

where $f(m,l) = m(l-1) - rl = l(m-r) - m$. Now recall that $m - r = \lceil \epsilon m \rceil = \epsilon m + O(1)$, so we have

$$f(m,l) = l\lceil \epsilon m \rceil - m = (l\epsilon - 1)m + O(1)$$

as l is fixed. Hence if $l \geqslant k > 1/\epsilon$, we have $l\epsilon - 1 > 0$ and so $f(m,l) \to \infty$. Then from (2.2.1), $n^{l-1}p^l \to \infty$. On the other hand, if $l \leqslant k - 1 < 1/\epsilon$, we have $l\epsilon - 1 < 0$ and so $f(m,l) \to -\infty$. Hence from (2.2.1), $n^{l-1}p^l \to 0$. $\square$

Now let us set some notation; first order $V(G)$ as $u_1, \ldots, u_n$ in some arbitrary way. For any pair i, j with $i < j$ and l a positive integer, let $P_{ij}^{(l)}$ be the number of paths of length exactly l between u_i and u_j in $G_{m,r}$. The quantities $P_{ij}^{(\leqslant l)}, P_i^{(>l)}$, etc are defined in the obvious way.

The next two results can be read out of Proposition 3.5 in [12], but are not formally stated there.

Proposition 2.2.3. *For any two vertices u_i and u_j, $\mathbb{E}\left[P_{ij}^{(<k)}\right] \to 0$ as $m \to \infty$.*

Proof. By re-ordering we may assume that $i = 1$, $j = 2$. First fix $l < k$ and consider $\mathbb{E}\left[P_{ij}^{(l)}\right]$; we can compute this expectation in the standard way by bounding the number of paths of length l and the probability of each one being present. As discussed in the previous section, this probability depends on the dimension of the path, so it makes sense to count by dimension. Let $P_{l,d}$ be the set of all possible paths between u_1 and u_2 of length l and dimension d, so that this set is only non-empty if $d \leqslant l+1$. For such l, d we have from Proposition 2.1.4 that $\mathbb{P}(T \subset G) \leqslant p^{d-1}$ for any $T \in P_{l,d}$.

Now each possible path of length l is just a sequence of $l-1$ vectors $u_3, \ldots, u_{l+1}$ forming the interior vertices of the path. For each $i = 3, \ldots, l+1$, we can choose u_i in one of two ways:

1. We can take $u_i \in \langle u_1, \ldots, u_{i-1}\rangle$, in which case we have $\dim(\langle u_1, \ldots, u_i\rangle) = \dim(\langle u_1, \ldots, u_{i-1}\rangle)$;

2. We can take $u_i \notin \langle u_1, \ldots, u_{i-1}\rangle$, in which case we have $\dim(\langle u_1, \ldots, u_i\rangle) = 1 + \dim(\langle u_1, \ldots, u_{i-1}\rangle)$.

There are at most 2^{i-1} choices in the first instance and $2^m - 2^{i-1}$ in the second. Note that as $\dim(\langle u_1, u_2\rangle) = 2$ and the whole path has dimension d, we need to make a choice as in item 2 exactly $d-2$ times, and thus a choice as in item 1 $(l-1)-(d-2) = l-d+1$ times. Accounting for the choice of coordinates in which the dimension changes, we have

$$|P_{l,d}| \leqslant \binom{l}{d-2}(2^m-1)^{d-2} \cdot 2^{l(l-d+1)} = c_{l,d} \cdot n^{d-2}$$

as l and d are constants. Summing over all possible values of d gives

$$\mathbb{E}\left[P_{ij}^{(l)}\right] \leqslant \sum_{d=2}^{l+1} |P_{l,d}| \cdot p^{d-1} \leqslant \sum_{d=2}^{l+1} c_{l,d} \cdot n^{d-2} p^{d-1} \to 0$$

as $m \to \infty$ from Lemma 2.2.2, noting that $d-1 \leqslant l < k$. Then observe that

$$\mathbb{E}\left[P_{ij}^{(<k)}\right] = \mathbb{E}\left[\sum_{l=1}^{k-1} P_{ij}^{(l)}\right] = \sum_{l=1}^{k-1} \mathbb{E}\left[P_{ij}^{(l)}\right] \to 0$$

as $m \to \infty$; this gives the result. $\square$

We can use this result to give a first moment argument immediately.

Proposition 2.2.4. $\mathbb{P}(\mathrm{diam}(G_{m,r}) < k) \to 0$ *as* $m \to \infty$.

Proof. Note that a graph has diameter strictly less than k if and only if $d(u,v) < k$ for each pair u, v of vertices (where $d(u,v) = \infty$ if there is no path between u and v). Therefore in the random case,

$$\{\mathrm{diam}(G_{m,r}) < k\} = \bigcap_{i<j} \{d_{G_{m,r}}(v_i, v_j) < k\} = \bigcap_{i<j} \{P_{ij}^{(<k)} > 0\},$$

as any two vertices are at distance less than k if and only if there is some path of length less than k between them. Then it follows that

$$\mathbb{P}(\mathrm{diam}(G_{m,r}) < k) = \mathbb{P}\left(\bigcap_{i<j} P_{ij}^{(<k)} > 0\right) \leqslant \mathbb{P}\left(P_{12}^{(<k)} > 0\right) \leqslant \mathbb{E}\left[P_{12}^{(<k)}\right] \to 0$$

as $m \to \infty$ from above, where we have used Markov's Inequality. $\square$

This result shows that with high probability, $G_{m,r}$ is either disconnected or has diameter at least k. Thus if we can show that $G_{m,r}$ is connected (with high probability) it must have large diameter; this would be enough to verify Conjecture 1.1.5. By comparison with $\mathcal{G}(n,p)$, we expect the diameter to be exactly k; i.e. that there exists some uv path of length exactly k for some $u, v \in V(G_{m,r})$.

In $\mathcal{G}(n,p)$ we can apply Janson's Inequality to the number of paths of length k, giving

$$\mathbb{P}\left(\bigcup_{i<j} \{P_{ij}^{(k)} = 0\}\right) \leqslant \sum_{i<j} \mathbb{P}\left(P_{ij}^{(k)} = 0\right) \leqslant \binom{n}{2} \exp(n^{\Theta(k\epsilon - 1)}) \to 0 \qquad (2.2.2)$$

as $n \to \infty$ because $1/\epsilon < k$. We cannot appeal to this result in $G_{m,r}$ though, since the edges do not appear independently of one another.

Now return to $G_{m,r}$. As we can change bases, each of the events $\{P_{ij}^{(k)} = 0\}$ (for $i < j$) has the same probability, and thus we will have a similar result to (2.2.2) if we can show that $\mathbb{P}(P_{12}^{(k)} = 0) = o(n^{-2})$. In the light of this, we will try to find other ways of examining paths of length k in $G_{m,r}$.

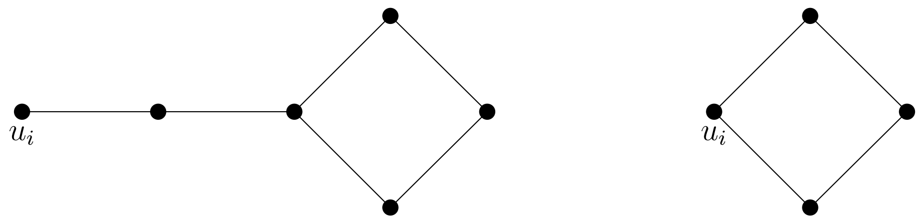

Figure 2.3: Two path-attached cycles; on the left $l = 6$ and $l' = 2$, while on the right $l = 4$ and $l' = 0$.

2.3 Short cycles

The idea we will use is as follows; given any two vertices, we will try to construct a path of length k by considering the successive neighbourhoods of our given two vertices. In the $\mathcal{G}(n,p)$ case, the low probability of seeing a short cycle means that there are a large number of vertices at distance $k/2$ from each, rendering it likely that there is some overlap and thus a path of length k. By adapting the methods for short paths we can make significant progress in the model $G_{m,r}$ as well. First we will need the following deterministic definition of a particular class of graphs.

Definition 2.3.1. For positive integers m, n with $n \geqslant m+3$, let $PC_{m,n} = ([n], E)$ be the graph on n vertices obtained by taking the path on $1, \ldots, n$ and adding the edge $m+1, n$. The resulting graph consists of a path of length m attached (at $m+1$) to a cycle of length $n-m$. Call $PC_{m,n}$ a *path-attached cycle*, and the leaf vertex 1 the *root*.

In the random graph model $G_{m,r}$, we set some notation similar to that for paths; for any vertex u_i and $l \geqslant 4$, let $C_i^{(l)}$ denote the number of path-attached cycles with exactly l vertices containing u_i as the root, with the length of the cycle at least 4 (note that these include cycles when the path is of length 0; see Figure 2.3). As before, the quantities $C_i^{(\leqslant l)}, C_i^{(>l)}$, etc. are then defined in the obvious way.

Proposition 2.3.2. *For any vertex u_i, $\mathbb{E}\left[C_i^{(<k)}\right] \to 0$ as $m \to \infty$.*

Proof. We will count the number of path-attached cycles in a similar manner to Proposition 2.2.3. As before, re-order such that $i = 1$ and fix $l < k$. Define $C(l', l, d)$ to be the set of possible graphs of total dimension d consisting of a path of length l' from u_1 to any vertex u' and then a cycle of length $l - l' \geqslant 4$ through u', not including any vertices from the path. The graphs in Figure 2.3 are examples with $l' = 2$ and $l' = 0$. As the total number of vertices is l, this set is only non-empty if $d \leqslant l$. As the cycle is of length at least 4, it cannot be of dimension 3, so by Proposition 2.1.6, $\mathbb{P}(H \subset G) \leqslant p^d$ for any $H \in C_{l',l,d}$.

Each possible $H \in C_{l',l,d}$ consists of an ordered list of $l-1$ vertices; the first l' describe the path and the remaining vertices the cycle. Counting all such combinations will give the size of $C_{l',l,d}$ up to a factor of 2, as reversing the list of vertices forming the cycle gives the same graph. As in Proposition 2.2.3, we have the option of raising the dimension of the list or not; as we start with a single vertex this time we need to raise the dimension exactly $d-1$ times. Hence we obtain

$$|C_{l',l,d}| \leqslant c_{l',l,d} \cdot n^{d-1} = O(n^{d-1})$$

for a constant $c_{l',l,d}$. Summing over l' and d gives

$$\mathbb{E}\left[C_i^{(l)}\right] \leqslant \sum_{d=3}^{l} \sum_{l'=0}^{l-4} |C_{l',l,d}| \cdot p^d \leqslant \sum_{d=3}^{l} O(n^{d-1}p^d) \to 0$$

as $m \to \infty$ from Lemma 2.2.2, noting that $d \leqslant l < k$. Summing over l from 4 to $k-1$ gives the result. □

Let us introduce more notation, some of which comes from basic graph theory. Let G be a graph and $u \in V(G)$; for any non-negative integer l, define the *lth neighbourhood* of u to be the set

$$\Gamma^l(u) = \{v \in V(G) \mid d(u,v) = l\},$$

i.e. the set of vertices at distance exactly l from u. Then $\Gamma^0(u) = \{u\}$ and $\Gamma^1(u) = \Gamma(u)$. We can then partition the component of G containing u into $\{\Gamma^i(u)\}_{i=0}^{n-1}$. If $x \in \Gamma^i(u)$, note that any neighbour of x is at distance at distance $i-1$, i or $i+1$ from u; for any such x define

$$\Gamma_x^+ := \Gamma(x) \cap \Gamma^{i+1}(u),$$

that is, the set of all neighbours of x at distance exactly $i+1$ from u. If we think of a process whereby we reveal $\Gamma^1(u), \Gamma^2(u), \ldots$ in turn, then Γ_x^+ is the set of new neighbours of x revealed at time $i+1$ (provided there is no overlap between vertices in $\Gamma^i(u)$). Proposition 2.3.2 gives the following result.

Proposition 2.3.3. *Fix a vertex u and define the events*

$$E_i = \{\Gamma_x^+ \cap \Gamma_y^+ = \emptyset \textit{ for all distinct } x, y \in \Gamma^i(u)\}$$

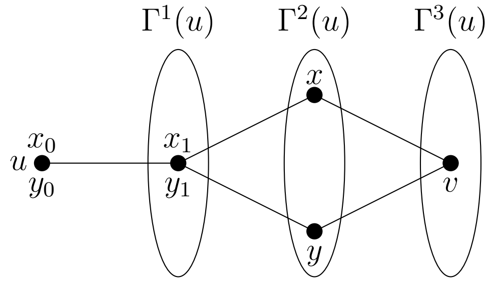

Figure 2.4: A common neighbour and corresponding cycle

and

$$F_j = \{e(\Gamma_x^+, \Gamma_y^+) = 0 \text{ for all distinct } x, y \in \Gamma^j(u)\}.$$

Let $n_1 = \lfloor (k-3)/2 \rfloor$ *and* $n_2 = n_1 - 1 = \lfloor (k-5)/2 \rfloor$. *Then* $\left(\bigcap_{i=1}^{n_1} E_i\right) \cap \left(\bigcap_{j=1}^{n_2} F_j\right)$ *occurs w.h.p. as* $m \to \infty$.

Proof. Fix $i = 1, \dots, n_1$ and consider the complementary event

$$E_i^{\mathrm{c}} = \{\text{there exist some } x, y \in \Gamma^i(u) \text{ with a common neighbour } v \in \Gamma^{i+1}(u)\}.$$

Suppose that E_i^{c} holds; then we can find paths $u = x_0, x_1, \dots, x_i = x$ and $u = y_0, y_1, \dots, y_i = y$ with $x_j, y_j \in \Gamma^j(u)$ for $j = 0, \dots, i$. In particular this means that $x_j \neq x_{j'}$ whenever $j \neq j'$. Let l be the last index for which $x_l = y_l$; this is well-defined as $x_0 = u = y_0$. We now have a cycle $x_l, \dots, x_i, v, y_i, \dots, y_l = x_l$ of length $2(i - l + 1) \geqslant 4$ and a path $u, x_1, \dots, x_l$ of length l attached to it (see Figure 2.4). The total number of vertices is $2(i+1) - l \leqslant 2i + 2 \leqslant k - 1$ by the choice of i, and so $C_1^{(<k)} \geqslant 1$ (where $u = u_1$ in the notation of the previous section). Hence

$$\mathbb{P}(E_i^{\mathrm{c}}) \leqslant \mathbb{P}\left(C_1^{(<k)} \geqslant 1\right) \leqslant \mathbb{E}\left[C_1^{(<k)}\right] \to 0$$

as $m \to \infty$ by Markov's Inequality and Proposition 2.3.2. Hence each E_i occurs with high probability. Similarly, for $j = 1, \dots, n_2$ consider

$$F_j^{\mathrm{c}} = \{\text{there exists some } x, y \in \Gamma^j(u), x' \in \Gamma_x^+, y' \in \Gamma_y^+ \text{ with } x'y' \in E(G)\}.$$

If F_j^{c} holds we can find paths $u = x_0, x_1, \dots, x_j = x, x'$ and $u = y_0, y_1, \dots, y_j = y, y'$ as before; letting l be as above we obtain a cycle $x_l, \dots, x_j, x', y', y_j, \dots, y_l$ with a path $u, x_1, \dots, x_l$ of length l attached to it (see Figure 2.5). The cycle has length at least 5 and the total number of vertices is at most $2j + 3 \leqslant k - 1$ by the choice of j. It then follows as above that $\mathbb{P}(F_j^{\mathrm{c}}) \to 0$ as $m \to \infty$. Taking the intersection over all of these events gives the result. $\square$

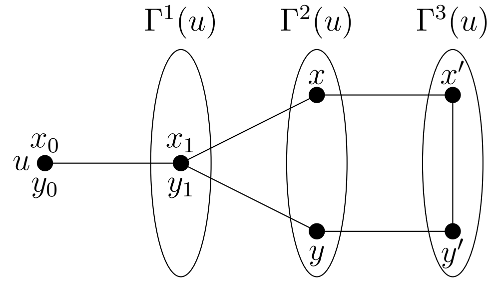

Figure 2.5: An edge within a neighbourhood

This tells us that there are a lot of new neighbours at each stage; new neighbourhoods of different vertices rarely overlap and there aren't many edges between vertices at the same level. What we have not considered is the case where two new neighbours of the same vertex x are adjacent, the reason being that arguing as in Proposition 2.3.3 would lead to consideration of unicyclic graphs where the cycle could possibly be a triangle. If this triangle is three dimensional then a similar argument will work, but two dimensional triangles were seen to be a special case in Section 2.1 and they are similarly ill-behaved here. Indeed, suppose $y \in \Gamma(x)$ so $[x, y] = 0$; then $[x, x + y] = [y, x + y] = 0$ so $x + y \in \Gamma(x)$ and y and $x + y$ are adjacent. This means that we can not expect there to be vanishingly few edges within individual neighbourhoods, although we need only consider edges that will form such a two dimensional triangle. To understand this further, we need to further investigate the degree and neighbourhood of a vertex.

Chapter 3

Diameters 2 and 3

In the last chapter of this part, we consider the first two non-trivial cases, where $1/2 < \epsilon < 1$ and $1/3 < \epsilon < 1/2$ and we expect to see diameter 2 and 3 respectively. We begin by considering the degree and neighbourhood of a vertex; these random variables are straightforward to calculate and supply the expected answer when $1/2 < \epsilon < 1$. We can't achieve the same success in the second case; we instead show how several standard tools fail for this model, suggesting that the dependencies between the basic variables are too complicated for the model to be viable.

3.1 Degrees and neighbourhoods

3.1.1 Prepatory results

In light of the discussion at the end of Section 2.3, we first consider the number of edges between a vertex and a subspace (as opposed to an arbitrary set). We will need some more notation.

Definition 3.1.1. For integers n, $k \geqslant 0$, we define $N_{n,k}$ as

$$N_{n,k} = \begin{cases} \prod_{i=0}^{k-1}(2^n - 2^i) & \text{if } n \geqslant k \\ 0 & \text{otherwise.} \end{cases}$$

The *Gaussian binomial coefficient* (over $\mathbb{F}_2$) is given by

$$\binom{n}{k}_2 = \frac{N_{n,k}}{N_{k,k}}$$

for $n \geqslant k \geqslant 0$.

Observe that if V is any n-dimensional vector space over $\mathbb{F}_2$ then (by Proposition 2.1.2) $N_{n,k}$ counts the number of *ordered* sets of k linearly independent vectors in V;

having selected i vectors ($0 \leqslant i < k$) we cannot choose any of the 2^i vectors in their span.

Now any k-dimensional subspace W of V is spanned by a linearly independent set of size k, and so is defined by one of the $N_{n,k}$ sets as above. But for each such subspace W there are exactly $N_{k,k}$ ordered bases, obtained by choosing a set of k linearly independent vectors in a k-dimensional space. This shows that $\binom{n}{k}_2$ counts the number of k-dimensional subspaces of V.

There is a useful alternative form of this expression, revealing its symmetry.

Lemma 3.1.2. *For $n \geqslant k \geqslant 0$,*

$$\binom{n}{k}_2 = \frac{N_{n,n}}{2^{k(n-k)} N_{k,k} N_{n-k,n-k}} = \binom{n}{n-k}_2$$

Proof. First observe that

$$\prod_{i=k}^{n-1} (2^n - 2^i) = \prod_{i=k}^{n-1} 2^k (2^{n-k} - 2^{i-k}) = 2^{k(n-k)} \prod_{i=0}^{n-k-1} (2^{n-k} - 2^i) = 2^{k(n-k)} N_{n-k,n-k}.$$

Therefore

$$\begin{aligned}\binom{n}{k}_2 = \frac{N_{n,k}}{N_{k,k}} &= \frac{\prod_{i=0}^{k-1}(2^n - 2^i)}{N_{k,k}} \cdot \frac{\prod_{i=k}^{n-1}(2^n - 2^i)}{2^{k(n-k)} N_{n-k,n-k}} \\ &= \frac{\prod_{i=0}^{n-1}(2^n - 2^i)}{2^{k(n-k)} N_{k,k} N_{n-k,n-k}} \\ &= \frac{N_{n,n}}{2^{k(n-k)} N_{k,k} N_{n-k,n-k}}.\end{aligned}$$

Now note that replacing k with $n-k$ in this last line leaves the expression invariant, giving the result. □

This notation eases the proof of the following result. Recall that in Section 1.3.1, we took a random linear map $F : V \to W$ by choosing the image of each basis vector under F uniformly at random from W, independently for each basis vector.

Proposition 3.1.3. *Let $(V_m)_{m=1}^{\infty}$ and $(W_m)_{m=1}^{\infty}$ be sequences of vector spaces over $\mathbb{F}_2$, with $\dim(V_m) = m$ and $\dim(W_m) = r(m) = r$ for each m. Let $(F_m)_{m=1}^{\infty}$ be a sequence of random linear maps, with $F_m : V_m \to W_m$ for all m, constructed as in Section 1.3.1. Then if $m - r \to \infty$, $\mathrm{rk}(F_m) = r$ w.h.p. as $m \to \infty$.*

Proof. Suppose that $\mathrm{rk}(F_m) = \dim(\mathrm{im}(F_m)) < r$. Then $F_m(V_m)$ is contained in a subspace of W_m of dimension $r-1$. From above, the number of such subspaces is given by

$$\binom{r}{r-1}_2 = \binom{r}{1}_2 = \frac{2^r - 1}{2^1 - 1} = 2^r - 1.$$

Now any such subspace U consists of 2^{r-1} vectors, so the uniformity of the map F_m gives

$$\mathbb{P}(F_m(x_i) \in U) = \frac{2^{r-1}}{2^r} = \frac{1}{2},$$

independently for each i. Hence $\mathbb{P}(F_m(V_m) \leqslant U) = \mathbb{P}(F_m(x_i) \in U \text{ for all } i) = 2^{-m}$. A simple union bound then gives

$$\mathbb{P}(\dim(F_m(V_m)) < r) = \mathbb{P}\left(\bigcup_{U:\dim(U)=r-1} \{F_m(V_m) \leqslant U\} \right) \leqslant 2^r 2^{-m} \to 0$$

as $m \to \infty$, by the assumed condition. □

3.1.2 The degree of a vertex

As mentioned at the start of this section, we first consider the number of edges between a vertex and a subspace.

Proposition 3.1.4. *Let $U \subseteq V(G_{m,r})$ be such that $U_0 := U \cup \{0\}$ is a d-dimensional subspace of V, and let $x \in V(G_{m,r}) \setminus U$. Then $\Gamma_U(x) \cup \{0\}$ is a subspace of U_0, and in particular*

$$d_U(x) = |\Gamma_U(x)| = 2^l - 1$$

for some $l = 0, \ldots, d$. Moreover,

$$\mathbb{P}(d_U(x) = 2^l - 1) = \binom{d}{l}_2 \cdot N_{r,d-l} \cdot p^d \tag{3.1.1}$$

for $l = 0, \ldots, d$.

Proof. Let $\mathcal{B}_U = \{x_1, \ldots, x_d\}$ be a basis of U_0; as $x \notin U$ we can extend $\mathcal{B}_U \cup \{x\}$ to a basis $\{x_1, \ldots, x_d, x_{d+1} = x, \ldots x_m\}$ of V. We will work with this basis; the random bilinear map $[\cdot,\cdot]$ induces a random linear map $F : U_0 \to W$ given by $F(x_j) = [x, x_j]$ for $j = 1, \ldots, d$. This is exactly the type of random linear map defined in Section 1.3.1.

For any $u \in U$, x is adjacent to u if and only if $F(u) = [x,u] = 0$, i.e. $u \in \ker F$. Hence $\Gamma_U(x) = (\ker F) \setminus \{0\}$, and so $\Gamma_U(x) \cup \{0\} = \ker F$ is a subspace of U_0. In particular, its dimension l is the nullity of the map F. This proves the first statement.

Now let $K \leqslant U_0$ be a subspace of dimension l, for some $l = 0, \dots, d$. Let $\{u_1, \dots, u_l\}$ be a basis for K and extend this to a basis $\{u_1, \dots, u_l, u_{l+1}, \dots u_d\}$ of U_0. From Proposition 1.3.2, we know that $F(u_1), \dots, F(u_d)$ are independent and define the map F.

Now consider $\mathbb{P}(\ker F = K)$. If we suppose that $\ker F = K$ then $\text{null}(F) = \dim(K) = l$, so by the Rank–Nullity Theorem $\text{rk}(F) = \dim(U) - l = d - l$. But also $F(U_0) = \langle F(u_1), \dots, F(u_d)\rangle = \langle F(u_{l+1}), \dots, F(u_d)\rangle$ as $F(u) = 0$ for $u \in K$. It follows that $\{F(u_{l+1}), \dots, F(u_d)\}$ is linearly independent.

Conversely, suppose that $F(u_i) = 0$ for $i \leqslant l$ and $\{F(u_{l+1}), \dots, F(u_d)\}$ is linearly independent. The first condition ensures that $K \leqslant \ker F$. As before we have that $F(U_0) = \langle F(u_{l+1}), \dots, F(u_d)\rangle$ and as this set is linearly independent it follows that $\text{rk}(F) = d - l$. Hence $\text{null}(F) = l$, i.e. $\dim(\ker F) = l$. But K is an l-dimensional subspace of $\ker F$ and so $\ker F = K$.

So we have

$$\mathbb{P}(\ker F = K) = \mathbb{P}\big(F(u_i) = 0 \text{ for } i \leqslant l, \{F(u_{l+1}), \dots, F(u_d)\} \text{ is l.i.}\big).$$

There are precisely $N_{r,d-l}$ ways of choosing an ordered linearly independent set of the correct size in the target space W, and the values $(F(u_i))_{i=1}^d$ are chosen uniformly and independently. Hence

$$\mathbb{P}(\ker F = K) = N_{r,d-l} \cdot 2^{-rd} = N_{r,d-l} \cdot p^d.$$

The result now follows because there are $\binom{d}{l}_2$ candidates for the subspace K. □

Note that there is a certain similarity between the form of this expression and the binomial distribution which holds in $\mathcal{G}(n, p)$; the binomial coefficient in the latter is replaced by a Gaussian binomial coefficient here, and the expressions in p are also quite similar. In fact, we can obtain the same exponential bound on the probability of there being no edges.

Corollary 3.1.5. *With x and U as above, $\mathbb{P}(d_U(x) = 0) \leqslant e^{-|U|p} = e^{-\mathbb{E}[d_U(x)]}$.*

Proof. To get no edges, we need to put $l = 0$ into (3.1.1); the right hand side reduces to $N_{r,d} \cdot p^d$. Using Definition 3.1.1 we see

$$\mathbb{P}(d_U(x) = 0) = \prod_{i=0}^{d-1}(2^r - 2^i) \cdot 2^{-rd} = \prod_{i=0}^{d-1}(1 - 2^{i-r}) \leqslant \prod_{i=0}^{d-1} e^{-2^{i-r}} = e^{-S},$$

where

$$S = \sum_{i=0}^{d-1} 2^{i-r} = 2^{-r} \sum_{i=0}^{d-1} 2^i = 2^{-r}(2^d - 1) = |U|p$$

as $U \cup \{0\}$ is a d-dimensional subspace and so has size 2^d. The expectation is easily computed as there are $|U|$ possible edges, each present with probability p. $\square$

A careful look at the result in Proposition 3.1.4 gives something much stronger when $d > r$.

Corollary 3.1.6. *For $d > r$, the minimum possible value of $d_U(x)$ is $2^{d-r} - 1$.*

Proof. From Definition 3.1.1, $N_{r,d-l} = 0$ if $r < d-l$, i.e. if $l < d-r$. From Proposition 3.1.4 it then follows that $\mathbb{P}(d_U(x) = 2^l - 1) = 0$ for such l and hence the minimum degree is $2^{d-r} - 1$. This corresponds to the map F in the proof of Proposition 3.1.4 having nullity at least $d - r$. $\square$

By considering a large vector space we can compute the distribution of the full degree of any vertex.

Proposition 3.1.7. *For any fixed vertex $x \in V$, we have that*

$$d(x) = 2(2^l - 1) = 2^{l+1} - 2$$

for some $l = 0, \ldots, m-1$. (Note that if $l = 0$, x is isolated and if $l = m-1$, x is adjacent to every other vertex). Further, the distribution of degrees is given by

$$\mathbb{P}(d(x) = 2^{l+1} - 2) = \binom{m-1}{l}_2 \cdot N_{r,m-1-l} \cdot p^{m-1} \tag{3.1.2}$$

for $l = 0, \ldots, m-1$.

Proof. Extend x to a basis $\{x = x_1, x_2, \ldots, x_m\}$ of V and let $U_0 = \langle x_2, \ldots, x_m \rangle$. Putting $U = U_0 \setminus \{0\}$, we see that $V(G) \setminus \{x\}$ is the disjoint union of U and $x + U$. From the Proposition 3.1.4, $d_U(x) = 2^l - 1$ for some $l = 0, \ldots, m-1$, where l is the nullity of the map F defined in the proof, and the probability of this degree is given by the right hand side of (3.1.2).

Let $v \in x + U$, so $v = x + u$ for some $u \in U$. Then $[x, v] = [x, x+u] = [x, u]$, so v is a neighbour of x if and only if u is. It follows that $\Gamma_{x+U}(x) = x + \Gamma_U(x)$, of size $2^l - 1$. Thus the full neighbourhood has size $2(2^l - 1) = 2^{l+1} - 2$ for some $l = 0, \ldots, m-1$, and (3.1.2) gives the correct distribution. $\square$

We can now obtain the following interesting results.

Corollary 3.1.8. *The minimum possible degree of any fixed vertex x is $2^{m-r} - 2 \sim \mathbb{E}[d(x)]$.*

Proof. As before, the right hand side of (3.1.2) vanishes if $r < m-1-l$, i.e. $l < m-1-r$. Putting this into the left hand side gives the minimum degree as $2^{(m-1-r)+1} - 2 = 2^{m-r} - 2$.

We can compute the expected degree of x as before; there are $2^m - 2$ possible neighbours, and each one is adjacent to x with probability $p = 2^{-r}$. Hence

$$\mathbb{E}[d(x)] = (2^m - 2)2^{-r} = (n-1)p$$

as in the binomial case. Then

$$\frac{2^{m-r} - 2}{2^{m-r} - 2^{1-r}} = \frac{1 - 2^{r-m}}{1 - 2^{1-2r-m}} \to 1$$

as $m \to \infty$, giving the result. □

Corollary 3.1.9. *For any fixed vertex x, $\mathbb{P}(d(x) = 2^{m-r} - 2) \to 1$ as $m \to \infty$.*

Proof. We have that $m - r = \lceil \epsilon m \rceil \to \infty$ as $m \to \infty$, so $m - 1 - r \to \infty$ as well. From Proposition 3.1.3, we have $\mathbb{P}(\mathrm{rk}(F) = r) \to 1$ as $m \to \infty$, where F is defined as in Proposition 3.1.4 with $d = m - 1$. But $\mathrm{rk}(F) = r$ if and only if $\mathrm{null}(F) = m - 1 - r$. From Propositions 3.1.4 and 3.1.7, this is equivalent to $d(x) = 2^{(m-1-r)+1} - 2 = 2^{m-r} - 2$. □

Corollary 3.1.10. *G has no components of size less than $2^{m-r} - 1 \sim n^\epsilon$. In particular, G has no isolated vertices for $m - r = \lceil \epsilon m \rceil > 1$.*

Proof. As each vertex has at least $2^{m-r} - 2$ neighbours the minimum size of a component is $2^{m-r} - 1 = 2^{\lceil \epsilon m \rceil} - 1$. The result follows as $n = 2^m - 1$. □

Note that this last result has some significance; it was shown by Erdős and Rényi in [6] that in $\mathcal{G}(n,p)$ the thresholds for connectivity and minimum degree at least 1 coincide, while Proposition 2.2.4 showed that proving connectivity suffices to show arbitrarily large diameter.

3.2 The diameter 2 case

We can now use Proposition 3.1.4 to confirm the conjecture in the case of diameter 2. This proof is entirely deterministic.

Proposition 3.2.1. *If ϵ is chosen such that $1/2 < \epsilon < 1$ then $G_{m,r}$ has diameter 1 or 2 for m sufficiently large.*

Proof. Let x_1, x_2 be an arbitrary pair of vertices and extend this pair to a basis $\{x_1, x_2, \ldots, x_m\}$ of V. Let $U_0 = \langle x_3, \ldots, x_m \rangle$ and $U = U_0 \setminus \{0\}$. Write $\Gamma_i = \Gamma_U(x_i) \cup \{0\}$ for $i = 1, 2$; by Proposition 3.1.4 Γ_1 and Γ_2 are both subspaces of U_0; by Corollary 3.1.6 they are both of dimension at least $(m-2) - r = \lceil \epsilon m \rceil - 2 \geqslant \epsilon m - 2$. The sum formula from linear algebra then gives

$$\dim(\Gamma_1) + \dim(\Gamma_2) - \dim(\Gamma_1 \cap \Gamma_2) = \dim(\Gamma_1 + \Gamma_2) \leqslant \dim(U_0) = m - 2$$

and hence

$$\begin{aligned}\dim(\Gamma_1 \cap \Gamma_2) &\geqslant \dim(\Gamma_1) + \dim(\Gamma_2) - m + 2 \\ &\geqslant 2(\epsilon m - 2) - m + 2 \\ &\geqslant (2\epsilon - 1)m - 2 \\ &\geqslant 1\end{aligned}$$

for m sufficiently large, since $2\epsilon - 1 > 0$. In this case there is a non-zero $u \in \Gamma_1 \cap \Gamma_2$. But then $u \in U$ is a neighbour of both x_1 and x_2, so $x_1 u x_2$ is a path of length 2. As x_1 and x_2 were arbitrary, each pair of vertices is connected by a path of length 2. This combined with the first moment analysis gives the result, and also gives that $\mathbb{P}(\mathrm{diam}(G_{m,r}) = 2) \to 1$ as $m \to \infty$. □

Having studied the neighbourhood of any vertex, we can now sketch how to build paths from two vertices as mentioned at the end of Section 2.3. With high probability, the neighbourhood of any vertex will consist of a vector subspace of dimension $m - r + O(1)$ and a translate (coset) of the same space; the necessary edges between them will be the only edges within this set. For any two vertices x and y, we can now expand their neighbourhoods as in Proposition 2.3.3; at distance $(k-1)/2$ (for k odd) we will expect to see around $2^{(k-1)(m-r)/2}$ *disjoint* subspaces. This indicates that we should consider the probability of there being no edge between two disjoint unions of subspaces of the appropriate size; if this is exponentially small, the probability of there being no edge between $\Gamma^{(k-1)/2}(x)$ and $\Gamma^{(k-1)/2}(y)$ should also be exponentially small. In this case we would conclude that with high probability an edge (and hence path of length k) exists.

The natural starting point is the case $k = 3$, where we are only considering 2 different subspaces. Unfortunately, this problem proves to be difficult, as explained in the next section.

3.3 The diameter 3 case

3.3.1 Model dependencies and first approaches

We now consider the case where $1/3 < \epsilon < 1/2$. As in the proof of Proposition 3.2.1, fix two vertices x_1 and x_2 and extend to a basis $\{x_1, x_2, \ldots, x_m\}$ of V. Let $U_0 = \langle x_3, \ldots, x_m \rangle$ and $U = U_0 \setminus \{0\}$. Write $\Gamma_i = \Gamma_U(x_i) \cup \{0\}$ for $i = 1, 2$; as before, these are both subspaces of U of dimension at least $(m-2) - r = \lceil \epsilon m \rceil - 2$. Note that

$$\mathbb{P}(\Gamma_1 \cap \Gamma_2 \neq \{0\}) \leqslant \mathbb{P}\left(P_{12}^{(2)} > 0\right) \leqslant \mathbb{E}\left[P_{12}^{(2)}\right] \to 0$$

by the choice of ϵ, so we can assume that $\Gamma_1 \cap \Gamma_2 = \{0\}$ (of course if $\Gamma_1 \cap \Gamma_2 \neq \{0\}$ then there is a path of length 2 between x_1 and x_2 anyway). We will hence treat Γ_1 and Γ_2 as disjoint subspaces. We also have that

$$\mathbb{P}\left(P_{12}^{(3)} = 0\right) \leqslant \mathbb{P}(e(\Gamma_1, \Gamma_2) = 0),$$

so it suffices to show that $\mathbb{P}(e(\Gamma_1, \Gamma_2) = 0) = o(n^{-2})$. We can recast this problem in terms of a random bipartite graph, as set out formally below.

Definition 3.3.1. Let $1/3 < \epsilon < 1/2$, and for each positive integer m put $r = r(m) = \lfloor (1-\epsilon)m \rfloor$. Let $U = \langle u_1, \ldots, u_d \rangle$ and $V = \langle v_1, \ldots, v_d \rangle$ be two disjoint vector spaces over $\mathbb{F}_2$, each of dimension $d = m - r + O(1) = \epsilon m + O(1)$, with given bases. Let W be a third vector space over $\mathbb{F}_2$ of dimension r.

Define a random bilinear map $B : U \times V \to W$ as follows: for each pair $1 \leqslant i, j \leqslant d$, select $B(u_i, v_j)$ uniformly at random from the 2^r elements of W, independently for each pair i, j, and then extend this in the unique way to the whole of $U \times V$.

Now create a random bipartite graph: the vertex classes are $U \setminus \{0\}$ and $V \setminus \{0\}$ and $\mathcal{E} = (U \setminus \{0\}) \times (V \setminus \{0\})$ is the set of all possible edges. For each $uv \in \mathcal{E}$, the edge uv is present if and only if $B(u, v) = 0$. Define for each uv the event $A_{uv} = \{B(u, v) = 0\} = \{uv \text{ is present}\}$. Let X be the number of edges, that is $X = \sum_{uv \in \mathcal{E}} 1_{A_{uv}}$. We wish to find a tight (i.e. exponential) bound on $\mathbb{P}(X = 0)$.

Table 3.1 contains the results of simulations (performed in R) for 9 values of the pair (d, r), chosen to produce an increasing sequence of means. 10 000 simulations were performed in each case and the proportion for which $X = 0$ gives an estimate of the probability of this event. These results seem to indicate an exponential bound. Before continuing, note that several results carry over from the previous analysis; each edge is present with probability $p = 2^{-r}$ (i.e. $\mathbb{P}(A_e) = 2^{-r}$ for each $e \in \mathcal{E}$) and

d	r	$\mathbb{E}[X]$	Sample mean	$e^{-\mathbb{E}[X]}$	Estimated probability
2	3	1.125	1.116	0.3247	0.2694
2	2	2.25	2.249	0.1054	0.0456
3	4	3.063	3.063	0.04677	0.0255
3	3	6.125	6.137	2.187×10^{-3}	5×10^{-4}
4	5	7.031	7.038	8.840×10^{-4}	1×10^{-4}
6	9	7.752	7.795	4.299×10^{-4}	2×10^{-4}
3	2	12.25	12.32	4.785×10^{-6}	0
4	4	14.06	14.10	7.811×10^{-7}	0
5	6	15.02	15.07	3.012×10^{-7}	0

Table 3.1: The results of 10 000 simulations

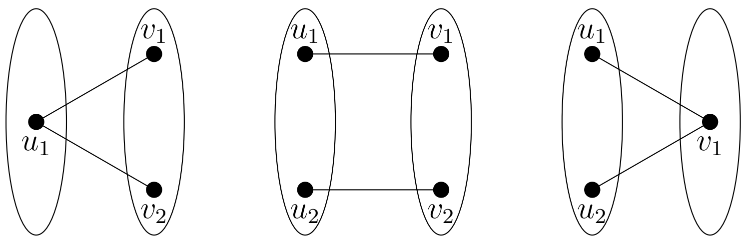

Figure 3.1: Configurations of two edges.

we can work with any bases of U and V. Let $N = |\mathcal{E}|$ be the total number of possible edges, so $N = (2^d - 1)^2 \sim 2^{2d}$. We start with a second moment analysis, which at least shows that $\mathbb{P}(X = 0) = o(1)$.

Lemma 3.3.2. *In the above notation,*

1. $\mu := \mathbb{E}[X] = Np \to \infty$ *as* $m \to \infty$;

2. $\mathbb{P}(X = 0) \leqslant \frac{1}{\mu} \to 0$ *as* $m \to \infty$.

Proof. It is trivial to see that $\mu = Np$. Now note that $Np \sim 2^{2d} 2^{-r} = 2^{f(m)}$, where

$$f(m) = 2d - r = 2\epsilon m - (1 - \epsilon)m + O(1) = (3\epsilon - 1)m + O(1) \to \infty$$

as $m \to \infty$ by the choice of ϵ. For the second statement, we need to calculate

$$\mathbb{E}[X(X - 1)] = \sum_{e \in \mathcal{E}} \sum_{f \neq e} \mathbb{P}(A_e \cap A_f).$$

There are only three possible configurations (two of which are equivalent) for two edges; see Figure 3.1. In the first case, extend $\{u_1\}$ to a basis of U and $\{v_1, v_2\}$ to a basis of V; working with these bases, $B(u_1, v_1)$ and $B(u_1, v_2)$ are chosen independently so $\mathbb{P}(A_{u_1v_1} \cap A_{u_2v_2}) = p^2$. Interchanging U and V gives the same probability in the third case; in the second case we can extend $\{u_1, u_2\}$ to a basis of U and get the same result. This gives that $\mathbb{P}(A_e \cap A_f) = p^2$ for *every* distinct pair e, f, i.e. the events $(A_e)_{e\in\mathcal{E}}$ are pairwise independent. Thus $\mathbb{E}[X(X-1)] = N(N-1)p^2$ and by Chebyshev's Inequality,

$$\begin{aligned}\mathbb{P}(X = 0) &\leqslant \frac{\operatorname{var}(X)}{\mathbb{E}[X]^2} \\ &= \frac{\mathbb{E}[X(X-1)] + \mathbb{E}[X] - \mathbb{E}[X]^2}{\mathbb{E}[X]^2} \\ &= \frac{N(N-1)p^2}{N^2p^2} + \frac{1}{\mu} - 1 \\ &= \frac{1}{\mu} + \left(1 - \frac{1}{N}\right) - 1 \\ &\leqslant \frac{1}{\mu},\end{aligned}$$

giving the result. □

Note that although the events $(A_e)_{e\in\mathcal{E}}$ are pairwise independent, they are certainly not mutually independent, which is the root cause of the difficulty in working with them.

3.3.2 Concentration

To gain more insight into the model dependencies, it may help to consider another interpretation of the structure of $\mathcal{E}$. First note that the target vector space W can be equipped with a multiplication in such a way that it becomes isomorphic to the field $F = \mathbb{F}_{2^r}$; this follows because F contains $\mathbb{F}_2$ as its characteristic field and can therefore be thought of as an r-dimensional vector space over $\mathbb{F}_2$. This allows us to carry the multiplication over the (linear) isomorphism with W, and thus consider each $B(u, v)$ as an element of a *field* rather than a vector space.

Now take a fixed pair of bases for U and V, say $\{u_1, \ldots, u_d\}$ and $\{v_1, \ldots, v_d\}$ and assign an ordering $Z_1, \ldots, Z_M$ to the list of variables $\{B(u_i, v_j)\}_{1\leqslant i,j\leqslant d}$ (where $M = d^2$). Then the vector $Z = (Z_1, \ldots, Z_M)$ can be interpreted as a random element of F^M where each component is chosen uniformly at random from F, independently of all other components. This will allow us to consider the problem geometrically.

Now observe that the random variable $B(u,v)$ corresponding to the edge uv can be written as a linear combination of the Z_is; if $u = \sum_{i=1}^{d} \lambda_i u_i$ and $v = \sum_{j=1}^{d} \rho_j v_j$ (for $\lambda_i, \rho_j \in \mathbb{F}_2$) then

$$B(u,v) = \sum_{i=1}^{d}\sum_{j=1}^{d} \lambda_i \rho_j B(u_i, v_j) = \sum_{l \in S} Z_l,$$

where $S \subseteq [M]$ is determined by which of the products $\lambda_i \rho_j$ equal 1. Hence each possible edge e can be assigned a set $S_e \subseteq [M]$ as above.

As an example, suppose $d = 2$ (so $M = 4$) and we order the variables as $Z_1 = B(u_1, v_1)$, $Z_2 = B(u_1, v_2)$, $Z_3 = B(u_2, v_1)$, $Z_4 = B(u_2, v_2)$. Now consider the edge e given by $u_1 + u_2, v_1$. As $B(u_1 + u_2, v_1) = B(u_1, v_1) + B(u_2, v_1) = Z_1 + Z_3$, we have $S_e = \{1, 3\}$.

Although we can not apply Janson's Inequality to this situation as the variables are not of the correct form, an estimate of the dependencies as above allows a similar inequality of Suen [31] (later improved by Janson [19]) to be applied. This inequality (Theorem 4 in [19]) is given below.

Theorem 3.3.3. *Let $\mathcal{I}$ be a finite or countable set indexing a set of indicator random variables $(X_i)_{i\in\mathcal{I}}$ and Γ a superdependency graph on $\mathcal{I}$, that is for any disjoint $A, B \subseteq \mathcal{I}$ with $e(A,B) = 0$, the two families $(X_i)_{i\in A}$ and $(X_j)_{j\in B}$ are independent from one another.*

Define $X = \sum_{i\in\mathcal{I}} X_i$, $\mu = \mathbb{E}(X)$,

$$\Delta := \frac{1}{2}\sum_{i\in\mathcal{I}}\sum_{j\sim i} \mathbb{E}[X_i X_j] = \frac{1}{2}\sum_{i\in\mathcal{I}}\sum_{j\sim i} \mathbb{P}(X_i = X_j = 1)$$

and

$$\Delta_0 := \frac{1}{2}\sum_{i\in\mathcal{I}}\sum_{j\sim i} \mathbb{E}[X_i]\mathbb{E}[X_j] = \frac{1}{2}\sum_{i\in\mathcal{I}}\sum_{j\sim i} \mathbb{P}(X_i = 1)\mathbb{P}(X_j = 1),$$

where $\sim$ denotes adjacency in Γ. Then

$$\mathbb{P}(X = 0) \leqslant e^{-\mu^2/\max\{32\Delta, 48\Delta_0, 4\mu\}}.$$

□

There are many variations on this result, but they all centre on estimating $\mathbb{P}(X = 0)$ by an exponential expression in μ and some term accounting for the dependence.

Let us fit our model into this framework. Our set of indicator variables is $(A_e)_{e\in\mathcal{E}}$ where $\mathcal{E}$ denotes the set of all N possible edges, so $V(\Gamma) = \mathcal{E}$. We have shown that these events are *pairwise* independent, i.e. $\mathbb{P}(A_e \cap A_f) = \mathbb{P}(A_e)\mathbb{P}(A_f)$ for $e \neq f$, and so $\Delta = \Delta_0$ regardless of our choice of Γ.

As described above, we assign a set $S_e \subseteq [M]$ to each variable, and then A_e depends on $\sum_{i \in S_e} Z_i$, where the variables $Z_1, \ldots, Z_M$ are independent, each chosen uniformly from the vector space W (these correspond to the choices for pairs of basis vectors). In the light of this, there is a natural choice of Γ, where we define edges by

$$e \sim f \text{ if and only if } S_e \cap S_f \neq \emptyset.$$

It is straightforward to check that is a superdependency graph; if $e(A, B) = 0$ then the corresponding elements of $\mathcal{E}$ are determined by disjoint sets of independent variables and so are independent. Unfortunately, this choice of Γ will not provide a tight bound.

Proposition 3.3.4. *With Γ defined as above, $E(\Gamma) \sim N^2/2$ and so $\Delta = \Delta_0 \sim \mu^2/2$.*

Proof. First note that $N = (2^d - 1)^2 \sim 4^d$. Fix bases $\{u_1, \ldots, u_d\}$ and $\{v_1, \ldots, v_d\}$ of U and V respectively. For each $u \in U$ we associate a unique vector $\lambda^u := (\lambda_1^u, \ldots, \lambda_d^u) \in \mathbb{F}_2^d$ with $u = \sum_{i=1}^d \lambda_i^u u_i$; let n_u denote the number of these coefficients equal to 1. We can similarly define ρ^v and n_v for $v \in V$.

For $1 \leqslant k, l \leqslant d$, consider the edge corresponding to the pair $(u, v) = (u_1 + \cdots + u_k, v_1 + \cdots + v_l)$. The event A_e is given by

$$A_e = \{B(u, v) = 0\} = \left\{ \sum_{i=1}^{k} \sum_{j=1}^{l} B(u_i, v_j) = 0 \right\},$$

so $(u, v) \sim (u', v')$ in Γ if and only if $\lambda_i^{u'} = \rho_j^{v'} = 1$ for some $i = 1, \ldots, k$, $j = 1, \ldots, l$. There are $2^{d-k} - 1$ non-zero vectors u' which don't satisfy this property (as we can only choose $\lambda_{k+1}^{u'}, \ldots, \lambda_d^{u'}$) and similarly $2^{d-l} - 1$ non-zero vectors v' not satisfying it. This means that (u, v) is adjacent to exactly $(2^d - 2^{d-k})(2^d - 2^{d-l}) - 1$ other vertices of Γ (we do not count (u, v) itself). Notice that by reordering the bases, this is true for any pair (u, v) for which $n_u = k$ and $n_v = l$. There are exactly $\binom{d}{k}\binom{d}{l}$ such pairs, and so by the Handshaking Lemma,

$$\begin{aligned}
2e(\Gamma) &= \sum_{e \in E} d_\Gamma(e) \\
&= \sum_{k=1}^{d} \sum_{l=1}^{d} \binom{d}{k} \binom{d}{l} [(2^d - 2^{d-k})(2^d - 2^{d-l}) - 1] \\
&= \left(\sum_{k=1}^{d} \binom{d}{k} (2^d - 2^{d-k}) \right)^2 - \left(\sum_{k=1}^{d} \binom{d}{k} \right)^2
\end{aligned}$$

Noting that the summand in the first bracket vanishes when $k = 0$, we see that

$$\sum_{k=1}^{d} \binom{d}{k} (2^d - 2^{d-k}) = \sum_{k=0}^{d} \binom{d}{k} (2^d - 2^{d-k}) = 2^d \cdot 2^d - (1+2)^d,$$

and hence $2e(\Gamma) = (4^d - 3^d)^2 - (2^d - 1)^2 \sim 4^{2d} \sim N^2$. Then by the pairwise independence of the events,

$$\Delta = \frac{1}{2} \sum_{e \in \mathcal{E}} \sum_{f \sim e} p^2 = e(\Gamma) p^2 \sim \frac{(Np)^2}{2} = \frac{\mu^2}{2}.$$

$\square$

Substituting this back into Theorem 3.3.3 gives only a constant bound as $\Delta \geqslant \mu$. By considering several small dependencies, we can show that no superdependency graph can provide a useful bound.

Proposition 3.3.5. *For any superdependency graph* Γ, $e(\Gamma) = \Omega(N^2)$.

Proof. Let $e = (u, v)$ be any possible edge in $\mathcal{E}$, and let $\{u = u_1, \ldots, u_d\}$ and $\{v = v_1, \ldots, v_d\}$ be bases for U and V respectively. Let $U' = \langle u_2, \ldots, u_d \rangle$ and $V' = \langle v_2, \ldots, v_d \rangle$.

For each $u' \in U'$, $v' \in V'$, let

$$T_{u',v'} = \{(u', v'), (u + u', v'), (u', v + v'), (u + u', v + v')\}.$$

Then there must be an edge between (u, v) and $T_{u',v'}$ in Γ as if $B(x, y) = 0$ for all pairs in $T_{u',v'}$,

$$B(u, v) = B((u + u') + u', (v + v') + v') = 0,$$

and so the corresponding families of events A_{uv} and $(A_{xy})_{(x,y) \in T_{u'v'}}$ are not independent. Note that if $T_{u',v'} \cap T_{u'',v''} \neq \emptyset$ then comparing the first components gives that either $u' = u''$ or $u' + u'' = u$; the second statement is a contradiction as $u' + u'' \in U'$ but $u \notin U'$. Hence $u' = u''$, and comparing the second components shows similarly that $v' = v''$, so the sets $(T_{u'v'})_{u' \in U, v' \in V}$ are disjoint.

Considering all non-zero elements of U' and V' gives a total of $(2^{d-1})^2$ pairs, and so $d_\Gamma(e) \geqslant 2^{2d-2}$. Hence

$$2e(\Gamma) = \sum_{e \in \mathcal{E}} d_\Gamma(e) \geqslant \sum_{e \in \mathcal{E}} 2^{2d-2} \geqslant \frac{1}{4} N (2^d - 1)^2 = \frac{N^2}{4},$$

and thus $e(\Gamma) \geqslant N^2/8 = \Omega(N^2)$. $\square$

It now follows as before that $\Delta = \Omega(\mu^2)$ and so substituting this result back into Theorem 3.3.3 gives no more than a constant bound. This would appear to indicate that the dependency between the events is too high for this Janson-type approach to work.

One other inequality which can be used to derive an exponential bound is Azuma's Inequality (see for example Theorem 7.2.1 of [1]), which works if changing one of the underlying variables does not affect the desired parameter by more than a constant amount. It is clear that this is not the case in our situation; this is perhaps best seen by considering the random matrix indexed by bases $\{u_1, \dots, u_d\}$ and $\{v_1, \dots, v_d\}$ of U and V with the (i, j)-entry given by $B(u_i, v_j)$. Note that if $u = u_1 + \dots + u_k$ and $v = v_1 + \cdots v_l$ then

$$B(u, v) = 0 \text{ if and only if } \sum_{i=1}^{k} \sum_{j=1}^{l} B(u_i, v_j) = 0,$$

which corresponds to the sum of the entries in the submatrix indexed by $\{1, \dots, k\}$ and $\{1, \dots, l\}$ equalling 0. Considering all possible vectors u and v shows that X is equal to the number of submatrices whose sum is 0.

Now note that any single entry lies in $(2^{d-1})^2$ such matrices, and so changing any entry will affect the presence of roughly a quarter of the possible edges.

3.3.3 Correlation

It may now be helpful to consider a slightly more geometric approach to the structure of our model. Recall that for each $e = uv \in \mathcal{E}$ we defined a set $S_e \subseteq [M]$ such that $B(u, v) = \sum_{l \in S_e} Z_l$. From this, we can define a natural map $f_e : F^M \to F$ given by

$$f_e : (z_1, \dots, z_M) \mapsto \sum_{l \in S_e} z_l.$$

As f is clearly linear, it belongs to the dual space $(F^M)^*$.

Note that edges of the form $e = B(u_i, v_j)$ have particularly simple representations; $S_e = \{l\}$ for some $l \in [M]$ and so f_e is a restriction to the lth co-ordinate. Observing that each coefficient in the map f_e is 0 or 1, we see that each f_e is contained in the $\mathbb{F}_2$-vector space generated the M simple edges.

Now recall that $Z = (Z_1, \dots, Z_M)$; we thus have that e is present in the graph model $G_{m,r}$ if and only if

$$0 = B(u, v) = \sum_{l \in S_e} Z_l = f_e(Z_1, \dots, Z_M) = f_e(Z),$$

i.e. $Z \in \ker f_e$. In order to ease notation, we will write $H_e = \ker f_e$ for any edge $e \in \mathcal{E}$, and if $E \subset \mathcal{E}$ is a set of possible edges we will put $H_E = \cap_{e \in E} H_e$ (note H_E is always a subspace of F^M). Note that we can also write $f_e(z) = \langle v_e, z \rangle$ where v_e is defined by $(v_e)_l = 1$ if $l \in S_e$ and 0 otherwise, and $\langle \cdot, \cdot \rangle$ is the usual dot product on F^M. In this formulation, it is clear that $H_e = \{z \in F^M \mid \langle v_e, z \rangle = 0\}$ is a hyperplane through the origin, and this may be a useful way of thinking about the problem. We start with the following lemma:

Lemma 3.3.6. *Let $E \subseteq \mathcal{E}$ be a set of possible edges. Then*

$$\mathbb{P}\left(\bigcap_{e \in E} A_e\right) = p^{\operatorname{codim}(H_E)}. \tag{3.3.1}$$

Proof. Note that

$$\bigcap_{e \in E} A_e = \bigcap_{e \in E} \{Z \in \ker f_e\} = \{Z \in H_E\},$$

so $\mathbb{P}(\cap_{e \in E} A_e) = \mathbb{P}(Z \in H_E)$. As Z takes every value in F^M with the same probability we need only consider the size of H_E; suppose this has dimension d. Then

$$\mathbb{P}\left(\bigcap_{e \in E} A_e\right) = \frac{|H_E|}{|F^M|} = \frac{|F|^d}{|F|^M} = (2^r)^{d-M} = p^{M-d} = p^{\operatorname{codim}(H_E)}.$$

□

This tells us that the important parameter in this case is the dimension of the subspace H_E. We can find a lower bound on this immediately.

Corollary 3.3.7. *For any set of possible edges E, $\dim(H_E) \geqslant M - |E|$.*

Proof. We will prove this by induction on $|E|$. Note that for each e, $\operatorname{im}(f_e) \leqslant F$, a one-dimensional space; as each f_e is non-zero we have $\operatorname{rk}(f_e) = 1$. Hence $\dim(H_e) = \operatorname{null}(f_e) = M - 1$ for each edge e, giving the result for $|E| = 1$ (note that is equivalent to each H_e being a hyperplane, and that putting this into (3.3.1) recovers the result that each edge appears with probability p). So take $|E| \geqslant 1$ and consider $e \notin E$; observing that $H_E + H_e \leqslant F^M$ we obtain from the sum formula

$$\dim(H_E) + \dim(H_e) - \dim(H_E \cap H_e) = \dim(H_E + H_e) \leqslant M,$$

so rearranging

$$\dim\left(H_{E \cup \{e\}}\right) \geqslant \dim(H_E) + (M-1) - M \geqslant M - |E| - 1 = M - (|E \cup \{e\}|).$$

The result now follows by induction. □

Having obtained a precise formulation for the probabilities of sets of edges, it is somewhat natural to consider whether such values are correlated (as in $\mathcal{G}(n,p)$). Unfortunately, this turns out not always to be the case, as the following two results show.

Proposition 3.3.8. *Let E_1 and E_2 be sets of possible edges, and for $i = 1, 2$ let C_i be the event that all edges in E_i are present, i.e. $C_i = \bigcap_{e \in E_i} A_e$. Then $\mathbb{P}(C_1 \cap C_2) \geqslant \mathbb{P}(C_1)\mathbb{P}(C_2)$.*

Proof. From Lemma 3.3.6, $\mathbb{P}(C_i) = p^{\operatorname{codim}(H_{E_i})}$. Now as $H_{E_1 \cap E_2} = H_{E_1} \cap H_{E_2}$, we get from the sum formula that

$$\dim(H_{E_1}) + \dim(H_{E_2}) - \dim(H_{E_1 \cap E_2}) = \dim(H_{E_1} + H_{E_2}) \leqslant M,$$

so rearranging and adding M,

$$M - \dim(H_{E_1 \cap E_2}) \leqslant 2M - \dim(H_{E_1}) - \dim(H_{E_2}).$$

Hence

$$\mathbb{P}(C_1 \cap C_2) = p^{\operatorname{codim}(H_{E_1 \cap E_2})} \geqslant p^{\operatorname{codim}(H_{E_1})} p^{\operatorname{codim}(H_{E_2})} = \mathbb{P}(C_1)\mathbb{P}(C_2).$$

□

However, we cannot extend the correlation to *unions* of such events; the result below gives an example showing that both positive and negative correlation are possible.

Proposition 3.3.9. *Let $u \in U \backslash \{0\}$ and $v_1, \dots, v_l$ be linearly independent vectors of V ($l \geqslant 2$). Define events $C, C_1, \dots C_l$ by $C_i = A_{uv_i}$ for $i = 1, \dots, l$ and $C = A_{uv}$, where $v = \sum_{i=1}^{l} v_i$. Then C is positively correlated with $\bigcup_{i=1}^{l} C_i$ if l is odd and negatively correlated with $\bigcup_{i=1}^{l} C_i$ if l is even.*

Proof. By the assumption of linear independence, the events $C_1, \dots, C_l$ are all independent, each occurring with probability p. As C also has probability p,

$$\mathbb{P}(C)\mathbb{P}\left(\bigcup_{i=1}^{l} C_i\right) = p\left(1 - \mathbb{P}\left(\bigcap_{i=1}^{l} C_i^{\mathrm{c}}\right)\right) = p(1-(1-p)^l),$$

which can be expanded as

$$\sum_{i=1}^{l} (-1)^{i+1} \binom{l}{i} p^{i+1}.$$

On the other hand, the inclusion-exclusion formula gives

$$\begin{aligned}
\mathbb{P}\left(C\cap\bigcup_{i=1}^{l}C_i\right) &= \mathbb{P}\left(\bigcup_{1=i}^{l}(C\cap C_i)\right)\\
&= \sum_{i=1}^{l}(-1)^{i+1}\sum_{S\subseteq[l]^{(i)}}\mathbb{P}\left(\bigcap_{j\in S}(C\cap C_j)\right)\\
&= \sum_{i=1}^{l}(-1)^{i+1}\sum_{S\subseteq[l]^{(i)}}\mathbb{P}\left(C\cap\bigcap_{j\in S}C_j\right).
\end{aligned}$$

Note that for sets S with $|S|<l$, $v_1+\cdots+v_l$ is linearly independent from $\{v_j \mid j\in S\}$ and so C is independent from $(C_j)_{j\in S}$. However, the bilinearity of the map B means that if each C_i holds, so does C, so the above formula gives

$$\mathbb{P}\left(C\cap\bigcup_{i=1}^{l}C_i\right) = \sum_{i=1}^{l-1}(-1)^{i+1}\binom{l}{i}p^{i+1}+(-1)^{l+1}p^l.$$

Hence

$$\mathbb{P}\left(C\cap\bigcup_{i=1}^{l}C_i\right)-\mathbb{P}(C)\mathbb{P}\left(\bigcup_{i=1}^{l}C_i\right) = (-1)^{l+1}p^l(1-p)\begin{cases}>0 & \text{if } l+1 \text{ is even}\\ <0 & \text{if } l+1 \text{ is odd.}\end{cases}$$

This gives the result. □

3.3.4 Fixed moments

These examples, while discouraging, would appear not to be fatal; similar results hold for a random linear map as in Section 1.3.1, and yet Corollary 3.1.5 shows that an exponential bound remains true in that case. One positive result in this direction concerns the fixed moments of X, as derived below.

To prove this, we will consider the case where we have an existing set E of edges and add another edge $e\notin E$. If we can control the dimension at each change, we may be able to build sets of edges with a given dimension. The behaviour of the edges is associated with the corresponding maps f_e, as set out in the following simple lemma.

Lemma 3.3.10. *Let $E=\{e_1,\dots,e_l\}$ be a set of possible edges and $e\notin E$ another edge. To ease notation, write $f_i=f_{e_i}$ for $i=1,\dots,l$. Then*

$$\dim\left(H_{E\cup\{e\}}\right) = \begin{cases}\dim(H_E) & \textit{if } f_e\in\langle f_1,\dots,f_l\rangle\\ \dim(H_E)-1 & \textit{otherwise,}\end{cases}$$

where the span is taken in the dual space $(F^M)^*$.

Proof. Note that we can assume without loss of generality that the set $\{f_1, \ldots, f_l\}$ is linearly independent; otherwise we can reorder to find an $l' < l$ such that $\langle f_1, \ldots, f_l \rangle = \langle f_1, \ldots, f_{l'} \rangle$. But then for each i, $f_i \in \langle f_1, \ldots, f_{l'} \rangle$ so

$$\bigcap_{i=1}^{l'} H_{e_i} \subseteq \bigcap_{i=1}^{l} H_{e_i} = H_E \subseteq \bigcap_{i=1}^{l'} H_{e_i},$$

giving $H_E = \cap_{i=1}^{l'} H_{e_i}$.

Now recall that for $A \leqslant F^M$, the *annihilator* of A is given by

$$A^0 = \{f \in (F^M)^* \mid f(v) = 0 \text{ for all } v \in A\} \leqslant (F^M)^*.$$

We have that $\dim(A) + \dim(A^0) = M$ and that $(A \cap B)^0 = A^0 + B^0$ for $B \leqslant F^M$. By definition, $f_i \in H_E^0$ for each i, so by the linear independence assumption, $\dim(H_E^0) \geqslant l$. But from Corollary 3.3.7, $\dim(H_E) \geqslant M - l$; comparing with the sum formula shows that we must have equality in both cases, and in particular that $\{f_1, \ldots, f_l\}$ is a basis of H_E^0. Observing that $H_{E\cup\{e\}} = H_E \cap H_e$, we see

$$\dim(H_{E\cup\{e\}}) = M - \dim((H_E \cap H_e)^0) = M - \dim(H_E^0 + H_e^0).$$

Now by definition, $H_e^0 = \langle f_e \rangle$; if $f_e \in \langle f_1, \ldots, f_l \rangle = H_E^0$ then $H_e^0 \subseteq H_E^0$ and so $\dim(H_E^0 + H_e^0) = \dim(H_E^0) = l$. Hence $\dim(H_E \cap H_e) = M - l = \dim(H_E)$.

Otherwise, $H_E^0 \cap H_e^0 = \emptyset$ and so $\dim(H_E^0 + H_e^0) = \dim(H_E^0) + \dim(H_e^0) = l + 1$. Thus in this case, $\dim(H_E \cap H_e) = M - l - 1 = \dim(H_E) - 1$. □

This allows us to produce an estimate of the value of the fixed moments of X. While this result seems encouraging, an attempt to mimic the proof of Janson's Inequality given in [18] will only work if a similar result can be found for moments growing with μ. The proof contains ideas similar to that of Proposition 2.2.3.

Proposition 3.3.11. *For each fixed integer l, $\mathbb{E}[(X)_l] \sim \mu^l$.*

Proof. By the linearity of expectation, $\mathbb{E}[(X)_l]$ is the sum over all ordered l-tuples $(A_1, \ldots, A_l)$ of $\mathbb{P}(A_1 \cap \ldots \cap A_l)$, which in our case is the sum over all ordered sets E containing l edges of the probability that all edges in E are present. From Lemma 3.3.6, this probability depends on the dimension (or codimension) of H_E, so we will define $F_{l,k}$ to be the number of ordered sets of l edges E such that $\operatorname{codim}(H_E) = k$. Then formally

$$\mathbb{E}[(X)_l] = \sum_{k=1}^{l} F_{l,k} p^k = \left((N)_l - \sum_{k=1}^{l-1} F_{l,k} \right) p^l + \sum_{k=1}^{l-1} F_{l,k} p^k,$$

as there are $(N)_l$ l-tuples in total.

We will now try to estimate $F_{l,k}$, using Lemma 3.3.10 to keep track of the codimension as we add edges to our ordered set. In keeping with the previous result, we will work with the associated maps f_e rather than the edge e. The first map f_{e_1} we choose gives $\operatorname{codim}(H_{e_1}) = 1$. For each $i = 2, \ldots, l$, we can choose f_{e_i} in one of two ways:

1. We can take $f_{e_i} \in \langle f_{e_1}, \ldots, f_{e_{i-1}} \rangle$, in which case we have
$$\operatorname{codim}(\langle f_{e_1}, \ldots, f_{e_i} \rangle) = \operatorname{codim}(\langle f_{e_1}, \ldots, f_{e_{i-1}} \rangle);$$
2. We can take $f_{e_i} \notin \langle f_{e_1}, \ldots, f_{e_{i-1}} \rangle$, in which case we have
$$\operatorname{codim}(\langle f_{e_1}, \ldots, f_{e_i} \rangle) = 1 + \operatorname{codim}(\langle f_{e_1}, \ldots, f_{e_{i-1}} \rangle).$$

As all of the possible maps lie in an $\mathbb{F}_2$-vector space, there are at most 2^{i-1} choices in the first instance. As the total number of edges is N, there are at most $N - 2^{i-1}$ in the second. Note that as the entire set has dimension d, we need to make a choice as in item 2 exactly k times (including the first edge, which is automatically a choice as in item 2 as $\langle \emptyset \rangle = \{0\}$). Hence $F_{l,k} \leqslant C_l \cdot N^k$ for some constant C_l depending only on the fixed integer l. Hence for $k < l$,
$$\frac{F_{l,k} p^k}{\mu^l} \leqslant C_l \frac{(Np)^k}{\mu^l} = C_l \mu^{k-l} \to 0$$
as $m \to \infty$ because $\mu \to \infty$. Note also that
$$\frac{F_{l,k} p^l}{\mu^l} \leqslant C_l \frac{N^k p^l}{N^l p^l} = C_l N^{k-l} \to 0$$
as $m \to \infty$. Putting this together gives
$$\frac{\mathbb{E}[(X)_l]}{\mu^l} = \frac{(N)_l p^l}{\mu^l} - \sum_{k=1}^{l-1} \frac{F_{l,k} p^l}{\mu^l} + \sum_{k=1}^{l-1} \frac{F_{l,k} p^k}{\mu^l} \to 1$$
as $m \to \infty$, as l is fixed. □

In order to get at the intended result, we believe it will be necessary to further examine the structure of the associated maps f_e (or equivalently the corresponding hyperplanes). In the case of a random linear map, the analogous functions are very structured; they form an $\mathbb{F}_2$-vector space, and we can compute exactly the probability that the analogous random vector lies outside of all of these hyperplanes. These computations become much more difficult in a bilinear map because the corresponding functions have a complicated global structure; if we could obtain a better understanding of these we may be able to give an estimate of the corresponding probability.

Part II

Graphs of Finite Racks, Quandles and Kei

Chapter 4

Introduction to Part II

4.1 Racks and quandles

We begin this part by defining the crucial algebraic objects that we will work with throughout.

Definition 4.1.1. A *rack* is a pair $(X, \triangleright)$, where X is a non-empty set and $\triangleright : X \times X \to X$ is a binary operation such that:

1. For any $y, z \in X$, there exists $x \in X$ such that $z = x \triangleright y$;

2. Whenever we have $x, y, z \in X$ such that $x \triangleright y = z \triangleright y$, we have $x = z$;

3. For any $x, y, z \in X$, $(x \triangleright y) \triangleright z = (x \triangleright z) \triangleright (y \triangleright z)$.

Definition 4.1.2. A *quandle* is a rack $(X, \triangleright)$ satisfying the additional condition that $x \triangleright x = x$ for every $x \in X$.

Definition 4.1.3. Let $(X, \triangleright)$ be a rack. If X is finite, we call $|X|$ the *order* of the rack.

Note that conditions 1 and 2 of Definition 4.1.1 are equivalent to the statement that for each y, the map $x \mapsto x \triangleright y$ is a bijection on X.

As mentioned by Blackburn in [2], racks originally developed from correspondence between J.H. Conway and G.C. Wraith in 1959, while quandles were introduced independently by Joyce [21] and Matveev [24] in 1982 as invariants of knots. Fenn and Rourke [8] provide a history of these concepts, while Nelson [26] gives an overview of how these structures relate to other areas of mathematics.

As a first example, note that for any set X, if we define $x \triangleright y = x$ for all $x, y \in X$, we obtain a rack, known as the *trivial rack* T_X. Let us consider two more interesting classical examples of quandles, which (at least in a combinatorial sense) are more widely studied than racks (see [11], [5]).

Definition 4.1.4. Let G be a group and define a binary operation $\triangleright : G \times G \to G$ by $x \triangleright y := y^{-1}xy$. The resulting quandle $(G, \triangleright)$ is known as a *conjugation quandle.*

It is straightforward to verify that $(G, \triangleright)$ is in fact a quandle; for example, item 3 of Definition 4.1.1 holds as

$$\begin{aligned}
(x \triangleright z) \triangleright (y \triangleright z) &= (y \triangleright z)^{-1}(x \triangleright z)(y \triangleright z) \\
&= (z^{-1}yz)^{-1}(z^{-1}xz)(z^{-1}yz) \\
&= z^{-1}(y^{-1}xy)z \\
&= (x \triangleright y) \triangleright z.
\end{aligned}$$

For the next definition, as is standard in algebra, we will write the operation of an abelian group additively. As stated in the introduction, we write homomorphisms and group actions on the *right.*

Definition 4.1.5. Let A be an Abelian group and $\tau \in \mathrm{Aut}(A)$ be an automorphism. Define a binary operation $\triangleright : A \times A \to A$ by $x \triangleright y = (x)\tau + y - (y)\tau = (x - y)\tau + y$. Then $(G, \triangleright) = (A, \tau)$ is an *Alexander quandle* or *affine quandle.*

Again, it is straightforward to check that $(G, \triangleright)$ is a quandle; for example, note that

$$(x \triangleright y) \triangleright z = ((x - y)\tau + y) \triangleright z = ((x - y)\tau + y - z)\tau + z,$$

while on the other hand, as τ is an automorphism,

$$\begin{aligned}
(x \triangleright z) \triangleright (y \triangleright z) &= ((x - y)\tau + z) \triangleright ((y - z)\tau + z) \\
&= ((x - z)\tau + z - (y - z)\tau - z)\tau + (y - z)\tau + z \\
&= ((x - y)\tau + y - z)\tau + z,
\end{aligned}$$

and so item 3 of Definition 4.1.1 holds.

As is standard with algebraic objects, we now define subracks, subquandles and notions of isomorphism.

Definition 4.1.6. Let $(X, \triangleright)$ be a rack (or quandle). Then a *subrack* (or *subquandle*) of $(X, \triangleright)$ is a rack (or quandle) $(Y, \triangleright|_{Y \times Y})$ where $Y \subseteq X$.

The notation $\triangleright|_{Y \times Y}$ in the above context will always be abbreviated to $\triangleright$, with the restriction to the subset Y left implicit.

$\triangleright$	ι	(1 2)	(1 3)	(2 3)	(1 2 3)	(1 3 2)
ι	ι	ι	ι	ι	ι	ι
(1 2)	(1 2)	(1 2)	(2 3)	(1 3)	(2 3)	(1 3)
(1 3)	(1 3)	(2 3)	(1 3)	(1 2)	(1 2)	(2 3)
(2 3)	(2 3)	(1 3)	(1 2)	(2 3)	(1 3)	(1 2)
(1 2 3)	(1 2 3)	(1 3 2)	(1 3 2)	(1 3 2)	(1 2 3)	(1 2 3)
(1 3 2)	(1 3 2)	(1 2 3)	(1 2 3)	(1 2 3)	(1 3 2)	(1 3 2)

Table 4.1: The conjugation quandle on S_3.

Definition 4.1.7. Let $(X, \triangleright)$ and $(X', \triangleright')$ be racks. A map $\phi : X \to X'$ is a *rack homomorphism* if $(x \triangleright y)\phi = (x)\phi \triangleright' (y)\phi$ for all $x, y \in X$. A bijective homomorphism is an *isomorphism*, and an isomorphism from a rack $(X, \triangleright)$ to itself is an *automorphism*.

We will usually only be concerned with racks and quandles up to isomorphism. If $(X, \triangleright)$ is a rack of order n, then it is clearly isomorphic to a rack on $[n]$, so we will almost always take $[n]$ to be our underlying ground set. We will denote the set of all racks on $[n]$ by $\mathcal{R}_n$.

Let us demonstrate some of these principles with a concrete example, namely the conjugation quandle on S_3. As with groups, we can express $\triangleright$ completely via a table; here, the (i, j) entry is equal to $i \triangleright j$. Notice that the sets $\{(1\ 2), (1\ 3), (2\ 3)\}$, $\{(1\ 2\ 3), (1\ 3\ 2)\}$ and $\{\iota\}$ all form subquandles under $\triangleright$, as does any union of these sets. In a general conjugation quandle, the set of elements of any given cycle type will form a subquandle (as conjugation preserves cycle type). Now if we identify the six elements of S_3 with the set $\{1, \ldots, 6\}$ (in the order given in Table 4.1), we obtain the following quandle on $\{1, \ldots, 6\}$, where we have omitted the labels and written the table as a matrix:

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 2 & 2 & 4 & 3 & 4 & 3 \\ 3 & 4 & 3 & 2 & 2 & 4 \\ 4 & 3 & 2 & 4 & 3 & 2 \\ 5 & 6 & 6 & 6 & 5 & 5 \\ 6 & 5 & 5 & 5 & 6 & 6 \end{pmatrix}$$

Now consider the following question: for a given set X, when does a matrix or table like the above actually define a rack (or quandle)? Any such matrix M with entries from X can be viewed as a set of maps $(f_y)_{y \in X}$ given by the columns, i.e.

$(x)f_y = M_{xy}$. We would like to define the operation $\triangleright$ by setting $x \triangleright y = (x)f_y = M_{xy}$; the following well-known result (for example, [8], [2]) gives the correct conditions for a collection of maps to define a rack.

Proposition 4.1.8. *Let X be a set and $(f_x)_{x\in X}$ be a collection of functions each with domain and co-domain X. Define a binary operation $\triangleright : X \times X \to X$ by $x \triangleright y := (x)f_y$. Then $(X, \triangleright)$ is a rack if and only if f_y is a bijection for each $y \in X$ and for all $y, z \in X$ we have*

$$f_{(y)f_z} = f_z^{-1} f_y f_z. \tag{4.1.1}$$

Further, $(X, \triangleright)$ is a quandle if in addition $(x)f_x = x$ for all $x \in X$.

Proof. As noted earlier, each f_y is a bijection, so it remains to show that item 3 in Definition 4.1.1 is equivalent to (4.1.1). Fix arbitrary elements $x, y, z \in X$; then $(x \triangleright y) \triangleright z = (x)f_y f_z$, while $(x \triangleright z) \triangleright (y \triangleright z) = (x)f_z f_{(y)f_z}$, so (as x was arbitrary) we have equality if and only if $f_y f_z = f_z f_{(y)f_z}$. This proves the result for racks; the result for quandles follows easily from Definition 4.1.2. $\square$

This means that we can just as well define a rack on a set X by the set of maps $(f_y)_{y\in X}$, providing they are all bijections and satisfy (4.1.1). We will move freely between the two definitions, with $x \triangleright y = (x)f_y$ for all $x, y \in X$ unless otherwise stated.

The following corollary to Proposition 4.1.8 is immediate but will be used many times.

Corollary 4.1.9. *Let $(X, \triangleright)$ be a rack and suppose $(x)f_y = x$ for some $x, y \in X$. Then f_x and f_y commute in* Sym(X). $\square$

We will now show how to extend Proposition 4.1.8 to elements of the group generated by $(f_y)_{y\in X}$, known as the *operator group* of the rack. This is again a standard result (see [8], [2]).

Lemma 4.1.10. *Let $(X, \triangleright)$ be a rack and let G be its operator group. Then for any $y \in X$ and $g \in G$, $f_{(y)g} = g^{-1} f_y g$.*

Proof. Let $g \in G$, so $g = f_{i_1}^{\epsilon_1} \cdots f_{i_m}^{\epsilon_m}$, where m is a positive integer, $i_j \in X$ for $1 \leqslant j \leqslant m$ and $\epsilon_j \in \{-1, 1\}$ for $1 \leqslant j \leqslant m$. We will prove the result by induction on m.

First suppose $m = 1$ and write $i_1 = z$. If $\epsilon_1 = 1$ then $g = f_z$ and the result is true from Proposition 4.1.8, so suppose $\epsilon_1 = -1$, i.e. $g = f_z^{-1}$. Write $x = (y)f_z^{-1}$; from

(4.1.1) we have $f_{(x)f_z} = f_z^{-1} f_x f_z$, i.e. $f_y = g f_{(y)g} g^{-1}$. This shows the result for $m = 1$, so take $m > 1$ and assume the result for smaller m.

Now write $g_1 = f_{i_1}^{\epsilon_1} \cdots f_{i_{m-1}}^{\epsilon_{m-1}}$ and $g_2 = f_{i_m}^{\epsilon_m}$, so by the inductive hypothesis the result is true for both g_1 and g_2. Applying the result to g_2 with $(y)g_1$ in place of y gives $f_{(y)g} = f_{(y)g_1 g_2} = g_2^{-1} f_{(y)g_1} g_2$, and then we can apply the result to g_1 to get $f_{(y)g} = g_2^{-1} g_1^{-1} f_y g_1 g_2 = g^{-1} f_y g$. Thus the result follows by induction. □

We have now established some basic working results about racks and quandles. In the next section we will see that a rack is closely related to its operator group.

4.2 Structure and an upper bound

Recall that $\mathcal{R}_n$ denotes the set of all racks on $[n]$. In this section we will give an upper bound on $|\mathcal{R}_n|$, the number of racks on $[n]$; we adapt a method of Blackburn used in [2] to give a stronger result. The following two results of [2] show how a rack is related to its operator group; the proofs are quite short and are thus included here. Recall that for a group G and an element $\pi \in G$, the *centraliser* of π in G is the subgroup

$$C_G(\pi) = \{g \in G \mid g^{-1} \pi g = \pi\}.$$

If $G \leqslant \mathrm{Sym}(X)$ then there is a natural action of G on X; for any $x \in X$ and $g \in G$, the pair (x, g) is mapped to the element $(x)g$[1]. Then for $\alpha \in X$ the *stabiliser* of α under G is the subgroup

$$G_\alpha = \{g \in G \mid (\alpha)g = \alpha\}.$$

Theorem 4.2.1. *Let X be a set and $G \leqslant \mathrm{Sym}(X)$. Let I index the set of orbits of G, and let $\{\alpha_i \mid i \in I\}$ be a complete set of orbit representatives. Let $(\pi_i)_{i \in I}$ be a sequence of elements from G such that*

$$G_{\alpha_i} \leqslant C_G(\pi_i) \tag{4.2.1}$$

for each $i \in I$. Define, for each $i \in I$ and $g \in G$, the map $f_{(\alpha_i)g} = g^{-1} \pi_i g \in \mathrm{Sym}(X)$. Then this set of bijections satisfies (4.1.1)*, and thus defines a rack on X, with operator group contained in G.*

Proof. We must first check that the maps are well-defined; suppose $x = (\alpha_i)g = (\alpha_i)h$ for some i (indexing the orbit containing x) and some $g, h \in G$. Then $(\alpha_i)gh^{-1} = \alpha_i$,

[1] We use 'orbits of G' as shorthand for 'orbits of G under this natural action'.

so $gh^{-1} \in G_{\alpha_i} \leqslant C_G(\pi_i)$. Thus $(hg^{-1})\pi_i(gh^{-1}) = \pi_i$, and so $g^{-1}\pi_i g = h^{-1}\pi_i h$, i.e. $f_{(\alpha)g} = f_{(\alpha)h}$. Thus the maps are well-defined.

Now let $y, z \in X$, so there exist $i, j \in I$, $g, h \in G$, such that $y = (\alpha_i)g$ and $z = (\alpha_j)h$. Then $(y)f_z = ((\alpha_i)g)h^{-1}\pi_j h$, and so

$$f_{(y)f_z} = (gh^{-1}\pi_j h)^{-1}\pi_i(gh^{-1}\pi_j h) = (h^{-1}\pi_j^{-1}h)(g^{-1}\pi_i g)(h^{-1}\pi_j h) = f_z^{-1} f_y f_z.$$

This shows that we have defined a rack. Finally, the operator group is exactly $\langle g^{-1}\pi_i g \mid i \in I, g \in G\rangle$, which by definition is contained in G. □

This result shows that given a group G and a complete set of orbit representatives $(\alpha_i)_{i\in I}$, we can define a rack just by defining f_{α_i} for each i, provided (4.2.1) is satisfied. The next result shows that every rack is constructed in this way.

Theorem 4.2.2. *Let $(X, \triangleright)$ be a rack with associated maps $(f_y)_{y\in X}$ and operator group G. Let I index the set of orbits of G, and let $\{\alpha_i \mid i \in I\}$ be a complete set of orbit representatives. Define $\pi_i = f_{\alpha_i}$ for each $i \in I$. Then $G_{\alpha_i} \leqslant C_G(\pi_i)$ for each i, and $f_{(\alpha_i)g} = g^{-1}\pi_i g$ for all $g \in G$ and $i \in I$.*

Proof. Applying Lemma 4.1.10 to $\alpha_i \in X$ gives that $f_{(\alpha_i)g} = g^{-1}f_{\alpha_i}g = g^{-1}\pi_i g$ for any $g \in G$. Now fix $i \in I$ and let $g \in G_{\alpha_i}$; then $g^{-1}\pi_i g = f_{(\alpha_i)g} = f_{\alpha_i} = \pi_i$, and so $G_{\alpha_i} \leqslant C_G(\pi_i)$. This proves the theorem. □

It is this last result that allows us to prove the upper bound. We will first need two enumerative results from group theory; the first is due to Pyber [28].

Theorem 4.2.3. *Let X be a set of order n. Then the number of subgroups of* $\mathrm{Sym}(X)$ *is at most $24^{(1/6+o(1))n^2}$.* □

The second is due to Maróti [23].

Theorem 4.2.4. *Let X be a set of order $n > 2$ and $G \leqslant \mathrm{Sym}(X)$. Then the number of conjugacy classes of G is at most $3^{(n-1)/2}$.* □

We can now show the result; the bound given is an improvement on Blackburn's result, which has the constant equal to $(\log_2 24)/6 + (\log_2 3)/2 \approx 1.5566$. We obtain the improvement by considering G-orbits of size one separately.

Theorem 4.2.5. $|\mathcal{R}_n| \leqslant 2^{(c+o(1))n^2}$, *where $c = (\log_2 24)/6 + (\log_2 3)/4 \approx 1.1604$.*

Proof. From Theorem 4.2.2, any rack on $X = [n]$ is determined by the operator group G and the elements $\pi_i = f_{\alpha_i}$, where $\{\alpha_i \mid i \in I\}$ is a complete set of orbit representatives for G. We will bound the number of choices for the group and the maps, considering separately the cases where G has r orbits of size 1, for $r = 1, \ldots, n$. From Theorem 4.2.3, there are at most $24^{(1/4+o(1))n^2}$ possibilities for the operator group G.

First take such a group $G \leqslant \mathfrak{S}_n$ and let X_0 be the set of elements of $[n]$ contained in G-orbits of size one, i.e.

$$X_0 = \{x \in [n] \mid (x)g = x \text{ for all } g \in G\}.$$

Write $r = |X_0|$ and put $X_1 = [n] \setminus X_0$; then each element of X_1 is contained in an orbit of size at least two, and so there are at most $|X_1|/2 = (n-r)/2$ non-trivial orbits. Let s be the total number of orbits, so $s \leqslant r + (n-r)/2 = (n+r)/2$. Note also that as each element of G fixes each element of X_0, we can regard G as a subgroup of $\mathrm{Sym}(X_1)$, and so we can replace n with $n - r$ when using Theorem 4.2.4.

Now G has s orbits $O_1, \ldots, O_s$; write $n_j = |O_j|$ for each j, where $n_j = 1$ for $1 \leqslant j \leqslant r$. Let $\alpha_1, \ldots, \alpha_s$ be a corresponding complete set of orbit representatives; we must now choose the elements $\pi_1, \ldots, \pi_s \in G$ subject to the condition that $G_{\alpha_i} \leqslant C_G(\pi_i)$ for each i. From the Orbit-Stabiliser Theorem we have

$$\frac{|G|}{n_i} = |G_{\alpha_i}| \leqslant |C_G(\pi_i)| = \frac{|G|}{|\mathrm{ccl}_G(\pi_i)|}$$

for all i, and so each π_i lies in a conjugacy class C_i of size at most n_i. Now G has at most $3^{(n-r)/2}$ conjugacy classes, and so there are at most $3^{(n-r)s/2} \leqslant 3^{(n-r)(n+r)/4}$ choices for the conjugacy classes $C_1, \ldots, C_s$. Within each C_i there are at most n_i choices for π_i, and these choices determine the rack entirely. Thus the number of racks on $[n]$ with operator group contained in G is at most

$$3^{(n^2-r^2)/4} \cdot \prod_{i=1}^{s} n_i \leqslant 3^{n^2/4} \cdot \prod_{i=1}^{s} n_i.$$

Consider this last product, and recall that $\sum_{i=1}^{s} n_i = n$. Suppose we had some $n_i \geqslant 4$; then if we replaced the ith orbit with two orbits, one of size 2 and one of size $n_i - 2$, the sum of orbit sizes is still n but the product has increased by a factor of $2(n_i - 2)/n_i = 2 - 4/n_i \geqslant 1$. Hence the product is maximised for some sequence in which every orbit has size at most 3, and thus $\prod_{i=1}^{s} n_i \leqslant 3^s \leqslant 3^n$.

It follows that the total number of racks on $[n]$ is at most

$$24^{(1/6+o(1))n^2} \cdot 3^{n^2/4} \cdot 3^n = 2^{(c+o(1))n^2},$$

where $c = (\log_2 24)/6 + (\log_2 3)/4 \approx 1.1604$. □

4.3 Further definitions and graphical representations

We now give some further definitions. The terminology used in the next two definitions is not standard.

Definition 4.3.1. Let $(X, \triangleright)$ be a rack. A subrack $(Y, \triangleright)$ is said to be *stable* if for all $y \in Y$, $z \in X \setminus Y$, $(y)f_z = y$ and $(z)f_y = z$.

Definition 4.3.2. A rack $(X, \triangleright)$ is *reducible* if X can be partitioned into two non-empty sets A and B such that both $(A, \triangleright)$ and $(B, \triangleright)$ are stable subracks. A rack is *irreducible* if it is not reducible.

Definition 4.3.3. A rack $(X, \triangleright)$ is *connected* (see [25]), or *indecomposable* (see [11]), if its operator group G acts transitively on X, i.e. if for all $y, z \in X$ there exists a sequence $(i_1, \ldots, i_m)$ of elements from X and a sequence $(\epsilon_1, \ldots, \epsilon_m) \in \{-1, 1\}^m$ such that $z = (y)f_{i_1}^{\epsilon_1} \cdots f_{i_m}^{\epsilon_m}$.

The following result can be seen as the reverse of Definition 4.3.2, constructing a larger rack (or quandle) from two smaller ones. Again the terminology is non-standard.

Proposition 4.3.4. *Let A and B be disjoint sets and $R = (A, \triangleright_A)$ and $S = (B, \triangleright_B)$ be racks. Then there is a rack $R \oplus S = (A \cup B, \triangleright)$ such that $(A, \triangleright)$ and $(B, \triangleright)$ are stable subracks, isomorphic to R and S respectively. Further, if R and S are quandles, so is $R \oplus S$.*

Proof. Let $(g_a)_{a \in A}$ and $(g_b)_{b \in B}$ be the associated maps of the two racks. For any $a \in A$ and $b \in B$ define bijections f_a and f_b on the whole of $A \cup B$ as follows: put $f_a|_A = g_a$, $f_a|_B = \iota$, $f_b|_A = \iota$ and $f_b|_B = g_b$. It is now straightforward to check (4.1.1) in all cases; note also if R and S are quandles then $(x)f_x = x$ for all $x \in A \cup B$. This shows that $(A \cup B, \triangleright)$ (defined with the maps above) is a rack (or quandle), and the remaining statements are true by construction. □

Now observe that any rack on X can be represented by a directed multigraph on X; we give each vertex a colour and then put a directed edge of colour i from vertex j to vertex k if and only if $(j)f_i = k$. We then remove all loops from the graph; i.e. if $(j)f_i = j$ we don't have an edge of colour i incident to j.

It will be helpful to recast the representation of racks by directed multigraphs in a slightly different setting. Let V be a finite set and let $\sigma \in \mathrm{Sym}(V)$; then we can define a simple, loopless directed graph G_σ on V by setting $\overrightarrow{uv} \in \overrightarrow{E}(G_\sigma)$ if and only if $u \neq v$ and $(u)\sigma = v$. By considering the decomposition of σ into disjoint cycles, we see that G_σ consists of a collection of disjoint directed cycles (some of which may be degenerate 'cycles' of length 2) and some isolated vertices. We can now extend this definition to the case of multiple permutations in a natural way.

Definition 4.3.5. Suppose $\Sigma = \{\sigma_1, \dots, \sigma_k\} \subseteq \mathrm{Sym}(V)$ is a set of permutations on a set V. Define a directed, loopless multigraph $G_\Sigma = (V, \overrightarrow{E})$ with a k-edge-colouring by putting a directed edge of colour i from u to v if and only if $u \neq v$ and $(u)\sigma_i = v$.

We also define the *reduced* graph G^0_Σ to be the directed graph on V obtained by setting $e = \overrightarrow{uv} \in \overrightarrow{E}(G^0_\Sigma)$ if and only if there is a directed edge from u to v in G_Σ.

Note the edge set of the reduced graph contains at most two edges between any $u, v \in V$, namely $\overrightarrow{uv}$ and $\overrightarrow{vu}$. Also observe that if $\Sigma' \subseteq \Sigma$, then $G_{\Sigma'}$ is a subgraph of G_Σ.

Now let us return specifically to racks.

Definition 4.3.6. Let $R = (X, \triangleright)$ be a rack, and let $(f_y)_{y \in X}$ be the associated maps. For any $S \subseteq X$, define $\Sigma_S = \{f_y \mid y \in S\}$. Then by G_S we mean the directed multigraph G_{Σ_S} in the sense of Definition 4.3.5; G_S thus has an associated $|S|$-edge-colouring, although if $|S| = 1$ we may not necessarily consider G_S as being coloured. We will also write $G_R = G_X$, indicating the graph for the whole rack.

The following two observations are straightforward but crucial.

Lemma 4.3.7. *Let V be a finite set and let $\Sigma \subseteq \mathrm{Sym}(V)$ be a family of permutations. Let $u, v \in V$ be distinct; then there is a uv-path in G_Σ if and only if there is a directed uv-path in G_Σ.*

Proof. We need only prove the 'only if' statement. We may thus assume that there is a uv-path in G_Σ, i.e. that the graph distance $d(u, v) < \infty$. We will prove the result by induction on $d = d(u, v)$.

First suppose that $d = 1$, so that there is an edge of some colour $\sigma \in \Sigma$ between u and v. If $(u)\sigma = v$ then we already have a directed path, so suppose that $(v)\sigma = u$. Then, considering the cycle decomposition of σ, there is a directed cycle $v = v_0, u = v_1, \dots, v_r, v_0$ in G_σ, and so $u = v_1, \dots, v_r, v_0 = v$ is a directed uv-path. This proves the base case, so take $d > 1$ and assume the result for $d - 1$.

Then there is a path $u = u_0, \dots, u_{d-1}, u_d = v$ of length d in G_Σ; by the inductive hypothesis there is a directed path between u and u_{d-1}. $d(u_{d-1}, v) = 1$, so from the above argument, there is a directed path from u_{d-1} to $u_d = v$; concatenating these paths gives a directed uv-walk. Hence there exists a directed uv-path and the result follows by induction. $\square$

Lemma 4.3.8. *Let $\Sigma \subseteq \mathrm{Sym}(V)$ be a family of permutations and $U \subseteq V$. Then U is an orbit of the natural action of $\langle\Sigma\rangle$ on V if and only if U spans a component in G_Σ.*

Proof. Let $u, v \in V$. Then u and v are in the same orbit of the natural action if and only if there exists a sequence $(\sigma_{i_1}, \ldots, \sigma_{i_m})$ of elements from Σ and a sequence $(\epsilon_1, \ldots, \epsilon_m) \in \{-1, 1\}^m$ such that $v = (u)\sigma_{i_1}^{\epsilon_1} \cdots \sigma_{i_m}^{\epsilon_m}$. But this is exactly equivalent to there being a uv-path in G_Σ with edges successively coloured $i_1, \ldots, i_m$, and the value of ϵ_i indicating the direction of the edge. Thus the partition of V into orbits of $\langle\Sigma\rangle$ coincides with the partition of V into components of G_Σ. $\square$

We can use this last result to recast some of the earlier definitions using graphical language. As a transitive group action has only one orbit, we see that a rack $R = (X, \triangleright)$ is connected (in the sense of Definition 4.3.3) if and only if G_R is connected in the graphical sense. A subset $Y \subseteq X$ forms a subrack if and only if for all $y, z \in Y$, $(z)f_y \in Y$; as each f_y is a bijection we also have that $(z)f_y^{-1} \in Y$ for all z and thus Y and $X \setminus Y$ are separated in the graph G_Y. From Definition 4.3.1, a subrack $(Y, \triangleright)$ is stable if and only if we also have that $X \setminus Y$ is an independent set in G_Y. Thus K is reducible if and only if we can partition X into non-empty sets A and B such that each edge of G_A joins two elements of A while each edge of G_B joins two elements of B, and Proposition 4.3.4 states that $K \oplus L$ is represented by the graph $G_K \cup G_L$.

4.4 Kei

We now introduce a more restricted structure than a quandle, which will be of interest throughout.

Definition 4.4.1. A quandle $(X, \triangleright)$ is a *kei* (or *involutive quandle*) if $(x \triangleright y) \triangleright y = x$ for any $x, y \in X$, or equivalently if $f_y^2 = \iota$ for all $y \in X$.

Kei were first studied by Takasaki [32] in 1943; a recent paper by Stanovský [30] gives a thorough survey of the history of research on these structures.

As an example, observe that the conjugation quandle on S_3 (given in Table 4.1) is not a kei, but the subquandles on the sets $\{(1\ 2), (1\ 3), (2\ 3)\}$ and $\{\iota, (1\ 2), (1\ 3), (2\ 3)\}$ are. Because of the condition that any map f_y must have order at most 2, we see that a subquandle $(Y, \triangleright)$ of a conjugation quandle is a kei if all elements of Y have

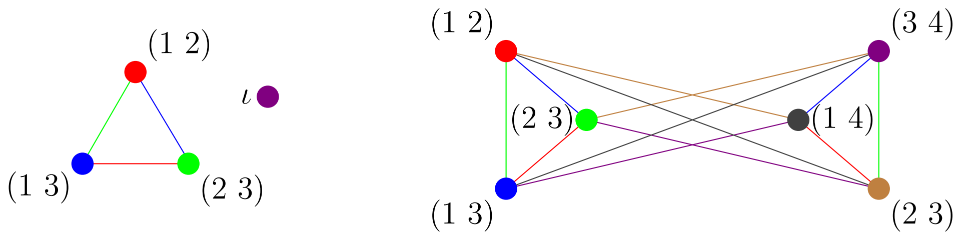

Figure 4.1: Graphical representations of two kei. Both of these are subquandles of conjugation quandles; S_3 on the left and S_4 on the right.

order at most 2 (the converse of this is not true; for example any singleton set is a kei under the inherited map $\triangleright$).

One advantage of working with a kei $K = (X, \triangleright)$ is the simple form of G_K. Each map f_i can be written as a disjoint union of transpositions, and thus the graph G_i consists only of isolated vertices and degenerate 'cycles' of length 2. We can thus replace each such cycle with a single undirected edge; observe that in G_K, each vertex is incident to at most one edge of each colour. We also note that the reduced graph G_K^0 is similarly undirected. Figure 4.1 gives two examples of undirected multigraphs representing kei.

The remainder of this part of the thesis will study racks, quandles and kei via their graphical representations, a departure from the papers discussed in [30], which rely heavily on algebra. Little algebra will be used here, with graph-theoretical concepts dominating. In Chapter 5, we investigate lower bounds for the number of kei on $[n]$ by expanding and generalising an argument of Blackburn in [2]; this provides some justification for why a lower bound of $2^{(1/4+o(1))n^2}$ seems feasible, while results about commuting involutions in 5.1 will prove useful throughout. In Chapter 6 we give a full classification (and subsequent enumeration) of kei in which each orbit has size p, where p is an odd prime; the kei in this case show markedly different behaviour to the case where $p = 2$. In Chapter 7 we prove an upper bound of $2^{(1/4+o(1))n^2}$ for the number of (isomorphism classes of) racks of order n, as well as showing that almost all racks are similar to kei constructed in Chapter 5. In Chapter 8 we discuss some issues surrounding connected racks, and we finish with an appendix including various miscellaneous results.

Chapter 5

Lower bounds

Theorem 4.2.5 is an improvement of Blackburn's upper bound on the number of isomorphism classes of racks of order n; the same paper [2] also gives a lower bound of $2^{(1/4+o(1))n^2}$ for the number of isomorphism classes of racks of order n. The bound is obtained by considering a partition $\mathcal{P}$ of the ground set, say $[n]$, into sets of size two, and considering only kei for which each map preserves $\mathcal{P}$; it transpires that in this case the maps f_x and f_y commute when x and y are in the same part of $\mathcal{P}$ and thus the number of choices for such a kei is large. In this chapter we will generalise this construction to different partitions of the ground set.

In Section 5.1, we examine the structure of the multigraph G_Σ in the case where Σ is a family of mutually commuting involutions, attempting to replicate the key property as described above. We show that in this case each component must have size a power of 2; in Section 5.2 we consider partitions of this form and show that the lower bound $2^{(1/4+o(1))n^2}$ remains true and is best possible. This provides some intuition for the correctness of the bound overall; it will be shown in Chapter 7 that the number of isomorphism classes of racks of order n is equal to $2^{(1/4+o(1))n^2/4}$.

5.1 Mutually commuting involutions

5.1.1 A Permutation Theory Lemma

We will first need a lemma concerning permutations. Let m be a non-negative integer and $S \subseteq [m]$. For any $v \in \{0,1\}^m$, define v_S to be the element of $\{0,1\}^m$ obtained by changing only the coordinates in S, i.e.

$$\{i \in [m] \mid (v_S)_i \neq v_i\} = S.$$

By considering which coordinates change at each step, it is straightforward to see that for any S, T, $(v_S)_T = v_{S\triangle T} = (v_T)_S$, where $S\triangle T = (S \cup T) \setminus (S \cap T)$ is the set difference. In particular, $(v_S)_S = v_\emptyset = v$ for any S; this means that any map $f : \{0,1\}^m \to \{0,1\}^m$ of the form $v \mapsto v_S$ is an involution with no fixed points (unless $S = \emptyset$). We will prove the following lemma.

Lemma 5.1.1. *Define involutions $\sigma_1, \ldots, \sigma_m$ of $\{0,1\}^m$ by $(v)\sigma_i = v_{\{i\}}$ for all i and v. Now let σ be another involution such that $\sigma\sigma_i = \sigma_i\sigma$ for $i = 1, \ldots, m$. Then there exists some $S \subseteq [m]$ such that $(v)\sigma = v_S$ for all v.*

Proof. Choose an arbitrary v_0 and write $u_0 = (v_0)\sigma$; let $S := \{i \in [m] \mid (u_0)_i \neq (v_0)_i\}$. Now take an arbitrary $v \in \{0,1\}^m$; we will prove that $(v)\sigma = v_S$ by induction on the Hamming distance $d = d(v, v_0)$. The base case $d = 0$ has just been established, so take $d > 0$ and suppose the result is true for smaller d.

Now the set of coordinates on which v and v_0 differ is non-empty, so choose such a coordinate i and let $u = v_{\{i\}}$. Then $d(u, v_0) = d - 1$, so by the inductive hypothesis,

$$(u)\sigma = u_S = (v_{\{i\}})_S = v_{S\triangle\{i\}}. \tag{5.1.1}$$

Now by our assumptions, $(v)\sigma = (v)\sigma_i\sigma\sigma_i$, and also $(v)\sigma_i = v_{\{i\}} = u$. Thus

$$(v)\sigma = (u)\sigma\sigma_i = (v_{S\triangle\{i\}})\sigma_i = v_{S\triangle\{i\}\triangle\{i\}} = v_S$$

where we have applied (5.1.1). The result now follows by induction. $\square$

5.1.2 Graphical representations of commuting families

Let $\Sigma = \{\sigma_1, \ldots, \sigma_k\}$ be a set of involutions from $\mathrm{Sym}(V)$. We would like to know what happens when we add a new involution $\sigma = \sigma_{k+1}$ to Σ such that σ commutes with each σ_i; how does the graph $G_{\Sigma\cup\{\sigma\}}$ relate to G_Σ? An important example concerns the cubes Q_m and is related to Section 5.1.1.

Write $V = \{0,1\}^m$ and let Σ, σ and G_Σ be as above; to ease notation, we will write $G_i := G_{\sigma_i}$ for each i in this setting. Now suppose the conditions of Lemma 5.1.1 are met, i.e. $(v)\sigma_i = v_{\{i\}}$ for $i = 1, \ldots, m$. Then the edges of G_i are exactly the pairs uv where u and v differ in the ith coordinate, so the reduced graph G^0_Σ contains all edges of the form

$$\{uv \mid d(u,v) = 1\},$$

i.e. G^0_Σ contains a subgraph isomorphic to the cube Q_m. Lemma 5.1.1 says that in the case where $\sigma = \sigma_{k+1}$ commutes with all the existing involutions, the edges of G_{k+1} are of the form vv_S for some fixed $S \subseteq [m]$ and all $v \in \{0,1\}^m$.

We now consider the case where the graph G_Σ may have multiple components. If G and H are edge-coloured multigraphs, we call an isomorphism $\phi : G \to H$ *colour-preserving* if the edges xy and $(x)\phi(y)\phi$ have the same colour for all $x, y \in V(G)$.

Lemma 5.1.2. *Let $\Sigma = \{\sigma_1, \ldots, \sigma_k\} \subseteq \mathrm{Sym}(V)$ be a set of involutions, and let $\sigma \in \mathrm{Sym}(V)$ be an involution such that $\sigma\sigma_i = \sigma_i\sigma$ for $i = 1, \ldots, k$. Suppose that $C, D \subseteq V$ span distinct components of G_Σ, and there exist $v_0 \in C$ and $w_0 \in D$ such that $(v_0)\sigma = w_0$. Then $\sigma|_C$ is a colour-preserving isomorphism between $G_\Sigma[C]$ and $G_\Sigma[D]$, and further, $G_\sigma[C \cup D]$ is a complete matching from C to D.*

Proof. Define $S \subseteq C$ by
$$S := \{v \in C \mid (v)\sigma \in D\},$$
so by assumption $v_0 \in S$. Write $T = (S)\sigma \subseteq D$. Now suppose that $S \neq C$; as C spans a component of G_Σ there exist $v \in S$ and $u \in C \setminus S$ such that u and v are joined by an edge of some colour $i \leqslant k$, i.e. $(v)\sigma_i = u$. Now as $v \in S$, $(v)\sigma \in D$, and as D is a component of G_Σ, $(v)\sigma\sigma_i \in D$. As σ and σ_i commute, $D \ni (v)\sigma_i\sigma = (u)\sigma$, and so by definition $u \in S$, which is a contradiction. Hence $S = C$. Similarly, by considering σ on T, we obtain $T = D$, and so $(C)\sigma = D$.

Now suppose $u, v \in C$ are joined by an edge of colour $i \leqslant k$ (so as before, $(u)\sigma_i = v$). Then $(v)\sigma = (u)\sigma_i\sigma = (u)\sigma\sigma_i$, i.e. $sigma_i$ maps $(u)\sigma$ onto $(v)\sigma$ and thus $(u)\sigma$ and $(v)\sigma$ are joined by an edge of colour i. By considering σ on D, we obtain the reverse implication, which is to say that $\sigma|_C$ is a colour-preserving isomorphism.

The last statement follows since G_σ is a matching and each vertex in C has a neighbour in D. $\square$

We can now use these two lemmas to determine exactly what form the coloured multigraph G_Σ must have in the case where all the involutions in Σ commute. All of the components are (multi)cubes with some additional edges; we will define a set of abstract edge-coloured multigraphs and work with isomorphisms.

Definition 5.1.3. Let k and m be non-negative integers and $\mathcal{F} = (F_1, \ldots, F_k)$ be a sequence in $\mathcal{P}(m)$. Then denote by $Q^m_{\mathcal{F}}$ the k-edge-coloured multigraph on $\{0,1\}^m$ where the edges of colour r are exactly the edges
$$E_r = \{vv_{F_r} \mid v \in \{0,1\}^m\},$$
where we will take $E_r = \emptyset$ if $F_r = \emptyset$. In other words, we put edges of colour r only between vertices whose coordinates differ in exactly the set F_r.

A sequence $\mathcal{F}$ is said to *contain all the singleton sets* if for any $j \in [m]$, there exists some r such that $F_r = \{j\}$.

Now we can state the result.

Proposition 5.1.4. *Let V be a finite set and $\Sigma = \{\sigma_1, \ldots, \sigma_k\} \subseteq \mathrm{Sym}(V)$ be a set of involutions such that $\sigma_i\sigma_j = \sigma_j\sigma_i$ for all i, j. Suppose the components of G_Σ are spanned by the vertex sets $C_1, \ldots, C_t$, and write $n_l = |C_l|$ for $l = 1, \ldots, t$. Then for $1 \leqslant l \leqslant t$, $n_l = 2^{m_l}$ for some non-negative integer m_l, and there exists a colour-preserving isomorphism $\phi_l : G_\Sigma[C_l] \to Q^{m_l}_{\mathcal{F}_l}$, where $\mathcal{F}_l$ contains all the singleton sets.*

Proof. We will prove this by induction on k. If $k = 0$ then G_Σ has no edges and the result holds as each component is trivially isomorphic to $Q^0_{(\emptyset)}$. So take $k \geqslant 0$ and assume the result for k; we will show it for $k+1$. Write $\Sigma' = \Sigma \cup \{\sigma_{k+1}\}$.

Now the components of G_Σ are spanned by the sets $C_1, \ldots, C_t$; consider an arbitrary C_i. By the inductive hypothesis, there is a colour-preserving isomorphism $\phi := \phi_i$ from $G_\Sigma[C_i]$ to $Q^{m_i}_{\mathcal{F}}$, where $\mathcal{F} := \mathcal{F}_i = (F_1, \ldots, F_k)$ is a sequence from $\mathcal{P}(m_i)$ containing all the singleton sets.

First consider the case where σ_{k+1} joins a vertex $v \in C_i$ with a vertex $w \in C_j$, $j \neq i$; from Lemma 5.1.2, σ_{k+1} is a colour-preserving isomorphism from $G_\Sigma[C_i]$ to $G_\Sigma[C_j]$ (see Figure 5.1). Now define a bijection $\psi : C_i \cup C_j \to \{0,1\}^{m_i+1}$ by setting $(x)\psi = (x)\phi 0$ if $x \in C_i$ and $(x)\psi = (x)\sigma_{k+1}\phi 1$ if $x \in C_j$; we thus identify C_i with the 'bottom half' of a hypercube and C_j with the 'top half'. Now form a k-edge-coloured multigraph H on $\{0,1\}^{m_i+1}$ by putting an edge of colour i between vertices u and v if and only if there is an edge of colour i between $\psi^{-1}(u)$ and $\psi^{-1}(v)$ in G_Σ; thus $\psi : G_{\Sigma'}[C_i \cup C_j] \to H$ is a colour-preserving isomorphism.

Now fix some $r \leqslant k$. There is an edge $xy = (x,y)$[1] of colour r in G_Σ if and only if $y = (x)\sigma_r$; as $\sigma_{k+1}|_{C_i}$ is a colour-preserving isomorphism, the set of edges coloured r in $G_{\Sigma'}[C_i \cup C_j]$ is exactly

$$\big\{\big(x, (x)\sigma_r\big), \big((x)\sigma_{k+1}, (x)\sigma_r\sigma_{k+1}\big) \mid x \in C_i\big\},$$

and thus the set of edges coloured r in H is exactly

$$\begin{aligned} E'_r &= \big\{\big((x)\psi, (x)\sigma_r\psi\big), \big((x)\sigma_{k+1}\psi, (x)\sigma_r\sigma_{k+1}\psi\big) \mid x \in C_i\big\} \\ &= \big\{\big((x)\phi 0, (x)\sigma_r\phi 0\big), \big((x)\phi 1, (x)\sigma_r\phi 1\big) \mid x \in C_i\big\}, \end{aligned}$$

noting that if $x \in C_i$, $(x)\sigma_{k+1} \in C_j$ so $(x)\sigma_{k+1}\psi = (x)\sigma_{k+1}^2\phi 1 = (x)\phi 1$. Now for any $x \in C_i$, the edge $\big((x)\phi, (x)\sigma_r\phi\big) \in Q^{m_i}_{\mathcal{F}}$ is of colour r, and thus by definition of this

[1]This is a slight abuse of notation; the pair (x, y) is *not* ordered.

Figure 5.1: Suppose that $G_\Sigma[C_i]$ is isomorphic to $Q^2_{\mathcal{F}}$, where $\mathcal{F} = (\{1\}, \{2\}, \{1, 2\})$. Adding the edges of colour $k+1$ (here coloured violet) merges C_i and another set C_j to form component of $G_{\Sigma'}$ isomorphic to $Q^3_{\mathcal{F}'}$, where $\mathcal{F}' = (\{1\}, \{2\}, \{1, 2\}, \{3\})$.

graph, $(x)\sigma_r\phi = ((x)\phi)_{F_r}$. Thus we may write

$$\begin{aligned} E'_r &= \big\{\big(u0, u_{F_r}0\big), \big(u1, u_{F_r}1\big) \mid u \in \{0,1\}^{m_i}\big\} \\ &= \{vv_{F_r} \mid v \in \{0,1\}^{m_i+1}\}, \end{aligned}$$

noting that $F_r \in \mathcal{P}(m_i)$ and thus $m_i + 1 \notin F_r$. Now by construction, the set of edges of $G_{\Sigma'}[C_i \cup C_j]$ coloured $k+1$ is exactly the set

$$\big\{\big(x, (x)\sigma_{k+1}\big) \mid x \in C_i\big\},$$

and thus the set of edges of H coloured $k+1$ is exactly the set

$$\begin{aligned} E'_{k+1} &= \big\{\big((x)\psi, (x)\sigma_{k+1}\psi\big) \mid x \in C_i\big\} \\ &= \big\{\big((x)\phi0, (x)\phi1\big) \mid x \in C_i\big\} \\ &= \{uu_{m_i+1} \mid u \in \{0,1\}^{m_i+1}\}. \end{aligned}$$

So if we set $F_{k+1} = \{m_i + 1\}$ and put $\mathcal{F}' = (F_1, \ldots, F_k, F_{k+1})$ then $H = Q^{m_i+1}_{\mathcal{F}'}$, and $\mathcal{F}'$ contains all the singleton sets. We then have that $G_{\Sigma'}[C_i \cup C_j]$ is isomorphic to H by ψ, and that $|C_i \cup C_j| = 2^{m_i+1}$. Now consider the case where $(C_i)\sigma_{k+1} = C_i$, in which case we can restrict σ_{k+1} to C_i. Extend the isomorphism ϕ to include the edges of colour $k+1$ and let $H = \phi(G_{\Sigma'}[C_i])$, so H is the graph $Q^{m_i}_{\mathcal{F}}$ with some additional edges of colour $k+1$.

For any $j \in [m_i]$, there exists some r_j such that $F_{r_j} = \{j\}$; as above, the pair $\big((x)\phi, (x)\sigma_{r_j}\phi\big)$ is an edge of colour r_j in $Q^{m_i}_{\mathcal{F}}$ for each $x \in C_i$, and thus $(x)\sigma_{r_j}\phi = ((x)\phi)_{F_{r_j}} = ((x)\phi)_{\{j\}}$. Now define $\tau_j = \phi^{-1}\sigma_{r_j}\phi$ for each j and let $v \in \{0,1\}^{m_i}$ be arbitrary; then $v = (x)\phi$ for some $x \in C_i$ and

$$(v)\tau = (x)\phi\phi^{-1}\sigma_{r_j}\phi = (x)\sigma_{r_j}\phi = ((x)\phi)_{\{j\}} = v_{\{j\}}.$$

Now define $\tau := \phi^{-1}\sigma_{k+1}\phi$; then τ commutes with each of $\tau_1, \ldots, \tau_m$ as σ_{k+1} commutes with each σ_{r_j}. From Lemma 5.1.1, there is some $S \subseteq [m_i]$ such that $(v)\tau = v_S$ for all

$v \in \{0,1\}^{m_i}$. Then for any $x \in C_i$, $(x)\sigma_{k+1}\phi = (x)\phi\tau = ((x)\phi)_S$, and thus the set of edges of H coloured $k+1$ is exactly

$$\begin{aligned} E'_{k+1} &= \{\big((x)\phi, (x)\sigma_{k+1}\phi\big) \mid x \in C_i\} \\ &= \{\big((x)\phi, ((x)\phi)_S\big) \mid x \in C_i\} \\ &= \{uu_s \mid u \in \{0,1\}^{m_i}\}. \end{aligned}$$

So if we put $F_{k+1} = S$ and $\mathcal{F}' = (F_1, \ldots, F_k, F_{k+1})$ then $H = Q_{\mathcal{F}'}^{m_i+1}$ and in the graph $G_{\Sigma'}$, C_i is still a component and is isomorphic to H; note that $\mathcal{F}'$ contains all the singletons.

The result now follows by induction. □

Note that any construction of the form described above gives rise to a family of commuting involutions; this follows from the straightforward fact mentioned at the beginning, that $(v_S)_T = (v_T)_S$. We have thus determined exactly the graphical representation of such a commuting family.

5.1.3 Extremal families

We will now use the classification from the previous subsection to derive some extremal results. We will use the following notation: for any positive integer n and any sequence $\mathcal{M} = (m_1, \ldots, m_t)$ of non-negative integers such that $\sum_{i=1}^{t} 2^{m_i} = n$, write $s(\mathcal{M}) = \sum_{i=1}^{t} m_i$.

Lemma 5.1.5. *For n and $\mathcal{M}$ as defined above, $s(\mathcal{M}) \leqslant \lfloor n/2 \rfloor$. Further, if $s(\mathcal{M}) = \lfloor n/2 \rfloor$ and n is even then $m_i = 1$ or 2 for each i, while if $s(\mathcal{M}) = \lfloor n/2 \rfloor$ and n is odd then there exists a single value k such that $m_k = 0$, with $m_i = 1$ or 2 for $i \neq k$.*

Proof. For a fixed sequence $\mathcal{M} = (m_1, \ldots, m_t)$ and any $0 \leqslant j \leqslant \lfloor \log_2 n \rfloor$, let p_j denote the number of terms of $\mathcal{M}$ equal to j; then we can rearrange our previous sums to see that

$$s(\mathcal{M}) = \sum_{j=1}^{\lfloor \log_2 n \rfloor} j p_j \quad \text{and} \quad n = \sum_{j=0}^{\lfloor \log_2 n \rfloor} 2^j p_j.$$

Now observe that for any non-negative integer j, $j \leqslant 2^{j-1}$, with equality if and only if $j = 1$ or 2. Hence

$$\frac{n}{2} - s(\mathcal{M}) = \frac{p_0}{2} + \sum_{j=3}^{\lfloor \log_2 n \rfloor} (2^{j-1} - j) p_j \geqslant 0, \tag{5.1.2}$$

showing the first statement as $s(\mathcal{M})$ is an integer. Now $n = p_0 + \sum_{j=1}^{\lfloor \log_2 n \rfloor} 2^j p_j$, and so n and p_0 have the same parity; hence if n is even we can achieve equality in (5.1.2) if and only if $p_0 = 0$ and $p_j = 0$ for $j \geqslant 3$. This is equivalent to the statement that $m_i = 1$ or 2 for each i.

If n is odd then we can't achieve equality in (5.1.2) as we must have $p_0 \geqslant 1$; however we can uniquely minimise the right hand side by setting $p_0 = 1$ and $p_j = 0$ for $j \geqslant 3$. Then $s(\mathcal{M}) = (n-1)/2 = \lfloor n/2 \rfloor$ and there is precisely one term in the sequence, say m_k, equal to 0, with $m_i = 1$ or 2 for $i \neq k$. □

We also have a stability result.

Corollary 5.1.6. *Suppose the sequence $\mathcal{M}$ is allowed to depend on n, with $s(\mathcal{M}) \sim n/2$. Then, in the notation of the lemma, $p_j = p_j(n) = o(n)$ for $j \neq 1, 2$.*

Proof. Using (5.1.2),

$$0 \leqslant \frac{1}{n}\left(p_0 + \sum_{j=3}^{\lfloor \log_2 n \rfloor} p_j\right) \leqslant \frac{p_0}{n} + \sum_{j=3}^{\lfloor \log_2 n \rfloor} (2^{j-1} - j)\frac{2p_j}{n} = 1 - \frac{2s(\mathcal{M})}{n} \to 0$$

as $n \to \infty$, and thus $p_j/n \to 0$ for $j \neq 1, 2$. This gives the result. □

While not strictly relevant to later work, the following result classifies extremal families of commuting involutions.

Proposition 5.1.7. *Let Σ be a family of mutually commuting involutions in $\mathfrak{S}_n$. Then $|\Sigma| \leqslant 2^{\lfloor n/2 \rfloor}$. Further, if $|\Sigma| = 2^{\lfloor n/2 \rfloor}$, then Σ is isomorphic to a group G formed by the following procedure.*

Let U be a maximal subset of $[n]$ of even cardinality (so $|U| = n$ if n is even and $|U| = n-1$ if n is odd). Partition U as $U = P_1 \cup \cdots \cup P_t$ where each P_i has size 2 or 4. Define a permutation group G_i on P_i by setting

- *$G_i = \{\iota, (x\ y)\}$ if $P_i = \{x, y\}$ is of order 2;*
- *$G_i = \{\iota, (w\ x)(y\ z), (w\ y)(x\ z), (w\ z)(x\ y)\}$ if $P_i = \{w, x, y, z\}$ is of order 4.*

Then $G = G_1 \cdots G_t \cong G_1 \times \cdots \times G_t$.

Proof. Put $k = |\Sigma|$ and consider the k-coloured multigraph $G = G_\Sigma$ on $[n]$. By Proposition 5.1.4, G is determined by a sequence of non-negative integers $\mathcal{M} = (m_1, \ldots, m_t)$ giving the component sizes (so $|C_i| = 2^{m_i}$ and $2^{m_1} + \cdots + 2^{m_t} = n$) and t sequences $(Q_{\mathcal{F}_i}^{m_i})_{i=1}^t$ of length k with elements from $\mathcal{P}(m_i)$ such that $C_i \cong Q_{\mathcal{F}_i}^{m_i}$

for each i. For a fixed $\mathcal{M}$ and $i \in [k]$, the edges coloured i correspond to the sets $((\mathcal{F}_1)_i, \dots, (\mathcal{F}_t)_i)$; the total number of choices for any colour is therefore

$$\prod_{i=1}^{t} |\mathcal{P}(m_i)| = \prod_{i=1}^{t} 2^{m_i} = 2^{s(\mathcal{M})},$$

where $s(\mathcal{M})$ is defined in Lemma 5.1.5. Thus $k = |\Sigma| \leqslant 2^{s(\mathcal{M})}$, with equality if and only if every subset of $[m_i]$ is present (for all i) and we give a different colour to each combination.

From Lemma 5.1.5, we have $s(\mathcal{M}) \leqslant \lfloor n/2 \rfloor$ and so $|\Sigma| \leqslant 2^{\lfloor n/2 \rfloor}$. We also have a description of the extremal sequences, which we can relate to extremal component sizes; if n is even then every component has size 2 or 4, while if n is odd there is exactly one isolated vertex and every other component has size 2 or 4. We can now read off the form given in the statement by setting U to be the set of all non-isolated vertices and P_i equal to $V(C_i)$ for all i. The form of the components follows by considering them as graphs on $\{0, 1\}$ or $\{0, 1\}^2$. $\square$

Corollary 5.1.8. *Let Σ be a family of mutually commuting involutions in $\mathfrak{S}_n$ such that $\langle \Sigma \rangle$ is transitive. Then $n = 2^m$ for some m and $|\Sigma| \geqslant n$. Further, if $|\Sigma| = n$, then Σ is a group.*

Proof. From Lemma 4.3.8, the components of G_Σ correspond exactly to the orbits of $\langle \Sigma \rangle$. Hence for $\langle \Sigma \rangle$ to be transitive, G_Σ must be connected, and as each component has size a power of 2, we must have $n = 2^m$ for some m.

In the notation of Proposition 5.1.7, this means $\mathcal{M} = (m)$ and so $|\Sigma| \leqslant 2^m = n$. To have equality, we must have $\mathcal{F} = \mathcal{P}(m)$, and so after relabelling $[n]$ to $\{0, 1\}^m$, $\Sigma = \{\sigma_S \mid S \subseteq [m]\}$ where $(v)\sigma_S = v_S$ for any $v \in \{0, 1\}^m$. It is straightforward to see that these maps form a subgroup of $\mathfrak{S}_n$; $\sigma_\emptyset = \iota$ and for any S, T, $\sigma_S \sigma_T = \sigma_{S \Delta T}$. $\square$

Definition 5.1.9. Let Σ be a group as in Corollary 5.1.8. A *framework* of such a group is any set $\Sigma' \subseteq \Sigma$ such that $G_{\Sigma'}$ is connected.

It is straightforward to see that Σ' is a framework if and only if Σ' generates the group Σ. Note that the set $\Sigma' = \{\sigma_{\{i\}} \mid i \in [n]\}$ is a framework, as the resulting graph $G_{\Sigma'}$ is isomorphic to the cube. The concept of a framework will be used in the next section.

5.2 Constructing lower bounds

5.2.1 A construction

We begin with a definition.

Definition 5.2.1. Let $\mathcal{P}$ be a partition of $[n]$. Then:

1. $\mathcal{R}_\mathcal{P}$ denotes the set of all racks R such that the orbits of the operator group of R are exactly the parts of $\mathcal{P}$;

2. $\mathcal{K}_\mathcal{P} \subseteq \mathcal{R}_\mathcal{P}$ denotes the set of all kei K such that the orbits of the operator group of K are exactly the parts of $\mathcal{P}$;

3. $\mathcal{IK}_\mathcal{P} \subseteq \mathcal{K}_\mathcal{P}$ denotes the set of irreducible kei satisfying item 2.

From Lemma 4.3.8, we can also define $\mathcal{R}_\mathcal{P}$ (or $\mathcal{K}_\mathcal{P}$) to be the set of racks (or kei) such that the vertex sets of the components of G_R (or G_K) are exactly the parts of $\mathcal{P}$.

We now describe how to construct an infinite family of kei. We will do this by using Theorem 4.2.1 with $X = [n]$; we will first define the operator group G and then the maps f_{α_i}, where $(\alpha_i)_{i \in I}$ is a complete set of orbit representatives.

Construction 5.2.2. Let n be a positive integer and let $(k_1, \ldots, k_m)$ be a sequence of non-negative integers such that $m > k := \max\{k_1, \ldots, k_m\}$ and $n = \sum_{i=1}^m 2^{k_i}$. Let $\mathcal{P} = (P_i)_{i=1}^m$ be a partition of $[n]$ into m parts, with $|\mathcal{P}_i| = 2^{k_i}$ for each i.

1. For each i, let $H_i \leqslant \mathrm{Sym}(P_i)$ be a transitive, abelian group of order 2^{k_i} such that each element of H_i has order 1 or 2, and put $G = H_1 \cdots H_m \cong H_1 \times \cdots \times H_m$.

2. For each j, choose an arbitrary $\alpha_j \in P_j$, and define

$$f_{\alpha_j} = \pi_j = \phi_j^{(1)} \cdots \phi_j^{(m)}, \quad \phi_j^{(i)} \in H_i, \tag{5.2.1}$$

such that $\phi_j^{(j)} = \iota$ and the following condition holds:

3. For each $i \in [m]$, there exist k_i distinct elements $j_1, \ldots, j_{k_i} \in [m] \setminus \{i\}$ such that the set $\{\phi_{j_1}^{(i)}, \ldots, \phi_{j_{k_i}}^{(i)}\}$ forms a framework of H_i.

Proposition 5.2.3. *Each object resulting from Construction 5.2.2 is a kei in $\mathcal{K}_\mathcal{P}$. Further, $|\mathcal{K}_\mathcal{P}| \geqslant 2^M$, where $M = \sum_{i=1}^m k_i(m - k_i - 1)$.*

Proof. As mentioned at the end of Section 5.1.3, we can always select the groups H_i as in item 1. Note that the orbits of the group G are exactly the parts P_i forming $\mathcal{P}$, as H_i acts transitively on P_i but fixes all other elements. Hence $I = [m]$ is an index set for the orbits of G.

For item 2, note that the condition (4.2.1) is automatically satisfied as each H_i is an abelian group, and thus by construction so is G. There is another consequence of being abelian; let $x \in P_j$, so for some $g \in G$ we have $x = (\alpha_j)g$. Then from Lemma 4.1.10, $f_x = f_{(\alpha_j)g} = g^{-1}\pi_j g = \pi_j = f_{\alpha_j}$, and so if x and y are in the same block $\mathcal{P}_j$, $f_x = f_y = f_{\alpha_j} = \pi_j$.

Note that we have already constructed a rack on $[n]$. As $\phi_j^{(j)} = \iota$ for each j, we have a quandle; for $x \in P_j$, $(x)f_x = (x)\pi_j = (x)\phi_j^{(j)} = x$. As every element of each H_i has order 1 or 2, so does every element of G; this ensures that the structure is in fact a kei, say K.

Now from the proof of Theorem 4.2.1, the operator group of K is exactly $\langle g^{-1}\pi_j g \mid j \in I, g \in G\rangle$, which may not be the whole group G; hence we must ensure that $\mathcal{P}$ is still the set of orbits for the action induced by operator group. Condition 3 gives that P_i is a component in the graph $G_{P_{j_1} \cup \cdots \cup P_{j_{k_i}}}$ and so is automatically a component of G_K; thus from Lemma 4.3.8 P_i is an orbit of the action induced by the operator group. This means we have free choice in the remaining $m - k_i - 1$ maps $\phi_l^{(i)}$, $l \neq j_r, i$ for any r, and as each such map is an element of H_i we have

$$|\mathcal{K}_{\mathcal{P}}| \geqslant \prod_{i=1}^{m} |H_i|^{m-k_i-1} = \prod_{i=1}^{m} 2^{k_i(m-k_i-1)},$$

from which the result follows. □

Suppose we take k to be constant in Construction 5.2.2 and put $k_i = k$ for each i, so $n = 2^k m$. Let $\epsilon > 0$; using the notation from Proposition 5.2.3,

$$M = mk(m-k-1) = \frac{n}{2^k}k\left(\frac{n}{2^k} - k - 1\right) = n^2\left(\frac{k}{4^k} - \frac{k^2}{2^k n} - \frac{k}{2^k n}\right) \geqslant \left(\frac{k}{4^k} - \epsilon\right)n^2$$

for sufficiently large n. Hence $|\mathcal{K}_{\mathcal{P}}| \geqslant 2^{(c_k+o(1))n^2}$, where c_k is the constant $k/4^k$. Clearly $c_k \leqslant c_1 = 1/4$ for any k; we will now show that the asymptotic upper bound of $n^2/4$ for M is also true if we allow different values of k_i, all constant.

Corollary 5.2.4. *Let n be even and k be a constant, and let $(m_n)_{n=1}^{\infty}$ be a sequence of positive integers. Let $(k_i)_{i=1}^{\infty}$ be a sequence of sequences, with each k_i having co-domain $[k]$, such that for each n, $\max\{k_1(n), \ldots, k_{m_n}(n)\} < m_n$ and $\sum_{i=1}^{m_n} 2^{k_i(n)} = n$. For* *each* $0 \leqslant j \leqslant k$ *and each* n*, let* $p_j(n)$ *be the number of elements of* $(k_1(n), \ldots, k_{m_n}(n))$ *equal to* j*. Let* $\epsilon > 0$*; then for* n *sufficiently large,*

$$M(n) = \sum_{i=1}^{m_n} k_i(n)(m_n - k_i(n) - 1) \leqslant \left(\frac{1}{4} + \epsilon\right) n^2.$$

Further, $M(n) \sim n^2/4$ *if and only if* $p_1(n) \sim n/2$*.*

Proof. As k is constant, $n = \sum_{i=1}^{m_n} 2^{k_i(n)} \leqslant m_n 2^k$ and so $m_n \geqslant n/2^k \to \infty$ as $n \to \infty$. Now

$$(m_n - k - 1) \sum_{i=1}^{m_n} k_i(n) \leqslant M(n) \leqslant (m_n - 1) \sum_{i=1}^{m_n} k_i(n),$$

and so $M(n) \sim m_n \sum_{i=1}^{m_n} k_i(n)$; in particular, there exists some positive integer n_0 such that for $n \geqslant n_0$, $M(n) \leqslant (1 + 4\epsilon) m_n \sum_{i=1}^{m_n} k_i(n)$. For $n \geqslant n_0$ we will consider the sequence $\mathcal{M}_n = (k_1(n), \ldots, k_{m_n}(n))$; in the language of Lemma 5.1.5 we seek to maximise $m_n \sum_{i=1}^{m_n} k_i(n) = m_n s(\mathcal{M}_n)$ with the constraint that $\sum_{i=1}^{m_n} 2^{k_i(n)} = n$. Now fix such an n; to ease notation we will omit the argument in the functions k_i, $1 \leqslant i \leqslant m$, and p_j, $1 \leqslant j \leqslant k$.

As in the proof of Lemma 5.1.5, $n = \sum_{j=0}^{k} 2^j p_j$ and $m_n = \sum_{j=0}^{k} p_j$; hence

$$n \geqslant p_0 + 2 \sum_{j=1}^{k} p_j = p_0 + 2(m_n - p_0)$$

and thus $p_0 \geqslant 2m_n - n$.

Now suppose (without loss of generality) that $\mathcal{M}'_n = (k_1, \ldots, k_{m_n - p_0})$ is the subsequence of $\mathcal{M}_n$ consisting of the non-zero terms; clearly $s(\mathcal{M}_n) = s(\mathcal{M}'_n) = \sum_{j=1}^{k} j p_j$. We may also write $n = \sum_{i=1}^{m_n - p_0} 2^{k_i} + p_0$, so

$$\sum_{i=1}^{m_n - p_0} 2^{k_i} = n - p_0 \leqslant n - (2m_n - n) = 2(n - m_n).$$

It follows from Lemma 5.1.5 that $s(\mathcal{M}'_n) \leqslant n - m_n$ and thus

$$m_n \sum_{i=1}^{m_n} k_i = m_n s(\mathcal{M}_n) \leqslant m_n(n - m_n) \leqslant \frac{n^2}{4}.$$

Hence for $n \geqslant n_0$, $M(n) \leqslant (1 + 4\epsilon) n^2/4$, proving the first statement.

For the second statement, first suppose that $p_1(n) \sim n/2$ and put $\epsilon' = 1 - (1 - \epsilon)^{1/3} > 0$; then there exists a positive integer n_1 such that for $n \geqslant n_1$, $p_1(n) \geqslant (1 - \epsilon') n/2$. Now take $n \geqslant n_1$; then $m_n = \sum_{j=0}^{k} p_j(n) \geqslant p_1(n) \geqslant (1 - \epsilon') n/2$ and

$s(\mathcal{M}_n) = \sum_{j=1}^{k} jp_j \geqslant p_1(n) \geqslant (1-\epsilon')n/2$. Now as $M(n) \sim m_n \sum_{i=1}^{m_n} k_i(n)$ there exists some positive integer n_2 such that for $n \geqslant n_2$,

$$M(n) \geqslant (1-\epsilon')m_n \sum_{i=1}^{m_n} k_i(n) = (1-\epsilon')m_n s(\mathcal{M}_n).$$

It follows that for $n \geqslant \max\{n_1, n_2\}$, $M(n) \geqslant (1-\epsilon')^3 n^2/4 = (1-\epsilon)n^2/4$. We can combine this with the first result to see that $M(n) \sim n^2/4$.

Now suppose that $M(n) \sim n^2/4$; then for sufficiently large n we have that

$$(1-\epsilon)\frac{n^2}{4} \leqslant M(n) \leqslant (1+\epsilon)m_n s(\mathcal{M}_n) \leqslant (1+\epsilon)m_n(n-m_n) \leqslant (1+\epsilon)\frac{n^2}{4}.$$

It follows that

$$\frac{1-\epsilon}{1+\epsilon} \leqslant \frac{1-\epsilon}{1+\epsilon}\left(\frac{n^2}{4m_n(n-m_n)}\right) \leqslant \frac{s(\mathcal{M}_n)}{n-m_n} \leqslant 1,$$

and as ϵ is arbitrary we conclude that $m_n \sim n/2$ and $s(\mathcal{M}_n) \sim n = m_n \sim n/2$. From Corollary 5.1.6, the latter statement implies that $p_j(n) = o(n)$ for all $j \neq 1, 2$; hence $n = 2p_1(n) + 4p_2(n) + o(n)$ and $m_n = p_1(n) + p_2(n) + o(n)$. We can rearrange this to see that $m_n = n/2 - p_2(n) + o(n)$; as $m_n \sim n/2$ we have that $p_2(n) = o(n)$ as well. Hence $p_1(n) \sim n/2$. □

This result also shows that, based on our construction with constant orbit sizes, the number of kei with a given orbit structure is maximised when almost every orbit has size 2. While Construction 5.2.2 doesn't produce every kei with the prescribed structure in all cases, we can show that it is asymptotically correct in the case where every orbit does indeed have size 2.

5.2.2 Orbits of size 2

Proposition 5.2.5. *Consider a partition $\mathcal{P}$ as in Construction 5.2.2, with $k_i = 1$ for all i. Then $|\mathcal{R}_\mathcal{P}| \leqslant 2^{n^2/4}$.*

Proof. Take a rack $R \in \mathcal{R}_\mathcal{P}$ with associated maps $f_1, \ldots, f_n$ and for each i write $P_i = \{u_i, v_i\}$. As the orbits of the operator group of R are precisely the parts of $\mathcal{P}$, for each j we have

$$f_j = \sigma_j^{(1)} \cdots \sigma_j^{(m)}, \quad \sigma_j^{(i)} \in \mathrm{Sym}(P_i).$$

Note that as the group $\mathrm{Sym}(P_i) \cong C_2$ is abelian, all of the maps $f_1, \ldots, f_n$ commute with each other. As each P_i spans a component in G_R, there exists some j such that (without loss of generality) $v_i = (u_i)f_{u_j}$, and hence $f_{v_i} = f_{u_j}^{-1} f_{u_i} f_{u_j}$. But as all the maps commute, this gives $f_{v_i} = f_{u_i}$ for each i, and so knowing one map from each pair determines the other. The total number of choices for any one map is at most 2^m, and as each set of maps gives rise to a distinct rack, we obtain $|\mathcal{R}_{\mathcal{P}}| \leqslant 2^{m^2} = 2^{n^2/4}$. $\square$

For every even n, let $\mathcal{P}_n$ be a partition of $[n]$ into pairs. We can combine Propositions 5.2.3 and 5.2.5 to see that for every $\epsilon > 0$ there exists a positive integer n_0 such that for every even $n \geqslant n_0$,

$$2^{(1/4-\epsilon)n^2} \leqslant |\mathcal{K}_{\mathcal{P}_n}| \leqslant |\mathcal{R}_{\mathcal{P}_n}| \leqslant 2^{n^2/4},$$

and thus $|\mathcal{R}_{\mathcal{P}_n}| = 2^{(1/4+o(1))n^2}$. For n odd, let $\mathcal{P}_n = \mathcal{P}'_n \cup \{n\}$, where $\mathcal{P}'_n$ is a partition of $[n-1]$ into pairs. If $n \geqslant n_0 + 1$ is odd, we have that $|\mathcal{K}_{\mathcal{P}'_n}| \geqslant 2^{(1/4-\epsilon)(n-1)^2}$; each kei $K' \in \mathcal{K}_{\mathcal{P}'_n}$ can be combined with the trivial kei $T_{\{n\}}$ on $\{n\}$ to construct a kei $K' \oplus T_{\{n\}} \in \mathcal{K}_{\mathcal{P}_n}$. This shows that for any n, $|\mathcal{R}_n| \geqslant 2^{(1/4+o(1))n^2}$; as there are (crudely) at most n^n choices for the partition $\mathcal{P}_n$ into pairs, we can't obtain a significantly larger lower bound by considering different partitions.

The fact that this lower bound is so large makes the counting of isomorphism classes easier, as noted by Blackburn in [2]. For any rack $R = ([n], \triangleright)$ there are at most $n!$ racks on $[n]$ isomorphic to R, one for each possible map $\phi : [n] \to [n]$. This means there are at least $2^{(1/4+o(1))n^2}/n!$ isomorphism classes, but as $n! = 2^{o(n^2)}$, we do actually have $2^{(1/4+o(1))n^2}$ distinct isomorphism classes. This shows the following corollary.

Corollary 5.2.6. *Let $\epsilon > 0$. Then there exists a positive integer n_0 such that for any $n \geqslant n_0$, there are at least $2^{(1/4-\epsilon)n^2}$ isomorphism classes of racks of order n.* $\square$

Another consequence of having such a large lower bound is that an upper bound on the number of irreducible racks on $[n]$ gives rise immediately to an upper bound on the total number of racks on $[n]$. Suppose $([n], \triangleright)$ is a rack; from Definition 4.3.2, we can partition $[n]$ into sets $U_1, \ldots, U_l$ such that each $(U_i, \triangleright)$ is an irreducible subrack. Now suppose we had a bound of $2^{(c+o(1))m^2}$ for the number of *irreducible* racks on $[m]$, where $1/4 \leqslant c \leqslant 1.1604\ldots$; then there would be at most $2^{(c+o(1))|U_i|^2}$ choices for $(U_i, \triangleright)$. There are (crudely) at most n^n ways of choosing the partition into the sets $U_1, \ldots, U_l$, and so the total number of racks on $[n]$ is at most

$$n^n \prod_{i=1}^{l} 2^{(c+o(1))|U_i|^2} \leqslant 2^{o(n^2)} 2^{(c+o(1))n^2} = 2^{(c+o(1))n^2}.$$

This restriction will be helpful in Section 6.2 but we do not need it to prove the overall upper bound for racks.

To end this chapter, consider Construction 5.2.2 again. By Corollary 5.1.8, we have to choose each $\mathcal{P}_i$ to have size a power of two to find a transitive, abelian subgroup of involutions H_i. In the next chapter we will investigate a different situation, where each $\mathcal{P}_i$ has size an odd prime.

Chapter 6

Kei with prime orbits

Recall that if $\mathcal{P}$ is a partition of $[n]$, then $\mathcal{K}_{\mathcal{P}}$ denotes the set of all kei K such that the orbits of the operator group of K are exactly the parts of $\mathcal{P}$, while $\mathcal{IK}_{\mathcal{P}}$ denotes the subset of $\mathcal{K}_{\mathcal{P}}$ consisting of irreducible kei. In Section 5.2, we were able to classify all kei $K \in \mathcal{K}_{\mathcal{P}}$ for a partition $\mathcal{P}$ of $[n]$ consisting only of pairs; we hence showed that $|\mathcal{K}_{\mathcal{P}}| = 2^{(1/4+o(1))n^2}$. In this section, we show how to obtain a similar classification when each part of the partition $\mathcal{P}$ has odd prime size p; the number of such kei is negligible compared to $2^{n^2/4}$. Let $m = n/p$; we can show exactly (Theorem 6.3.7) that

$$|\mathcal{IK}_{\mathcal{P}}| = \frac{(p!)^m}{p(p-1)} \cdot (2^m - 1),$$

and it follows (Corollary 6.3.8) that $|\mathcal{K}_{\mathcal{P}}| \leqslant (2p!m)^m = 2^{o(n^2)}$.

We begin by defining a form of quotient for racks. In Section 6.2.1 we analyse components of the subgraph G_S, where $(S, \triangleright)$ is a subrack; it then follows that there are only two possibilities for the subkei induced by a single part of $\mathcal{P}$; it is either trivial or isomorphic to a dihedral kei (these kei are discussed in Section 6.2.3). After some technical lemmas, we can then show that each such kei is in fact an expansion of a quotient as described in Section 6.1, allowing us to complete the classification.

6.1 A form of quotient

6.1.1 Taking quotients

Let $R = ([n], \triangleright)$ be a rack and $f_1, \ldots, f_n$ be the associated maps. Define a relation $\sim$ on $[n]$ by setting

$$i \sim j \text{ if and only if } f_i = f_j. \tag{6.1.1}$$

It is easy to see that this is an equivalence relation; write $[[i]]$ for the equivalence class of i and $\mathcal{C}$ for the set of equivalence classes. If we define an operation on $\mathcal{C}$ in the natural way, we will obtain a *reduced rack*, as described below.

Proposition 6.1.1. *In the above notation, the binary operation $\triangleright' : \mathcal{C} \times \mathcal{C} \to \mathcal{C}$ given by $[[i]] \triangleright' [[j]] = [[i \triangleright j]]$ is well-defined. Further, $(\mathcal{C}, \triangleright')$ is a rack, and if R is a quandle (or kei) then $(\mathcal{C}, \triangleright')$ is a quandle (or kei).*

Proof. Suppose $i \sim i'$ and $j \sim j'$; then by definition $f_i = f_{i'}$ and $f_j = f_{j'}$. Then $f_j^{-1} f_i f_j = f_{j'}^{-1} f_{i'} f_{j'}$, or equivalently $f_{(i)f_j} = f_{(i')f_{j'}}$. In the $\triangleright$ notation this becomes $f_{i \triangleright j} = f_{i' \triangleright j'}$ or $i \triangleright j \sim i' \triangleright j'$, so $[[i \triangleright j]] = [[i' \triangleright j']]$. Thus the operation is well-defined. The proof of the remaining statements simply involve straightforward checking of axioms, and are omitted. $\square$

Note that by setting $j = j'$ in the above proof we obtain that if $i \sim i'$ then $(i)f_j \sim (i')f_j$ for every j. The reverse is also true; if $(i)f_j \sim (i')f_j$ then $f_{(i)f_j} = f_{(i')f_j}$ and thus $f_j^{-1} f_i f_j = f_j^{-1} f_{i'} f_j$. Hence $f_i = f_{i'}$ and $i \sim i'$, and we conclude that each map f_j of the original rack permutes the equivalence classes. The following lemma is very straightforward but will be quite illuminating; note that as before, we define $f_C \in \mathrm{Sym}(\mathcal{C})$ by $(D)f_C = D \triangleright' C$ for $C, D \in \mathcal{C}$.

Lemma 6.1.2. *For any $C, D \in \mathcal{C}$, $|(D)f_C| = |D|$.*

Proof. Let $i \in D$, $j \in C$, so that $(D)f_C = D \triangleright' C = [[i]] \triangleright' [[j]] = [[i \triangleright j]] = [[(i)f_j]]$. But note that as $i \in D$ and f_j permutes equivalence classes, the equivalence class of $(i)f_j$ is the class D translated by f_j; as f_j is a bijection we have $|(D)f_C| = |D|$, as desired. $\square$

Now write $\mathcal{C} = \{C_1, \ldots, C_m\}$. Choose an equivalence class C_j and take $x \in C_j$; how do the maps f_x and f_{C_j} relate to each other? To make them comparable, we will first define an extension $\hat{f}_{C_j}$ of the latter to the whole of $[n]$. This requires more notation; for any class C_i and $1 \leqslant r \leqslant |C_i|$, let $C_i\langle r \rangle$ denote the rth smallest element of C_i. Then $\hat{f}_{C_j}|_{C_i}$ is defined by $C_i\langle r \rangle \mapsto ((C_i)f_{C_j})\langle r \rangle$. From Lemma 6.1.2, $|C_i| = |(C_i)f_{C_j}|$ and $\hat{f}_{C_j}$ is a bijection from C_i to $(C_i)f_{C_j}$, mimicking element-by-element the action of f_{C_j} on the equivalence class C_i. Hence for each j, $\hat{f}_{C_j}$ is a bijection on the whole of $[n]$.

Now consider the map $\hat{f}_{C_j}^{-1} f_x \in \mathfrak{S}_n$. For any $y \in C_i$, we have $(y)\hat{f}_{C_j}^{-1} \in (C_i)f_{C_j}^{-1}$ and so $(y)\hat{f}_{C_j}^{-1} f_x \in (C_i)f_{C_j}^{-1} f_{C_j} = C_i$; hence the composition map preserves the equivalence classes (but may permute the elements within a class). Then we can define, for any j and any $x \in C_j$, the map $\sigma_{ij} = \hat{f}_{C_j}^{-1} f_x|_{C_i} \in \mathrm{Sym}(C_i)$. In this notation we have

$$f_x = \hat{f}_{C_j} \sigma_{1j} \cdots \sigma_{mj}. \tag{6.1.2}$$

This relation makes it clear that given the reduced rack $(\mathcal{C}, \triangleright')$ and the complete set of maps σ_{ij}, $1 \leqslant i, j \leqslant m$, we have enough information to reconstruct the original rack $([n], \triangleright)$. At present the maps σ_{ij} act on different sets; it will be helpful to make them somewhat more comparable. To this end, we define for each i the canonical bijection $\rho_i : C_i \to [|C_i|]$ by $C_i\langle r \rangle \mapsto r$, and then set $\theta_{ij} = \rho_i^{-1} \sigma_{ij} \rho_i$. Then each θ_{ij} is a bijection on $[|C_i|]$, i.e. $\theta_{ij} \in S_{|C_i|}$. We can think of these permutations being stored in an $m \times m$ matrix Θ, where the jth column encodes the information in the map $\hat{f}_{C_j} f_x$ for any $x \in C_j$. As previously mentioned, we have enough information to reconstruct $R = ([n], \triangleright)$; let us write $R = [\mathcal{C}, \triangleright', \Theta]$ in this case.

We can now recast the action of any f_x in this notation. Take $i, j \in [m]$ and suppose $C_k = (C_i) f_{C_j}$; let $x \in C_j$ and $1 \leqslant r \leqslant |C_i|$. Then $(C_i\langle r \rangle)\hat{f}_{C_j} = C_k\langle r \rangle$, and hence from (6.1.2),

$$(C_i\langle r \rangle) f_x = (C_k\langle r \rangle) \sigma_{1j} \cdots \sigma_{mj} = (C_k\langle r \rangle) \sigma_{kj},$$

noting that for $i \neq k$, σ_{ij} fixes all elements in C_k, and that for a fixed j, all of the maps σ_{ij} commute with each other. Then

$$(C_k\langle r \rangle)\sigma_{kj} = (C_k\langle r \rangle)\rho_k \theta_{kj} \rho_k^{-1} = (r)\theta_{kj}\rho_k^{-1} = C_k\langle (r)\theta_{kj} \rangle,$$

so we see

$$(C_i\langle r \rangle) f_x = C_k\langle (r)\theta_{kj} \rangle. \tag{6.1.3}$$

Note that we can apply (6.1.3) repeatedly; for instance, if we also know that $(C_k) f_{C_p} = C_q$ then for any $c' \in C_p$ we have

$$(C_i\langle r \rangle) f_c f_{c'} = (C_k\langle (r)\theta_{kj} \rangle) f_{c'} = C_q\langle (r)\theta_{kj}\theta_{qp} \rangle \tag{6.1.4}$$

6.1.2 Racks with a given quotient

Now suppose we have a triple $[\mathcal{C}, \triangleright', \Theta]$, where $\mathcal{C} = \{C_1, \ldots, C_m\}$ is a partition of $[n]$, $(\mathcal{C}, \triangleright')$ is a rack, and Θ is an $m \times m$ matrix where $\theta_{ij} \in S_{|C_i|}$ for all i, j. We would first like to find conditions on these objects such that $R = [\mathcal{C}, \triangleright', \Theta]$ is a rack on $[n]$; if $(\mathcal{C}, \triangleright')$ is a quandle (or kei) then we would also like additional conditions ensuring that $[\mathcal{C}, \triangleright', \Theta]$ is a quandle (or kei).

For clarity, suppose that $\mathcal{C}$ is ordered such that $|C_1| \leqslant \cdots \leqslant |C_m|$. If concreteness is desired, we may set, for $1 \leqslant i \leqslant m$ and $1 \leqslant r \leqslant |C_i|$,

$$C_i\langle r\rangle := \sum_{j=1}^{i-1} |C_j| + r,$$

and so we can think of each C_i as using the next $|C_i|$ available numbers from $[n]$, in increasing order. For example, suppose $n = 9$ and $m = 5$, and $\mathcal{C}$ consists of two singletons, two pairs and a three-set. Then in the above setup, we will take

$$C_1 = \{1\},\ C_2 = \{2\},\ C_3 = \{3, 4\},\ C_4 = \{5, 6\},\ C_5 = \{7, 8, 9\},$$

so $C_3\langle 1\rangle = 3$ and $C_5\langle 2\rangle = 8$. If we had $f_{C_1} = (C_3\ C_4)$ then the extended map $\hat{f}_{C_1} = (3\ 5)(4\ 6)$.

Now from (6.1.3) we know the forms the maps $f_1, \ldots, f_n$ must take. We can check if they form a rack on $[n]$ by using Proposition 4.1.8, and the conditions are given in the next result.

Proposition 6.1.3. *Let $\mathcal{C} = \{C_1, \ldots, C_m\}$ be a partition of $[n]$, $(\mathcal{C}, \triangleright')$ be a rack, and Θ be an $m \times m$ matrix where $\theta_{ij} \in S_{|C_i|}$ for all i, j; given the triple $[\mathcal{C}, \triangleright', \Theta]$, define maps $f_1, \ldots, f_n$ on $[n]$ via* (6.1.3)*, or equivalently* (6.1.2)*. Then $f_1, \ldots, f_n$ form a rack $R = [\mathcal{C}, \triangleright', \Theta]$ on $[n]$ if and only if*

$$\theta_{\alpha k} = \theta_{\delta j}^{-1} \theta_{\beta i} \theta_{\alpha j} \tag{6.1.5}$$

for all i, j, $\alpha \in [m]$, where $(C_i)f_{C_j} = C_k$, $C_\beta = (C_\alpha)f_{C_j}^{-1}$ and $C_\delta = (C_\beta)f_{C_i}^{-1}f_{C_j}$ (see Figure 6.1).

Further, suppose that (6.1.5) *holds, so $([n], \triangleright)$ is a rack, and let $\mathcal{D}$ be the set of equivalence classes under $\sim$. Then $\mathcal{C}$ is a refinement of $\mathcal{D}$, with equality if and only if $\hat{f}_{C_j}\sigma_{1j} \cdots \sigma_{mj} \neq \hat{f}_{C_k}\sigma_{1k} \cdots \sigma_{mk}$ for $j \neq k$.*

Proof. By our construction, the maps $f_1, \ldots, f_n$ are all bijections; from Proposition 4.1.8, $([n], \triangleright)$ is a rack if and only if $f_{(x)f_y} = f_y^{-1} f_x f_y$ for all x, y.

Again by our construction, any element of $[n]$ can be expressed uniquely as $C_i\langle r\rangle$ for some $1 \leqslant i \leqslant m$ and $1 \leqslant r \leqslant |C_i|$; choose two elements $x = C_i\langle r\rangle$, $y = C_j\langle s\rangle$ and an arbitrary $\alpha \in [m]$. Now define β, γ and δ as in Figure 6.1, i.e. $C_\beta = (C_\alpha)f_{C_j}^{-1}$, $C_\gamma = (C_\beta)f_{C_i}^{-1}$ and $C_\delta = (C_\gamma)f_{C_j}$. For any $t \in [|C_\gamma|]$, we have from (6.1.4) that

$$(C_\gamma\langle t\rangle)f_x f_y = (C_\gamma\langle t\rangle)f_{C_i\langle r\rangle} f_{C_j\langle s\rangle} = C_\alpha\langle (t)\theta_{\beta i}\theta_{\alpha j}\rangle. \tag{6.1.6}$$

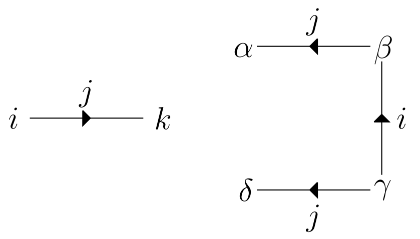

Figure 6.1: The relationships between i, j, k, α, β, γ and δ.

Now suppose $(C_i)f_{C_j} = C_k$; from (6.1.3) we have $(x)f_y = C_k\langle(r)\theta_{kj}\rangle = C_k\langle u\rangle$ say. As the maps $f_{C_1}, \ldots, f_{C_m}$ form a rack, we have $(C_\delta)f_{C_k} = (C_\delta)f_{C_j}^{-1}f_{C_i}f_{C_j} = C_\alpha$, so applying (6.1.4) again,

$$(C_\gamma\langle t\rangle)f_y f_{(x)f_y} = (C_\gamma\langle t\rangle)f_{C_j\langle s\rangle}f_{C_k\langle u\rangle} = C_\alpha\langle(t)\theta_{\delta j}\theta_{\alpha k}\rangle. \tag{6.1.7}$$

By comparing (6.1.6) and (6.1.7), we see that $f_y f_{(x)f_y} = f_x f_y$ for all x, y if and only if $\theta_{\beta i}\theta_{\alpha j} = \theta_{\delta j}\theta_{\alpha k}$ for all i, j, α (with the remaining variables defined as above), proving the first result.

Finally, let $x \in C_j$ and $y \in C_k$. Then from (6.1.2), $f_x = f_y$ (i.e. x and y are in the same equivalence class in $\mathcal{D}$) if and only if $\hat{f}_{C_j}\sigma_{1j}\cdots\sigma_{mj} = \hat{f}_{C_k}\sigma_{1k}\cdots\sigma_{mk}$. Thus the remaining statements follow. □

Proposition 6.1.4. *Let $[\mathcal{C}, \triangleright', \Theta]$ be a triple as in Proposition 6.1.3 and define maps $f_1, \ldots, f_n$ on $[n]$ via* (6.1.3)*, or equivalently* (6.1.2)*. Then if $(\mathcal{C}, \triangleright')$ is a rack, the rack $[\mathcal{C}, \triangleright', \Theta]$ is a quandle on $[n]$ if and only if*

$$\theta_{ii} = \iota \tag{6.1.8}$$

for all i. Further, if $(\mathcal{C}, \triangleright')$ is a quandle, the quandle $[\mathcal{C}, \triangleright', \Theta]$ is a kei on $[n]$ if and only if

$$\theta_{kj} = \theta_{ij}^{-1} \tag{6.1.9}$$

for all i, where $(C_i)f_{C_j} = C_k$.

Proof. First suppose that $(\mathcal{C}, \triangleright')$ is a quandle; recall that $([n], \triangleright)$ is a quandle if and only if $(x)f_x = x$ for all x. Take any $x \in [n]$; as before, $x = C_i\langle r\rangle$ for some $1 \leqslant i \leqslant m$ and $1 \leqslant r \leqslant |C_i|$. As $(C_i)f_{C_i} = C_i$ we have from (6.1.3) that

$$(x)f_x = (C_i\langle r\rangle)f_{C_i\langle r\rangle} = C_i\langle(r)\theta_{ii}\rangle,$$

which is equal to $C_i\langle r\rangle$ if and only if $(r)\theta_{ii} = r$. Hence $(x)f_x = x$ for all x if and only if $\theta_{ii} = \iota$ for all i.

Now suppose $(\mathcal{C}, \triangleright')$ is a kei; recall that $([n], \triangleright)$ is a kei if and only if $f_x^2 = \iota$ for all x. Take two elements $C_i\langle r\rangle$ and $x = C_j\langle s\rangle$ and write $C_k = (C_i)f_{C_j}$. Then $(C_k)f_{C_j} = C_i$, so applying (6.1.4) with $p = j$ and $q = i$ gives

$$(C_i\langle r\rangle)f_x^2 = (C_i\langle r\rangle)f_{C_j\langle s\rangle}^2 = C_i\langle (r)\theta_{kj}\theta_{ij}\rangle,$$

which is equal to $C_i\langle r\rangle$ if and only if $(r)\theta_{kj}\theta_{ij} = r$. Hence $f_x^2 = \iota$ for all x if and only if $\theta_{kj}\theta_{ij} = \iota$, or equivalently $\theta_{kj} = \theta_{ij}^{-1}$, for all i, j (where k is defined as above). □

Note that if $(\mathcal{C}, \triangleright')$ then (6.1.9) holds, so we can replace the term $\theta_{\delta j}^{-1}$ in (6.1.5) with $\theta_{\gamma j}$ (as $(C_\gamma)f_{C_j} = C_\delta$), giving $\theta_{\alpha k} = \theta_{\gamma j}\theta_{\beta i}\theta_{\alpha j}$.

Propositions 6.1.3 and 6.1.4 mean that if we start with an abstract kei K and take the elements of the ground set to be a partition $\mathcal{C}$ of $[n]$, we can construct a kei on $[n]$ such that the reduced kei is isomorphic to K, provided we satisfy (6.1.5), (6.1.8) and (6.1.9). As an example, suppose we have the trivial kei on two elements, which we will label $C_1 = \{1, 2\}$ and $C_2 = \{3, 4\}$, which we want to expand to a kei on $\{1, 2, 3, 4\}$; we need to find elements θ_{11}, θ_{12}, θ_{21} and θ_{22}, all in S_2, satisfying (6.1.5), (6.1.8) and (6.1.9). Both maps f_{C_1} and f_{C_2} are trivial, and then it is straightforward to see that (6.1.5) and (6.1.9) reduce to trivialities; so we can choose Θ to be any of

$$\begin{pmatrix} \iota & (1\,2) \\ \iota & \iota \end{pmatrix}, \quad \begin{pmatrix} \iota & \iota \\ (1\,2) & \iota \end{pmatrix}, \quad \begin{pmatrix} \iota & (1\,2) \\ (1\,2) & \iota \end{pmatrix} \quad \text{or} \quad \begin{pmatrix} \iota & \iota \\ \iota & \iota \end{pmatrix}.$$

In the last case, each $\sigma_{ij} = \iota$, and as $\hat{f}_{C_1} = \hat{f}_{C_2} = \iota$, we see that $\mathcal{C} \neq \mathcal{D}$; in this case the expanded kei is trivial and $\mathcal{D}$ is the partition consisting of just the element $\{1, 2, 3, 4\}$.

6.1.3 Relation to previous concepts

The kei generated by Construction 5.2.2 had the property that all the maps in the same block P_i acted in the same way across the entire ground set; in the language of the previous subsections this means $x \sim y$ whenever $x, y \in P_i$. As each map fixes P_i, it is straightforward to see that the reduced kei of any such K is trivial. We now generalise the argument given at the end of Section 6.1.2 to show that *all* kei expanded from trivial kei can be constructed from Construction 5.2.2. Recall that we denote the trivial kei on a set X by T_X.

Proposition 6.1.5. *Let $K = [\mathcal{C}, \triangleright', \Theta]$ be a kei on $[n]$, and suppose $(\mathcal{C}, \triangleright') = T_{\mathcal{C}}$. Then there is a partition $\mathcal{P}$ of $[n]$, refining $\mathcal{C}$, such that K is isomorphic to a kei as constructed in Construction 5.2.2.*

Proof. As before, write $\mathcal{C} = \{C_1, \ldots, C_m\}$. We have $(C_i)f_{C_j} = C_i$ for all i, j, so (6.1.9) becomes $\theta_{ij} = \theta_{ij}^{-1}$ for all i, j, and (6.1.5) gives $\theta_{\alpha i} = \theta_{\alpha j}\theta_{\alpha i}\theta_{\alpha j}$ for all i, j and α. Thus every element of Θ must be an involution, and the elements in any row must all commute with one another. In terms of the original maps σ_{ij}, this means that $\sigma_{ii} = \iota$ for all i, each map σ_{ij} is an involution, and for a fixed i, the maps $(\sigma_{ij})_{j=1}^{m}$ mutually commute.

Now these maps all act on the class C_i, so from Proposition 5.1.4 there is a partition $\mathcal{P}_i$ of C_i such that each element of $\mathcal{P}_i$ has size a power of 2 and is an orbit of the natural action of $\langle \sigma_{ij} \rangle_{j=1}^{m}$ on C_i. Now let $\mathcal{P} = \bigcup_{i=1}^{m} \mathcal{P}_i$, so that $\mathcal{P}$ is a partition of $[n]$; write $\mathcal{P} = \{P_1, \ldots, P_r\}$ and put $H_i = \langle \sigma_{ij}|_{P_i} \rangle_{j=1}^{m}$.

Now for any class C_j, f_{C_j} is the identity map, and so $\hat{f}_{C_j}$ is as well. Hence from (6.1.2), $f_c = \sigma_{1j} \cdots \sigma_{mj}$ for any $c \in C_j$. But as the component structure of each C_k is exactly $\mathcal{P}_k$, there exists $\sigma'_{ij} \in H_i$, $1 \leqslant i \leqslant r$, such that

$$f_c = \sigma'_{1j} \cdots \sigma'_{mj},$$

and further, the involutions $(\sigma'_{ij})_{j=1}^{r'}$ all commute for a fixed i. Comparing this to (5.2.1) gives that $\sigma'_{ij} = \phi_j^{(i)}$ for each i; now choose an arbitrary $\alpha_j \in C_j$ and define $\pi_j = f_{\alpha_j} = f_c$.

This shows item 2 of Construction 5.2.2; for item 1 note that the product of any two commuting involutions σ and τ is again an involution as $(\sigma\tau)^2 = \sigma^2\tau^2 = \iota$. Hence each H_i is an abelian group of involutions which is by definition transitive on P_i. For item 3, as each P_i is an orbit of the original natural action there is a framework as described. $\square$

We will use these quotients to help classify all kei in which every orbit has an odd prime size, as described in the rest of this chapter.

6.2 Subset actions and generators

In this section we introduce some concepts that will be needed in the classification result.

6.2.1 Orbits of subset actions

Definition 6.2.1. Let $([n], \triangleright)$ be a rack and $f_1, \ldots, f_n$ be the associated maps. For subsets S, T of $[n]$, we say that S is T-*invariant* if for every $i \in S$ and $j \in T$ we have $(i)f_j \in S$; as each map is a bijection this is equivalent to $(S)f_j = S$ for every $j \in T$.

Observe that $(S, \triangleright)$ is a subrack if and only if S is S-invariant. Also note that if S is $\{j\}$-invariant (which we may write as j-invariant) then $f_j|_S$ is a permutation of the elements in S; we shall write $\phi_j^{(S)} := f_j|_S \in \mathrm{Sym}(S)$ in this case.

Now for an arbitrary subset $T \subseteq [n]$, there is a canonical partition of $[n]$ into a set of T-invariant subsets, namely the partition into orbits of the group $\langle f_t \mid t \in T \rangle$ in its natural action on $[n]$. We will call such an orbit a *T-orbit*, and write $\mathcal{C}(T)$ for the partition of $[n]$ into T-orbits.

From Proposition 4.3.8, a set $U \subseteq [n]$ is a T-orbit if and only if U spans a component of G_T, so in graphical terms, $\mathcal{C}(T)$ is the partition of $[n]$ corresponding to the components of G_T.

We will now show that components of invariant sets are highly structured.

Lemma 6.2.2. *Let $T \subseteq [n]$ be j-invariant for some $j \in [n]$, and let C and D be distinct elements of $\mathcal{C}(T)$, which is to say that C and D span distinct components of G_T. Suppose that there exist $u_0 \in C$, $w_0 \in D$ such that $(u_0)f_j = w_0$. Then $G_T[C]$ and $G_T[D]$ are isomorphic in the following sense; for any $u, v \in C$ and $i \in T$,*

$$(u)f_i = v \textit{ if and only if } ((u)f_j)f_{(i)f_j} = (v)f_j. \tag{6.2.1}$$

The statement has a nice graphical interpretation; the map f_j restricted to C is a bijection onto D, and a directed edge of colour i between two vertices u and v in C is replaced by a directed edge of the transformed colour $(i)f_j$ between the transformed vertices $(u)f_j$ and $(v)f_j$ in D. The proof is very similar to the proof of Lemma 5.1.2; indeed the earlier lemma may in a certain case be seen as a corollary of this result.

Proof. Define $C' \subseteq C$ by

$$C' := \{u \in C \mid (u)f_j \in D\},$$

so by assumption $u_0 \in C'$. Write $D' = (C')f_j \subseteq D$. Now suppose that $C' \neq C$, so there exists $v_0 \in C \setminus C'$; as C spans a component of G_T we have from Lemma 4.3.7 that there is a directed u_0v_0-path in G_T. By considering the last vertex of this path lying in C', we see that there exist $u \in C'$ and $v \in C \setminus C'$ such that there is a directed edge from u to v of some colour $i \in T$, i.e. that $(u)f_i = v$. Now as T is j-invariant, $(i)f_j \in T$; by the definition of C', $(u)f_j \in D' \subseteq D$, and as D is a component of G_T, $(u)f_j f_{(i)f_j} \in D$. But $f_{(i)f_j} = f_j^{-1} f_i f_j$ and hence $D \ni (u)f_i f_j = (v)f_j$, and so by definition $v \in C'$, which is a contradiction. Hence $C' = C$. We can prove that $D' = D$ similarly; if $D' \neq D$ then there exist $w \in D'$ and $x \in D \setminus D'$ such that $(w)f_k = x$ for some $k \in T$. But as T is j-invariant, $(k)f_j^{-1} \in T$, and by Lemma

4.1.10, $f_{(k)f_j^{-1}} = f_j f_k f_j^{-1}$; thus $(w)f_k f_j^{-1} = (w)f_j^{-1} f_{(k)f_j^{-1}} \in C$, a contradiction to the definition of D'.

Now let $u, v \in C$ and suppose that there is a directed edge from u to v of colour $i \in T$ (so as before, $(u)f_i = v$). Then $(v)f_j = (u)f_i f_j = (u)f_j f_{(i)f_j}$, i.e. $f_{(i)f_j}$ maps $(u)f_j$ onto $(v)f_j$ and there is a directed edge from $(u)f_j$ to $(v)f_j$ of colour $(i)f_j$. By considering f_j^{-1} on D as above, we obtain the reverse implication, which is to say that $f_j|_C$ is an isomorphism in the sense described above. $\square$

Note that as C and D are T-invariant, we may write (6.2.1) as

$$(u)\phi_i^{(C)} = v \text{ if and only if } ((u)f_j)\phi_{(i)f_j}^{(D)} = (v)f_j,$$

or, substituting $(u)\phi_i^{(C)}$ for v, $((u)f_j)\phi_{(i)f_j}^{(D)} = ((u)\phi_i^{(C)})f_j$. Hence the statement (6.2.1) may be compactly written as

$$\phi_{(i)f_j}^{(D)} = f_j^{-1}|_D \circ \phi_i^{(C)} \circ f_j|_C$$

for all $i \in T$.

In many cases we will not assume full knowledge of the map f_j, so it is useful to consider a slightly weaker form of isomorphism.

Definition 6.2.3. Let $T \subseteq [n]$, and let $C, D \in \mathcal{C}(T)$ (so C and D span components of G_T). Then $G_T[C]$ and $G_T[D]$ are *edge-isomorphic* if there exists a bijection $\chi : C \to D$ and a permutation $\sigma \in \mathrm{Sym}(T)$ such that

$$\phi_{(i)\sigma}^{(D)} = \chi^{-1} \circ \phi_i^{(C)} \circ \chi$$

for all $i \in T$.

Now let us consider the case where we have prescribed the final component structure, i.e. we set $\mathcal{C}([n]) = \{B_1, \ldots, B_m\}$ for some fixed 'blocks' $B_1, \ldots, B_m$. Note that for any j, each B_i is j-invariant, and so we can write

$$f_j = \phi_j^{(1)} \cdots \phi_j^{(m)},$$

where we use $\phi_j^{(i)}$ as shorthand for $\phi_j^{(B_i)}$. Any multiplication of these elements happens block-wise, so that $f_k = f_i^{-1} f_j f_i$ if and only if $\phi_k^{(l)} = \left(\phi_i^{(l)}\right)^{-1} \phi_j^{(l)} \phi_i^{(l)}$ for all l. Also note that for any $T \subset [n]$, $\mathcal{C}(T)$ is a refinement of $\mathcal{C}([n])$, from the definition of orbits given above. We now show that for certain subsets T, all the components within a given B_i must be edge-isomorphic.

Lemma 6.2.4. *Let $t \subseteq [m]$, and put $T = \bigcup_{k \in t} B_k$, so that T is a disjoint union of some of the blocks. Let $C, D \in \mathcal{C}(T)$ be distinct orbits such that $C, D \subseteq B_i$ for some i. Then $G_T[C]$ and $G_T[D]$ are edge-isomorphic.*

Proof. As $\mathcal{C}(T)$ is a refinement of $\mathcal{C}([n])$, the block B_i has a partition $\mathcal{P}$ into T-orbits; define an equivalence relation on $\mathcal{P}$ by setting $C \sim D$ if and only if $G_T[C]$ and $G_T[D]$ are edge-isomorphic.

It is easy to see that this is indeed an equivalence relation; let $\mathcal{I}_1, \ldots, \mathcal{I}_r$ be the resulting equivalence (isomorphism) classes. Suppose that $r > 1$; as B_i is an $[n]$-orbit there exist (without loss of generality) T-orbits $C \in \mathcal{I}_1$, $D \in \mathcal{I}_2$ and vertices $u \in C$, $v \in D$ such that u and v are joined by a directed edge of some colour j, i.e. $(u)f_j = v$. Now as T is a disjoint union of $[n]$-orbits, T is j-invariant, and so from Lemma 6.2.2, $G_T[C]$ and $G_T[D]$ are edge-isomorphic, with $\chi = f_j|_C$ and $\sigma = f_j|_T = \phi_j^{(T)}$. But this is a contradiction as C and D are in different isomorphism classes. Hence $r = 1$ and thus any two T-orbits lying in B_i are edge-isomorphic. □

6.2.2 Generated subkei

For the rest of this chapter, we return our attention exclusively to kei; this means that whenever f_x is a map associated with a kei $([n], \triangleright)$, $f_x^{-1} = f_x$. We will thus omit the inverse notation.

Definition 6.2.5. Let $([n], \triangleright)$ be a kei and $f_1, \ldots, f_n$ be the associated maps. Let $X \subseteq [n]$ be a subkei under $\triangleright$ and let $S \subseteq X$. We will say that X is *generated* by S if for every $x \in X$ there exist a finite number of elements $j, i_1, \ldots, i_r \in S$ such that $x = (j)f_{i_1} \cdots f_{i_r}$.

Note that for such a sequence we have $f_x = f_{i_r} \cdots f_{i_1} f_j f_{i_1} \cdots f_{i_r}$, which is to say that every edge in G_X can be replaced with a walk in G_S. Hence if S generates X, two vertices in $[n]$ are in the same component of G_X if and only if they are in the same component of G_S; if S is much smaller than X then this can be an efficient way to determine the components of G_X.

With this in mind, we now consider the components of the graph $G_{ij} := G_{\{i,j\}}$, for two distinct elements i and j. As each vertex in this graph is incident to at most one edge of colour i and at most one edge of colour j, G_{ij} consists only of paths and cycles. Further, edges must alternate in colour along any path or around any cycle, and so each cycle is of even length. Note that it could be the case that two vertices are joined by two edges, one of each colour; we will treat this as a degenerate 'cycle' of length 2.

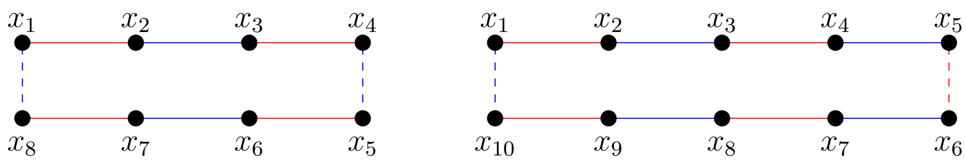

Figure 6.2: Edges of colour i are in red and those of colour j are in blue. Notice that the situation is different for even and odd k.

Proposition 6.2.6. *Let $([n], \triangleright)$ be a kei and let $i, j \in [n]$ with $i \neq j$. Suppose that the graph G_{ij} contains an ij-path with l vertices (with edges alternating in colour between j and i). Then the following hold:*

1. *For any component P of G_{ij} isomorphic to a path, $v(P)$ divides l;*
2. *For any component C of G_{ij} isomorphic to a cycle, $v(C)/2$ divides l.*

Proof. Let Q be the path $i = v_1, \ldots, v_l = j$ given in the statement. As i is fixed by f_i, the edge v_1v_2 is of colour j; similarly the edge $v_{l-1}v_l$ is of colour i. As the edges alternate in colour along the path, there must be an even number (at least two) of edges and so an odd number of vertices; write $l = 2l' + 1$ for some $l' \geqslant 1$. Hence there are l' edges of each colour and so from our construction $j = (i)(f_jf_i)^{l'}$. Hence

$$f_j = (f_if_j)^{l'} f_i (f_jf_i)^{l'} = (f_if_j)^{2l'} f_i = (f_if_j)^{l-1} f_i,$$

and thus we have $(f_if_j)^l = \iota$, and so $(f_jf_i)^l = \iota$ as well.

Now let C be a component of G_{ij} isomorphic to a cycle; label the vertices around the cycle as $x_1, \ldots, x_{2k}$ such that the edge x_1x_2 is of colour i. It is easy to see by construction that $(x_1)(f_if_j)^l = x_{2l+1}$, where we reduce the index modulo $2k$; hence $x_1 = x_{2l+1}$ and so $2l + 1 \equiv 1 \pmod{2k}$. Then $2k$ divides $2l$ and thus $k = v(C)/2$ divides l. Now let P be a component of G_{ij} isomorphic to a path; label the vertices along the path as $x_1, \ldots, x_k$ such that the edge x_1x_2 is of colour i. Now define vertices $x_{k+1}, \ldots, x_{2k}$ to be duplicates of the vertices of P by setting $x_s = x_{2k+1-s}$ for $k + 1 \leqslant s \leqslant 2k$ (see Figure 6.2). Now note that the edges of colour i on P are precisely the set $x_{2r-1}x_{2r}$ for $1 \leqslant r \leqslant \lfloor k/2 \rfloor$; as $x_s = x_{2k+1-s}$ we also have that the edges $x_{2(k-r)+1}x_{2(k-r)+2}$ are of colour i (for the same range of k). Relabelling the index to $k - r + 1$ and adjusting the range accordingly gives that the edges

$$\{x_{2r-1}x_{2r} \mid 1 \leqslant r \leqslant \lfloor k/2 \rfloor, \lceil k/2 \rceil + 1 \leqslant r \leqslant k\} \tag{6.2.2}$$

are all of colour i. Similarly, the edges of colour j on P are precisely the set $x_{2r}x_{2r+1}$ for $1 \leqslant r \leqslant \lfloor (k-1)/2 \rfloor$; as before the edges $x_{2(k-r)}x_{2(k-r)+1}$ are of colour j as well. Relabelling the index to $k-r$ and adjusting the range gives that the edges

$$\{x_{2r}x_{2r+1} \mid 1 \leqslant r \leqslant \lfloor (k-1)/2 \rfloor, \lceil (k+1)/2 \rceil \leqslant r \leqslant k-1\}$$

are all of colour j. Note also that by construction x_1 is not incident to an edge of colour j, so $(x_1)f_j = x_1 = x_{2k}$. So if we write $x_{2k+1} = x_1$, we have an 'edge' of colour j between x_{2k} and x_{2k+1}, and so the edges

$$\{x_{2r}x_{2r+1} \mid 1 \leqslant r \leqslant \lfloor (k-1)/2 \rfloor, \lceil (k+1)/2 \rceil \leqslant r \leqslant k\} \tag{6.2.3}$$

are all of colour j.

First suppose that k is even. Then from (6.2.2), each edge $x_{2r-1}x_{2r}$, $1 \leqslant r \leqslant k$, is of colour i, and from (6.2.3), the edges $x_2x_3, \dots, x_{k-2}x_{k-1}, x_{k+2}x_{k+3}, \dots, x_{2k}x_1$ are all of colour j. But there are an odd number of edges on P, and thus x_k is incident to an edge of colour i but not of colour j. Hence $(x_k)f_j = x_k = x_{k+1}$ and x_kx_{k+1} is of colour j as well.

Now suppose that k is odd. Then from (6.2.3), each edge $x_{2r}x_{2r+1}$, $1 \leqslant r \leqslant k$ is of colour j, and from (6.2.2), the edges $x_1x_2, \dots, x_{k-2}x_{k-1}, x_{k+2}x_{k+3}, \dots, x_{2k-1}x_{2k}$ are of colour i. But in this case x_k is not incident to an edge of colour i, so x_kx_{k+1} is of colour i as well.

So in both cases the duplicated path behaves like the cycle C, and thus $x_1 = x_{2l+1}$, reducing the index modulo $2k$. As we also have $x_{2k} = x_1$, we have $2l+1 \equiv 0, 1 \pmod{2k}$, and so as before we have that k divides l. □

6.2.3 Dihedral kei

We now turn our attention to a specific class of kei.

Definition 6.2.7. Let n be a positive integer and define a binary relation $\triangleright$ on $[n]$ by

$$i \triangleright j = 2j - i \pmod{n}.$$

The kei $D_n = ([n], \triangleright)$ is the *dihedral kei* on n elements. As before, we will usually consider the associated maps $(g_j)_{j=1}^n$, which in this case are given by $(i)g_j = 2j - i \pmod{n}$[1].

[1]These *specific* maps are labelled g_j to distinguish them from *general* f_j.

Proposition 6.2.8. *$D_n = ([n], \triangleright)$ is indeed a kei. Further, if n is odd, the following hold for any $i, j \in [n]$:*

1. *j is the only fixed point of g_j;*
2. *$g_i \neq g_j$ for $i \neq j$ (and so all of the associated maps are distinct);*
3. *If, for some k, $(k)g_i = (k)g_j$, then $i = j$.*

Proof. It is easy to see that each g_j is a bijection and that $(j)g_j = j$ and $g_j^2 = \iota$ for all j. For the rack condition (4.1.1), let k be arbitrary; then, working modulo n,

$$(k)g_j g_i g_j = (2j - k)g_i g_j = (2i - 2j + k)g_j = 4j - 2i - k. \tag{6.2.4}$$

On the other hand,

$$(k)g_{(i)g_j} = (k)g_{2j-i} = 4j - 2i - k,$$

and thus $(k)g_j g_i g_j = (k)g_{(i)g_j}$. As k was arbitrary all of the conditions to be a kei are fulfilled.

Now suppose that n is odd and $(i)g_j = i$ for some i, j. Then $2j - i \equiv i \pmod{n}$, i.e. $n|2(i - j)$. As n is odd, we have $n|(i - j)$ and by considering the range of i and j we have $i = j$. Hence j is the only fixed point of g_j. Now suppose $g_i = g_j$; then $i = (i)g_i = (i)g_j$, so $i = j$ and hence the maps $g_1, \ldots, g_n$ are all distinct. Finally, suppose $(k)g_i = (k)g_j$; then $2i - k \equiv 2j - k \pmod{n}$ and after cancelling the ks we again have $i = j$. $\square$

Note that properties 1 and 3 in Proposition 6.2.8 ensure that when n is odd the graphical representation of D_n is particularly nice. As each g_j has only one fixed point, there are $(n - 1)/2$ edges of each colour. Further, no pair of vertices is joined by two edges (of different colours); if this were the case then we would have $(k)g_i = (k)g_j$ for $i \neq j$. Thus there are $n(n - 1)/2$ distinct pairs joined by an edge, and so every pair is joined by precisely one edge. As each vertex is incident to at most one edge of any colour, we obtain a proper n-edge-colouring of K_n in this way (see Figure 6.3).

Proposition 6.2.9. *Let n be odd and suppose $i, j \in [n]$ are such that $i < j$ and $j - i$ is coprime with n. Then the subgraph G_{ij} of D_n is a path $i = v_1, \ldots, v_n = j$ with edges alternating in colour between j and i.*

Proof. As each map only has one fixed point, the vertex i is incident to an edge of colour i but not to an edge of colour j while vertex j is incident to an edge of colour j but not to an edge of colour i; every other vertex is incident to precisely one edge of

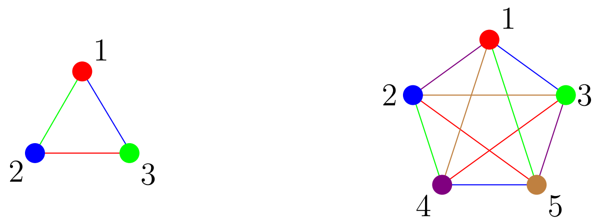

Figure 6.3: Graphical representations of D_3 (left) and D_5 (right).

each colour. So G_{ij} has degree sequence $(1, 1, 2, \ldots, 2)$; hence the graph consists of a collection of disjoint cycles and a single path P between vertices i and j. It remains to show that there are in fact no cycles in the graph; then every vertex is on P and we know that edges must alternate in colour.

Label the vertices of P as $i = v_1, \ldots, v_l = j$ for some odd $l \leqslant n$. Write $l = 2l' + 1$ for some integer l'; we will show by induction that $v_{2k+1} \equiv (2k+1)i - 2kj \pmod{n}$ for $0 \leqslant k \leqslant l'$. The base case $k = 0$ is clear from the definition, so take $k > 0$ and assume true for smaller k.

Now from our construction, there is an edge of colour j between v_{2k-1} and v_{2k} and an edge of colour i between v_{2k} and v_{2k+1}, so $v_{2k+1} = (v_{2k-1})g_j g_i$. Then from the first part of (6.2.4) and the inductive hypothesis we have (working modulo n)

$$v_{2k+1} = 2i - 2j + v_{2k-1} = 2i - 2j + (2k-1)i - (2k-1)j = (2k+1)i - 2kj,$$

and we have this result by induction. Then setting $k = l'$ we have $j = v_l \equiv li - (l-1)j \pmod{n}$, so that $li \equiv lj \pmod{n}$ and $n | l(j - i)$. As $j - i$ and n are coprime, we must have $n | l$; as $l \leqslant n$ we have $l = n$ as required. $\square$

This shows that D_n is connected for all odd n; for example, take $i = 1$ and $j = 2$.

Now consider the case where n is an odd prime p. Then for any $i, j \in [p]$ with $i \neq j$, p and $j - i$ are coprime, and so D_p is generated by *any* two distinct elements. As G_{ij} consists of a single component of size p for any $i \neq j$, we have from Proposition 5.1.4 that g_i and g_j don't commute for any $i \neq j$.

We will now show that any connected kei on p elements is isomorphic to D_p. The result is a consequence of the following theorem, which was first shown by Etingof, Soloviev and Guralnick in [7] and later recast in more familiar language (see [11] for example); recall from Definition 4.1.5 that for an abelian group A and $\tau \in \mathrm{Aut}(A)$, the Alexander quandle (A, τ) is defined by the binary operation $x \triangleright y = (x - y)\tau + y$.

Theorem 6.2.10. *Let p be a prime number. Then any connected quandle on p elements is isomorphic to an Alexander quandle $(\mathbb{Z}_p, \tau)$ for some $\tau \in Aut(\mathbb{Z}_p) \setminus \{\iota\}$.* $\square$

Corollary 6.2.11. *Let p be a prime number. Then any connected kei on p elements is isomorphic to D_p.*

Proof. First consider the Alexander quandles $(\mathbb{Z}_p, \tau)$ as described. $\mathrm{Aut}(\mathbb{Z}_p)$ is isomorphic to the multiplicative group of the field $\mathbb{F}_p$, so each automorphism τ can be realised as multiplication by a non-zero element c, where ι corresponds to multiplication by 1. Hence for $x, y \in \mathbb{Z}_p$ we have $x \triangleright y = c(x-y)+y$, and $(x \triangleright y) \triangleright y = c^2(x-y)+y$, so $(x \triangleright y) \triangleright y = x$ for all x (i.e. we have a kei) if and only if $c^2 = 1$, i.e. $c = \pm 1$.

Now let K be a connected kei on p elements; from Theorem 6.2.10 it is isomorphic to an Alexander quandle as described above, with $c \neq 1$. The only remaining kei has $c = -1$, and in this case $x \triangleright y = -(x-y)+y = 2y-x$. Hence $K \cong (\mathbb{Z}_p, -\iota) \cong D_p$. $\square$

6.3 Classifying kei with prime orbits

6.3.1 Kei such that all orbits have size p

Let p be an odd prime and suppose we have prescribed the final block structure by setting $\mathcal{C}([n]) = \{B_1, \ldots, B_m\}$ where each 'block' B_i is of size p. Throughout we will put $\phi_x^{(j)} = \phi_x^{(B_j)}$ for any $x \in [n]$, $j \in [m]$.

We begin by showing how the maps $(f_x)_{x \in B_i}$ can act on B_i itself, as well as on any other block B_j.

Lemma 6.3.1. *Let B_i and B_j be two distinct 'blocks', each of size p. Then $(B_i, \triangleright)$ is either trivial or isomorphic to D_p. Further:*

1. *If $(B_i, \triangleright) = T_{B_i}$ then $\phi_x^{(j)} = \iota$ for all $x \in B_i$;*
2. *If $(B_i, \triangleright) \cong D_p$ then either $\phi_x^{(j)} = \iota$ for all $x \in B_i$ or there exists a bijection $\chi_{ij} : B_i \to B_j$ such that*
$$\phi_x^{(j)} = \chi_{ij}^{-1} \circ \phi_x^{(i)} \circ \chi_{ij}$$
for all $x \in B_i$.

Proof. By applying Lemma 6.2.4 with the singleton set $\{i\}$ we see that any two B_i-orbits contained in B_i itself must have the same size. Let s_i be this common size; it is clear that $s_i | p$ and thus $s_i = 1$ or p. If $s_i = 1$ then each orbit is of size one and so the subkei $(B_i, \triangleright) = T_{B_i}$; If $s_i = p$ then there is only a single orbit and thus $(B_i, \triangleright)$ is connected; from Corollary 6.2.11 it follows that $(B_i, \triangleright)$ is isomorphic to the dihedral kei D_p.

First suppose that $(B_i, \triangleright)$ is trivial. Then $(y)f_x = y$ for all $x, y \in B_i$ and it follows that all the maps corresponding to elements of B_i mutually commute; thus (from Proposition 5.1.4) each B_i-orbit has size equal to a power of 2. But by Lemma 6.2.4, again applied to the singleton set $\{i\}$, any two B_i-orbits contained in B_j must have the same size, and this size divides p. Hence each orbit is of size 1 and each map f_x acts trivially on B_j.

Now suppose that $(B_i, \triangleright) \cong D_p$. Let $y, z \in B_i$ with $y \neq z$; from Proposition 6.2.9 there is a yz-path P of length p in G_{yz} (spanning all the vertices in B_i) with edges alternating in colour between z and y, and thus $(B_i, \triangleright)$ is generated by y and z. Thus the components of G_{B_i} coincide with the components of G_{yz}, and from Proposition 6.2.6 the remaining components of G_{yz} are paths with 1 or p vertices and cycles of length 2 or $2p$. But from Lemma 6.2.4 again, any two B_i-orbits (i.e. vertex sets of components of G_{B_i} and thus of G_{yz}) contained in B_j have the same size dividing p, and thus there are either p orbits of size 1 or 1 orbit of size p. In the first case, f_x acts trivially on B_j for each $x \in B_i$.

In the second case, we have from Proposition 6.2.6 that there is a path $w_1, \ldots, w_p$ in G_{yz} spanning all the vertices of B_j, with edges alternating in colour between z and y (say the first edge is of colour z). List the vertices along the original path P as $y = v_1, \ldots, v_p = z$, and define a bijection $\chi = \chi_{ij} : B_i \to B_j$ by $v_r \mapsto w_r$ for each r. Now note that by our construction

$$\phi_y^{(i)} = (v_2\ v_3) \cdots (v_{p-1}\ v_p)$$

while

$$\phi_y^{(j)} = (w_2\ w_3) \cdots (w_{p-1}\ w_p) = ((v_2)\chi\ (v_3)\chi) \cdots ((v_{p-1})\chi\ (v_p)\chi)$$

and similarly

$$\phi_z^{(i)} = (v_1\ v_2) \cdots (v_{p-2}\ v_{p-1})$$

while

$$\phi_z^{(j)} = (w_1\ w_2) \cdots (w_{p-2}\ w_{p-1}) = ((v_1)\chi\ (v_2)\chi) \cdots ((v_{p-2})\chi\ (v_{p-1})\chi),$$

so $\phi_y^{(j)} = \chi^{-1} \circ \phi_y^{(i)} \circ \chi$ and $\phi_z^{(j)} = \chi^{-1} \circ \phi_z^{(i)} \circ \chi$. Now for $0 \leqslant r \leqslant (p-1)/2$, we have that $v_{2r+1} = (y)(f_z f_y)^r$, so $f_{v_{2r+1}} = (f_y f_z)^r f_y (f_z f_y)^r$. In the block B_i this gives

$$\phi_{v_{2r+1}}^{(i)} = \left(\phi_y^{(i)} \phi_z^{(i)}\right)^r \phi_y^{(i)} \left(\phi_z^{(i)} \phi_y^{(i)}\right)^r$$

and in the block B_j, noting the cancellations,

$$\begin{aligned}\phi^{(j)}_{v_{2r+1}} &= \left(\phi^{(j)}_y\phi^{(j)}_z\right)^r \phi^{(j)}_y \left(\phi^{(j)}_z\phi^{(j)}_y\right)^r \\ &= \chi^{-1}\left(\phi^{(i)}_y\phi^{(i)}_z\right)^r \phi^{(i)}_y \left(\phi^{(i)}_z\phi^{(i)}_y\right)^r \chi \\ &= \chi^{-1}\circ\phi^{(i)}_{v_{2r+1}}\circ\chi\end{aligned}$$

Similarly, for $0 \leqslant r \leqslant (p-3)/2$, $v_{2r+2} = (v_{2r+1})f_z$, so pre- and post-multiplying all of the above by f_z gives the result for even-indexed vertices as well. Hence $\phi^{(j)}_x = \chi_{ij}^{-1}\circ\phi^{(i)}_x\circ\chi_{ij}$ for all $x \in B_i$. □

We will now introduce some useful terminology. In the case where $(B_i, \triangleright) = T_{B_i}$, we will say B_i *doesn't act* on itself; if $(B_i, \triangleright) \cong D_p$ then we will say B_i *acts* on itself. Similarly, if $\phi^{(j)}_x = \iota$ for each $x \in B_i$, we will say that B_i *doesn't act* on B_j, and if a map χ_{ij} exists as described above we will say B_i *acts* on B_j *by* χ_{ij}. We have shown that a block which doesn't act on itself can't act on any other blocks, and if a block B_i acts on itself and on another block B_j, the action on B_j is identical to the action of B_i on itself up to a relabelling of the vertices (in a graphical sense), where the relabelling is given by the map χ_{ij}.

We will now need two technical lemmas regarding how blocks can act on each other.

Lemma 6.3.2. *Let B_i and B_j be two distinct 'blocks', each of size p. Suppose B_i and B_j both act on themselves, and that B_i acts on B_j. Then for every $x \in B_i$ there exists $s \in B_j$ such that f_x and f_s agree on B_j, i.e. $\phi^{(j)}_x = \phi^{(j)}_s$.*

Proof. Let $x \in B_i$ and choose a pair of distinct vertices $q, r \in B_j$ such that $(q)f_x = r$; such a pair certainly exists as B_i acts on B_j. The fact that it transposes q and r is the only property of f_x that will be used in this argument.

As B_j acts on itself, there is a path $q = w_1, \ldots, w_p = r$ in G_{qr} (spanning all the vertices of B_j) with edges alternating in colour between r and q. We will show by induction that for $1 \leqslant k \leqslant (p+1)/2$, $(w_k)f_x = w_{p+1-k}$, and so in particular that $s = w_{(p+1)/2}$ is the only fixed point of f_x in B_j (see Figure 6.4). The base case $k = 1$ is true by construction, so take $k > 1$ and assume the result for smaller k. Note that as $(q)f_x = r$ we have $f_r = f_x f_q f_x$, which can also be expressed as $f_x = f_q f_x f_r = f_r f_x f_q$. If k is even then $p+2-k$ is odd, so the edge $w_{k-1}w_k$ is of colour r while the edge $w_{p+1-k}w_{p+2-k}$ is of colour q. Then as $p+2-k = p+1-(k-1)$ we have by the inductive hypothesis that

$$(w_k)f_x = (w_k)f_r f_x f_q = (w_{k-1})f_x f_q = (w_{p+2-k})f_q = w_{p+1-k}. \tag{6.3.1}$$

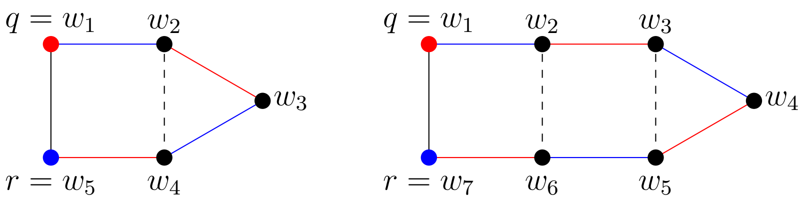

Figure 6.4: Edges of colour q are in red and those of colour r are in blue. If we know any map transposes q and r we can determine the rest of the map on the block B_j. Notice that the path between q and the central vertex is slightly different for $p \equiv 1$ or $3 \pmod 4$.

Similarly, if k is odd then $w_{k-1}w_k$ is of colour q and $w_{p+1-k}w_{p+2-k}$ is of colour r, so we can use the other identity for f_x to get the same result. Thus this result follows by induction.

Now note that s is at distance $(p-1)/2$ from q on the path. If $(p-1)/2$ is even (i.e. $p \equiv 1 \pmod 4$) then we have $s = (q)(f_r f_q)^{(p-1)/4}$, so

$$f_s = (f_q f_r)^{(p-1)/4} f_q (f_r f_q)^{(p-1)/4} = f_q (f_r f_q)^{(p-1)/2},$$

while if $(p-1)/2$ is odd (i.e. $p \equiv 3 \pmod 4$) then $s = (q)(f_r f_q)^{(p-3)/4} f_r$, so

$$f_s = f_r (f_q f_r)^{(p-3)/4} f_q (f_r f_q)^{(p-3)/4} f_r = (f_r f_q)^{(p-1)/2} f_r.$$

In both cases we have that $(q)f_s = r$, as q and r are at distance $p-1$ in G_{qr}. In the earlier argument we only used the fact that $(x)f_q = r$; thus we can apply the same argument to f_s in place of f_x and deduce that f_s and f_x agree on B_j. $\square$

Corollary 6.3.3. *Let B_i, B_j, x and s be as in Lemma 6.3.2. Then $s = (x)\chi_{ij}$, $\chi_{ji} = \chi_{ij}^{-1}$, and f_x and f_s agree on $B_i \cup B_j$.*

Proof. As f_x and f_s agree on B_j and $(B_j, \triangleright) \cong D_p$, s is the only fixed point of $\phi_x^{(j)} = \phi_s^{(j)}$. As mentioned after Lemma 6.3.1, the action of B_i on B_j is identical to the action of B_i on itself up to a relabelling of the vertices by the map χ_{ij}; as x is the only fixed point of $\phi_x^{(i)}$, $(x)\chi_{ij}$ is the only fixed point of $\phi_x^{(j)}$, and so $s = (x)\chi_{ij}$.

Now consider the action of B_j on B_i. If B_j doesn't act on B_i then $(y)f_t = y$ for all $y \in B_i$, $t \in B_j$ and so in particular f_x must commute with f_t for any $t \in B_j$. Choose such a $t \neq s$; we know that $\phi_s^{(j)}$ and $\phi_t^{(j)}$ don't commute as $(B_j, \triangleright) \cong D_p$. As f_x and f_s agree on B_j, $\phi_x^{(j)}$ and $\phi_t^{(j)}$ don't commute, which is a contradiction.

Hence B_j acts on B_i, and reversing i and j we have that there exists some $y = f_{(s)\chi_{ji}}$ such that f_s and f_y agree on B_i. But as f_x and f_s agree on B_j, we have that $(s)f_x = (s)f_s = s$, so f_x and f_s commute. So suppose $y \neq x$; then as above $\phi_s^{(i)} = \phi_y^{(i)}$

doesn't commute with $\phi_x^{(i)}$, a contradiction. So $(s)\chi_{ji} = x$ and as x was arbitrary we conclude that $\chi_{ji} = \chi_{ij}^{-1}$. We also have that f_s and f_x agree on B_i (as well as B_j), completing the proof. □

Lemma 6.3.4. *Let B_i and B_j be distinct 'blocks', each of size p. Suppose B_i and B_j both act on themselves, and that B_i acts on B_j by χ_{ij}. Let $k \in [m]$ be such that B_i acts on B_k by χ_{ik}. Then B_j acts on B_k, and more specifically:*

1. *For any $x \in B_i$, the maps f_x and $f_{(x)\chi_{ij}}$ agree on B_k, i.e. $\phi_x^{(k)} = \phi_{(x)\chi_{ij}}^{(k)}$;*
2. *For $k \neq i$, $\chi_{jk} = \chi_{ji}\chi_{ik} = \chi_{ij}^{-1}\chi_{ik}$.*

Proof. Note that Corollary 6.3.3 shows item 1 in the case where $k = i$, so we need only consider $k \neq i$.

Let $x \in B_i$ be arbitrary and put $s = (x)\chi_{ij}$. Choose a pair of distinct vertices $y, z \in B_i$ such that $(y)f_x = z$ and consider the path $y = v_1, \ldots, v_p = z$, with edges alternating in colour between z and y. As in the proof of Lemma 6.3.2, $(v_r)f_x = v_{p+1-r}$ for $1 \leqslant r \leqslant (p+1)/2$, so $x = v_{(p+1)/2}$. Now B_i acts on B_k, so from Lemma 6.3.1, $(v_1)\chi_{ik}, \ldots, (v_p)\chi_{ik}$ is a path, spanning all the vertices of B_k, with edges alternating in colour between z and y. Write $u_r = (v_r)\chi_{ik}$ for each r, so from the proof of Lemma 6.3.2, $(u_r)f_x = u_{p+1-r}$ for $1 \leqslant r \leqslant (p+1)/2$. We will show by *reverse* induction on r that $(u_r)f_s = u_{p+1-r}$ for $1 \leqslant r \leqslant (p+1)/2$, i.e. that $\phi_x^{(k)} = \phi_s^{(k)}$.

Let us start with the base case $r = (p+1)/2$. From Corollary 6.3.3, f_x and f_s agree on $B_i \cup B_j$, and so in particular f_x and f_s commute. Now the only fixed point of $\phi_x^{(k)}$ is $u_{(p+1)/2}$, which is to say that $G_x[B_k]$ consists of one isolated vertex and $(p-1)/2$ disjoint edges. Since f_x and f_s commute, the graph $G_{xs}[B_k]$ can consist only of cycles of length 2 or 4 and paths with 1 or 2 vertices, and so it follows that $u_{(p+1)/2}$ is isolated in this graph as well, i.e. $(u_{(p+1)/2})f_s = u_{(p+1)/2}$.

Now take $r < (p+1)/2$ and assume the result for larger r. As f_x and f_s agree on B_i, $(y)f_s = (y)f_x = z$, so $f_sf_yf_s = f_z$. This allows us to use an argument similar to (6.3.1) in the proof of Lemma 6.3.2; if r is even then $p - r$ is odd, so the edge u_ru_{r+1} is of colour y while the edge $u_{p-r}u_{p+1-r}$ is of colour z. Then from the inductive hypothesis,

$$(u_r)f_s = (u_r)f_yf_sf_z = (u_{r+1})f_sf_z = (u_{p-r})f_z = u_{p+1-r}.$$

Similarly, if r is odd then u_ru_{r+1} is of colour z and $u_{p-r}u_{p+1-r}$ is of colour y, so we get the same result. The claim then follows by reverse induction.

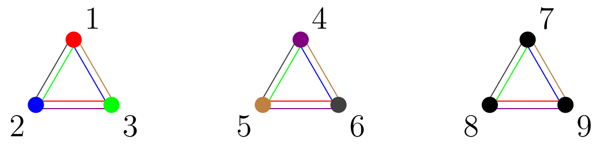

Figure 6.5: A graphical representation of an irreducible kei on $\{1, \dots, 9\}$, with blocks $B_1 = \{1, 2, 3\}$, $B_2 = \{4, 5, 6\}$ and $B_3 = \{7, 8, 9\}$. As there are no edges of colour 7,8 or 9, these vertices are given no colour.

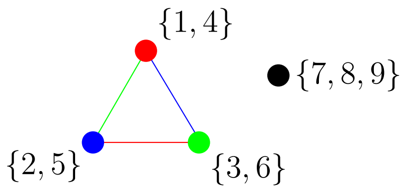

Figure 6.6: The reduced kei of the kei in Figure 6.5, on elements $C_1 = \{1, 4\}$, $C_2 = \{2, 5\}$, $C_3 = \{3, 6\}$ and $C_4 = \{7, 8, 9\}$. The last row of the corresponding matrix Θ is $((2\ 3), (1\ 3), (1\ 2), \iota)$, mimicking the actions of the original f_1, f_2 and f_3 on $\{7, 8, 9\}$ under the canonical bijection from $\{7, 8, 9\}$ to $\{1, 2, 3\}$.

This shows item 1. For item 2, note that as $u_{(p+1)/2}$ is the only fixed point of $\phi_s^{(k)} = \phi_x^{(k)}$, it follows that $(x)\chi_{ik} = u_{(p+1)/2} = (s)\chi_{jk} = (x)\chi_{ij}\chi_{jk}$. As x was arbitrary we have $\chi_{ik} = \chi_{ij}\chi_{jk}$, i.e. $\chi_{jk} = \chi_{ij}^{-1}\chi_{ik} = \chi_{ji}\chi_{ik}$. $\square$

We can now classify all irreducible kei with this block structure, where we will assign specific labels to these blocks for clarity. We use the notion of reduced kei from Section 6.1; each such kei has a very specific reduced form (see Figures 6.5 and 6.6 for an example when $p = 3$). Recall that for a partition $\mathcal{P}$ of $[n]$, $\mathcal{IK}_{\mathcal{P}}$ denotes the set of irreducible kei on $[n]$ with final block structure $\mathcal{P}$.

Proposition 6.3.5. *Let p be an odd prime and $n = mp$ for some integer m. Let $\mathcal{P} = \{B_1, \dots, B_m\}$ be a partition of $[n]$ where $B_i = \{(i-1)p + 1, \dots, ip\}$ for all i, and suppose $K = ([n], \triangleright) \in \mathcal{IK}_{\mathcal{P}}$. Then $K = [\mathcal{C}, \triangleright', \Theta]$ for some triple satisfying one of the two sets of conditions listed below. Either*

1a. $\mathcal{C} = \{C_1, \dots, C_p\}$ where each C_i contains precisely one element from each B_j, $1 \leqslant j \leqslant m$, with $i \in C_i$ for all i;

1b. $(\mathcal{C}, \triangleright') \cong D_p$;

1c. $\theta_{ij} = \iota$ for all i, j, *or there exists some $1 \leqslant r < m$ such that*

2a. *$\mathcal{C} = \{C_1, \ldots, C_{p+1}\}$ where each C_i, $1 \leqslant i \leqslant p$, contains precisely one element from each B_j, $1 \leqslant j \leqslant r$, and $i \in C_i$ for all i;*

2b. *$K' := (\{C_1, \ldots, C_p\}, \triangleright') \cong D_p$, and $(\mathcal{C}, \triangleright') = K' \oplus T_{\{C_{p+1}\}}$;*

2c. *$\theta_{ij} = \iota$ for $i \leqslant p$, $\theta_{p+1,p+1} = \iota$ and*

$$\theta_{p+1,j} = \rho_{p+1}^{-1} \circ \phi_j^{(E)} \circ \rho_{p+1}$$

for $j \leqslant p$, where $E = \bigcup_{i=r+1}^{m} B_i$ and $\rho_{p+1} : E \to [(m-r)p]$ is the canonical bijection as defined in Section 6.1.1.

Proof. From Lemma 6.3.1, if no block acts on itself then $f_x = \iota$ for all $x \in [n]$, and so $([n], \triangleright)$ is the trivial kei $T_{[n]}$ with block structure consisting of all the singletons, a contradiction. Hence there is at least one block which acts on itself; after renumbering if necessary we can take this to be B_1. First define two subsets of $[m]$; let

$$a = \{j \in [m] \mid B_1 \text{ acts on } B_j, B_j \text{ acts on itself}\},$$

so $1 \in a$, and let

$$e = \{k \in [m] \mid B_1 \text{ acts on } B_k, B_k \text{ doesn't act on itself}\}.$$

Put $A = \bigcup_{j \in a} B_j$ and $E = \bigcup_{k \in e} B_k$; we first show that $a \cup e = [m]$, and thus that $A \cup E = [n]$.

Suppose that this is not the case, so there exists some $t \in [m] \setminus (a \cup e)$. Then by definition B_1 does not act on B_t and thus B_j does not act on B_t for any $j \in a$ (or, noting that B_1 and B_j act on each other, we would have a contradiction to Lemma 6.3.4). Suppose now that B_t acts on B_j for some $j \in a \cup e$, and let $y, z \in B_1$ be distinct. As B_1 doesn't act on B_t, we have $(w)f_y = w = (w)f_z$ for any $w \in B_t$, and so f_w commutes with f_y and f_z. Now B_1 acts on B_j, so $(y)\chi_{1j}$ is the only fixed point of $\phi_y^{(j)}$. As in the proof of Lemma 6.3.4, this must also be the only fixed point of $\phi_w^{(j)}$, and as f_w commutes with f_z, it also follows that $(y)\chi_{1j}$ is the only fixed point of $\phi_z^{(j)}$. But this is a contradiction; $(z)\chi_{1j} \neq (y)\chi_{1j}$ is the only fixed point of $\phi_z^{(j)}$. Hence $A \cup E$ and $[n] \setminus (A \cup E)$ are both stable subkei, which is a contradiction to the irreducibility assumption.

Now (again up to renumbering) we may take $a = [r]$ (for some $1 \leqslant r \leqslant m$) and $e = \{r+1, \ldots, m\}$, so $A = [rp]$ and $E = \{rp+1, \ldots, n\}$. As B_1 acts on every other block and B_j acts on itself for $1 \leqslant j \leqslant r$, we have from Lemma 6.3.4 that each such B_j acts on every B_k, and that the corresponding maps χ_{jk} are related to the maps χ_{1k}. In particular, for every $i \in B_1 = [p]$ and $2 \leqslant j \leqslant r$, the maps f_i and $f_{(i)\chi_{1j}}$ agree on the whole of $[m]$. Now define, for $i \in [p]$, the set

$$C_i = \{i\} \cup \{(i)\chi_{1j} \mid 2 \leqslant j \leqslant r\},$$

and if $E \neq \emptyset$ (i.e. $r < m$), put $C_{p+1} = E$. The sets $C_1, \ldots, C_p$ partition $A = [rp]$ as each χ_{1j} is a bijection from $B_1 = [p]$ to B_j; hence each C_i contains precisely one element from each B_j.

Now for any i and any $x, y \in C_i$, we have $f_x = f_y$; note this is true for $i = p+1$ as E contains all elements x such that $f_x = \iota$. Now write, for $i \in [p]$, $g_i = \phi_i^{(1)}$; as $(B_1, \triangleright) \cong D_p$ the maps $(g_i)_{i=1}^p$ are all distinct and not equal to $g_{p+1} := \iota$. So for $i \neq j$, $x \in C_i$ and $y \in C_j$, $\phi_x^{(1)} = \phi_i^{(1)} = g_i \neq g_j = \phi_j^{(1)} = \phi_y^{(1)}$, and so $\mathcal{C}$ is the partition into the equivalence classes of $\sim$. Note that as $i \in C_i$ for $1 \leqslant i \leqslant p$, we may write $C_i = [[i]]$.

Now consider the reduced kei action $\triangleright'$ on $\mathcal{C}$. For $i, j \in [p]$, $[[i]] \triangleright' [[j]] = [[i \triangleright j]] = [[(i)g_j]]$, i.e. $C_i \triangleright' C_j = C_{(i)g_j}$. Hence the set $\{C_1, \ldots, C_p\}$ forms a subkei under $\triangleright'$, isomorphic to $(B_1, \triangleright) \cong D_p$ under the map $C_i \mapsto i$. If $r = m$, this forms the entire reduced kei and we have shown items 1a and 1b; otherwise we are in case 2 and must also consider C_{p+1}. As f_{C_i} is a bijection for $i \leqslant p$, $(C_{p+1})f_{C_i} = C_{p+1}$ for all such i; we also have, for $y \in C_{p+1}$, $C_i \triangleright C_{p+1} = [[i]] \triangleright' [[y]] = [[i \triangleright y]] = [[i]] = C_i$, as f_y acts as the identity. Thus $f_{C_{p+1}}$ acts as the identity on $\mathcal{C}$, and this gives us the form of the reduced kei $(\mathcal{C}, \triangleright')$, showing items 2a and 2b.

Now consider the matrix θ; we can determine its elements by finding the maps σ_{ij} as in (6.1.2). Fix $j \leqslant p$ and consider the map $\hat{f}_{C_j}$; for $i \in [p]$ and $1 \leqslant s \leqslant |C_i| = r$, $C_i\langle s\rangle$ is mapped onto $C_{(i)g_j}\langle s\rangle$ by definition. But note that each C_i contains one element from each of $B_1, \ldots, B_r$, and by definition of the blocks these elements are in increasing order. As $C_i\langle s\rangle$ is the sth smallest element of C_i, it is equal to the element of C_i from B_s, namely $(i)\chi_{1s}$ if $s \geqslant 2$ and i if $s = 1$. So for $s \geqslant 2$,

$$(i)\chi_{1s}\hat{f}_{C_j} = (C_i\langle s\rangle)\hat{f}_{C_j} = C_{(i)g_j}\langle s\rangle = (i)g_j\chi_{1s},$$

and as i was arbitrary, this gives $\chi_{1s}\hat{f}_{C_j} = g_j\chi_{1s}$; considering the domains and co-domains of these maps gives

$$\hat{f}_{C_j}|_{B_s} = \chi_{1s}^{-1} \circ g_j \circ \chi_{1s} = \chi_{1s}^{-1} \circ \phi_j^{(1)} \circ \chi_{1s} = \phi_j^{(s)} = f_j|_{B_s},$$

from Lemma 6.3.1. For $s = 1$ we simply have $(i)\hat{f}_{C_j} = (i)g_j$, and so $\hat{f}_{C_j}|_{B_1} = g_j = f_j|_{B_1}$; we can combine these results over all values of s from 1 to r to see $\hat{f}_{C_j}|_A = f_j|_A$. Hence $\hat{f}_{C_j} f_j|_A = \iota$ and as the classes $C_1, \ldots, C_p$ partition A, we have from (6.1.2) that $\sigma_{ij} = \iota$ for all $i \leqslant p$. As $j \leqslant p$ was arbitrary, $\theta_{ij} = \iota$ for all $1 \leqslant i, j \leqslant p$.

If $r = m$ (so $A = [n]$ and we are in case 1), this determines Θ entirely and we have shown item 1c; if $r < m$, we must also consider $C_{p+1} = E$, adding an extra row and column to the matrix. For $y \in E$, $f_y = \iota$, and so we see $\hat{f}_E = \iota$ as well; thus $\hat{f}_{C_{p+1}} f_y = \iota$ and $\sigma_{i,p+1} = \iota$ for all i. Now for $j \leqslant p$, $(E)f_{C_j} = E$, and so $\hat{f}_{C_j}|_{C_{p+1}} = \iota$. Hence

$$(\hat{f}_{C_j} f_j)|_{C_{p+1}} = \iota \circ f_j|_E = \phi_j^{(E)},$$

and thus $\sigma_{p+1,j} = \phi_j^{(E)}$. Defining the canonical bijection ρ_{p+1} on E as in Section 6.1.1, we obtain $\theta_{p+1,j}$, and so item 2c follows. □

This result shows that any kei $K \in \mathcal{IK_P}$ has a representation as a triple $[\mathcal{C}, \triangleright', \Theta]$ satisfying one of the two sets of conditions given. We will now work backwards to establish a one-to-one correspondence between $\mathcal{IK_P}$ and the triples satisfying one of these two sets of conditions.

First suppose we have a triple $[\mathcal{C}, \triangleright', \Theta]$ satisfying conditions 1a, 1b and 1c of Proposition 6.3.5. The conditions (6.1.5), (6.1.8) and (6.1.9) of Propositions 6.1.3 and 6.1.4 are trivially satisfied; further, as the maps $(f_{C_i})_{i=1}^p$ are all distinct, the extended maps $(\hat{f}_{C_i})_{i=1}^p$ are as well. Hence $K = [\mathcal{C}, \triangleright', \Theta]$ is a kei on $[n]$ and $\mathcal{C}$ is truly the set of equivalence classes of K under $\sim$. By examining the proof of Proposition 6.3.5, we can now obtain this particular triple from a kei $K \in \mathcal{IK_P}$ by setting $(j)\chi_{1s} = C_j\langle s\rangle$ for $s \geqslant 2$, and creating a copy of D_p on $B_1 = [p]$ (isomorphic to $(\mathcal{C}, \triangleright')$ using the natural bijection $C_i \mapsto i$).

Now suppose our triple satisfies conditions 2a and 2b, and that we also have $\theta_{ij} = \iota$ for $i \leqslant p$ and $\theta_{p+1,p+1} = \iota$. The only non-trivial elements of Θ are in the bottom row, where each entry is a member of $S_{(m-r)p}$, so it is easier to work with the original maps $\sigma_{p+1,j}$ lying in $\mathrm{Sym}(C_{p+1})$. The condition (6.1.8) is true; for (6.1.9), as $(C_{p+1})f_{C_j} = C_{p+1}$ for each j, we require $\sigma_{p+1,j} = \sigma_{p+1,j}^{-1}$ for each j. Thus $\sigma_{p+1,j}$ must be an involution for $1 \leqslant j \leqslant p$.

For (6.1.5), we need only put $\alpha = p + 1$; any other choice leads to a triviality. In this non-trivial case we have $\beta = \gamma = \delta = p + 1$ and we require

$$\sigma_{p+1,k} = \sigma_{p+1,j}\sigma_{p+1,i}\sigma_{p+1,j} \tag{6.3.2}$$

for all i, j, where $C_k = (C_i)f_{C_j}$. If $i = p+1$ then $k = p+1$ and $\sigma_{p+1,i} = \iota$ so (6.3.2) is true; similarly if $j = p+1$ then $i = k$ and this time $\sigma_{p+1,j} = \iota$, again satisfying (6.3.2). Hence any set of involutions $(\sigma_{p+1,j})_{j=1}^{p}$ (covering the remaining indices) satisfying (6.3.2) can be used to define a valid Θ and thus a kei K on $[n]$, and as $\hat{f}_{C_{p+1}}\sigma_{1,p+1}\cdots\sigma_{p+1,p+1} = \iota$, $\mathcal{C}$ is truly the set of equivalence classes under $\sim$. Note that as the maps $\sigma_{p+1,j}$ don't have to respect the block structure within C_{p+1}, we won't get a member of $\mathcal{K}_{\mathcal{P}}$ for each valid choice; it remains to find suitable restrictions.

Fix a $j \in [p]$ and take any $x \in C_j$; then $f_x = \hat{f}_{C_j}\sigma_{p+1,j}$. As $(C_{p+1})f_{C_j} = C_{p+1}$, $\hat{f}_{C_j}$ acts entirely on $A = \bigcup_{i=1}^{r} B_i$, and so $f_x|_{C_{p+1}} = \sigma_{p+1,j}$. So to obtain a kei $K \in \mathcal{IK}_{\mathcal{P}}$, we must have

$$\sigma_{p+1,j} = \phi_j^{(r+1)}\cdots\phi_j^{(m)}$$

for each j. Now as before, define $(j)\chi_{1s} = C_j\langle s\rangle$ for $2 \leqslant s \leqslant r$, and create a copy of D_p on $B_1 = [p]$ (isomorphic to $(\mathcal{C}, \triangleright')$ using the natural bijection $C_i \mapsto i$); then any K as in Proposition 6.3.5 will give the same reduced kei $(\mathcal{C}, \triangleright)$. We still need to consider the bijections $(\chi_{1l})_{l=r+1}^{m}$ determining how B_1 acts on the blocks $B_{r+1}, \ldots, B_m$.

As in the proof of Propostion 6.3.5, let $g_j := \phi_j^{(1)}$ for $j \in [p]$. Then after restricting to the block B_l, $r < l \leqslant m$ and noting the cancellations, (6.3.2) becomes

$$\chi_{1l}^{-1} \circ g_k \circ \chi_{1l} = \chi_{1l}^{-1} \circ g_j \circ g_i \circ g_j \circ \chi_{1l},$$

which is satisfied for *any* bijection χ_{1l} (note that $(C_j)f_{C_i} = C_k$ if and only if $(i)g_j = k$ as $(\mathcal{C}, \triangleright') \cong (B_1, \triangleright)$). Thus any choice of the bijections $(\chi_{1l})_{l=r+1}^{m}$ gives a valid Θ corresponding to K.

Finally, observe that any kei K as in Propostition 6.3.5 is genuinely irreducible. Suppose U and V form stable subkei and that U and V partition $[n]$, so $U = \bigcup_{i\in s} B_i$ for some $s \subseteq [m]$, while $V = \bigcup_{i\notin s} B_i$. Suppose without loss of generality that $1 \in s$; then as B_1 acts on every other block there are edges from U in $G_K[V]$, a contradiction.

We can thus count the number of irreducible kei on $\mathcal{P}$, but we will need the following result due to Elhamdadi, MacQuarrie and Restrepo [5].

Theorem 6.3.6. *Let p be a prime number. Then the number of automorphisms of D_p is $p(p-1)$.* □

Theorem 6.3.7. *Let p be an odd prime and $n = mp$ for some integer m. Let $\mathcal{P} = \{B_1, \ldots, B_m\}$ be a partition of $[n]$ with $|B_i| = p$ for all i. Then*

$$|\mathcal{IK}_{\mathcal{P}}| = \frac{(p!)^m}{p(p-1)} \cdot (2^m - 1).$$

Proof. We will label each $B_i = \{(i-1)p+1, \ldots, ip\}$ for clarity. From Proposition 6.3.5 and the following remarks, any $K \in \mathcal{IK}_{\mathcal{P}}$ is determined by the partition $\mathcal{C} = \{C_1, \ldots, C_p, C_{p+1}\}$, the set of blocks (of cardinality r) which act on themselves, the kei on $\{C_1, \ldots, C_p\}$ and, for $r < m$, the set of bijections defining the bottom row of Θ.

We can choose r to be any element of $[m]$; there are then $\binom{m}{r}$ ways to choose the set of blocks that act on themselves; for clarity we will consider them to be $B_1, \ldots, B_r$. There are $(p!)^{r-1}$ partitions $\{C_1, \ldots, C_p\}$ of $[rp]$ such that each C_i contains precisely one element from each B_j (the order does not matter), and from Theorem 6.3.6, there are $p!/p(p-1)$ ways of constructing a kei K' on this set isomorphic to D_p. Finally, there are $(p!)^{m-r}$ ways of choosing the bijections $(\chi_{1l})_{l=r+1}^{n}$. Thus summing over all values of r from 1 to m gives

$$|\mathcal{IK}_{\mathcal{P}}| = \sum_{r=1}^{m} \binom{m}{r} \frac{p!}{p(p-1)} (p!)^{r-1} (p!)^{m-r} = \frac{(p!)^m}{p(p-1)} \cdot (2^m - 1).$$

□

This result allows us to bound the number of kei on such a partition.

Corollary 6.3.8. *Let p be an odd prime and $n = mp$ for some integer m. Let $\mathcal{P} = \{B_1, \ldots, B_m\}$ be a partition of $[n]$ where $|B_i| = p$ for all i. Then $|\mathcal{K}_{\mathcal{P}}| \leqslant (2p!m)^m$.*

Proof. Let $K = ([n], \triangleright) \in \mathcal{K}_{\mathcal{P}}$. By definition, there is a partition of $[n]$ into subsets $U_1, \ldots, U_l$ such that for each i, $(U_i, \triangleright)$ is an irreducible subkei. The block structure ensures that each B_j is B_j-invariant, and so each $(B_j, \triangleright)$ is a subkei; thus each B_j is contained entirely in some U_i. Hence each U_i is a disjoint union of some of the blocks and thus there exists a partition $\{u_1, \ldots, u_l\}$ of $[m]$ such that $U_i = \bigcup_{k \in u_i} B_k$ for each i. As noted at the end of Section 5.2, there at most m^m possibilities for this partition as each element of $[m]$ is contained in one of the $l \leqslant m$ sets $u_1, \ldots, u_l$.

Now we can apply Proposition 6.3.7 (with the number of blocks m replaced by $|u_i|$) to U_i to see that there are at most $(2p!)^{|u_i|}$ irreducible kei on that set. This gives

$$|\mathcal{K}_{\mathcal{P}}| \leqslant m^m \prod_{i=1}^{l} (2p!)^{|u_i|} = (2p!m)^m,$$

proving the result. □

We note that this upper bound is very small compared to the theoretical upper bound of $2^{(c+o(1))n^2}$ for some constant c. We have that $2p!m = 2n(p-1)! \leqslant 2np^p$, so

$$(2p!m)^m = (2p!m)^{n/p} \leqslant (2n)^{n/p} p^n \leqslant (2n^2)^n = 2^{o(n^2)}.$$

6.3.2 Kei such that all orbits have odd prime size

We now consider a natural extension to this question, where each block of the partition $\mathcal{C}([n])$ has an odd prime size but not all of the sizes are equal. The structure of any kei in $\mathcal{K}_\mathcal{P}$ is very similar to that described in the last subsection.

Theorem 6.3.9. *Let $(p_1, \ldots, p_m)$ be a sequence of odd primes and $n = \sum_{i=1}^m p_i$. Let $\mathcal{P} = \{B_1, \ldots, B_m\}$ be a partition of $[n]$ where $|B_i| = p_i$ for each i. Then $|\mathcal{K}_\mathcal{P}| \leqslant (2n^3)^n = 2^{o(n^2)}$.*

Proof. Take the natural partition of $[m]$ formed by including i and j in the same part if and only if $|B_i| = |B_j|$; suppose the resulting sets are $v_1, \ldots, v_r$ and write $V_i = \bigcup_{k \in v_i} B_k$, so each V_i consists of $|v_i|$ blocks, all of the same size. Now take a kei $K \in \mathcal{K}_\mathcal{P}$.

Let $i \in [m]$ and suppose that $i \in v_k$; from the first part of Lemma 6.3.1, $(B_i, \triangleright)$ is either trivial or isomorphic to D_{p_i}, and from the proof of the same result, the components of G_{B_i} can only have size 1, p_i, $2p_i$ or a power of 2. Now take $j \notin v_k$; from Lemma 6.2.4, all the components of G_{B_i} contained within B_j have a common size equal to 1 or p_j. But p_j is distinct from p_i and thus this size must be 1, i.e. B_i acts trivially on any block of a different size.

It follows that each block in V_k acts trivially on V_l for any $l \neq k$, and so the partition $\{V_1, \ldots, V_r\}$ of $[n]$ is a partition into stable subkei and $K = (V_1, \triangleright) \oplus \cdots \oplus (V_r, \triangleright)$. Hence K is determined exactly by the subkei $(V_i, \triangleright)$ for $i = 1, \ldots, r$; as each block contained within each V_i has the same odd prime size, we have from Corollary 6.3.8 and the following discussion that there are at most $(2|V_i|^2)^{|V_i|}$ possibilities for each $(V_i, \triangleright)$. As there are again at most m^m ways of forming the partition $\{v_1, \ldots, v_r\}$, we have that

$$|\mathcal{K}_\mathcal{P}| \leqslant m^m \prod_{i=1}^{r} (2n^2)^{|V_i|} \leqslant n^n (2n^2)^n = (2n^3)^n.$$

□

Chapter 7

Counting racks

So far, we have proved various partial results related to the enumeration of finite racks and kei, depending on the sizes of the orbits involved. In Section 5.2, we established a lower bound of $2^{(1/4+o(1))n^2}$ for the case where every orbit has size two, and in Section 6.2 we proved a much smaller upper bound for the case where every orbit has size equal to an odd prime. We first showed an overall upper bound of $2^{(c+o(1))n^2}$ (with $c \approx 1.1604$) in Section 4.2. In this chapter, we prove an upper bound of $2^{(1/4+o(1)n^2}$ for the number of racks of order n; by comparing it to the lower bound, we establish an accurate asymptotic result on the number of (isomorphism classes of) racks on $[n]$.

The principle behind the argument is as follows: we choose a set $T \subseteq [n]$ and reveal the maps $(f_j)_{j \in T}$. We then consider the components of the graph G_T; an important lemma shows that if V is a set of vertices such that each component contains precisely one element from V then revealing the maps $(f_v)_{v \in V}$ determines the entire rack R. The difficulty is in finding a set T which is not too big, but such that G_T has relatively few components.

We will actually need to consider two different sets T and W. We choose a threshold Δ and first consider the set of vertices $S_{>\Delta}$ that have degree strictly greater than Δ in the underlying graph G_R^0; we will choose probabilistically a relatively small set $W \subseteq [n]$ such that any vertex with high degree in G_R^0 also has high degree in G_W^0. Because the degree of any vertex in $S_{>\Delta}$ is so high, the number of components of G_W^0 contained in $S_{>\Delta}$ is small; this allows us to determine the maps $(f_s)_{s \in S_{>\Delta}}$.

For the vertices in $S_{\leqslant\Delta}$ (the set of vertices with degree at most Δ in G_R^0), we will construct greedily a set T of a given size by adding vertices one at a time and revealing their corresponding maps, each time choosing the vertex whose map joins up the most components. It follows that for every $j \in S_{\leqslant\Delta} \setminus T$, f_j can only join up a limited number of components of G_T; we will reveal the restriction of f_j to these components.

Because of the complex nature of this argument, we will 'store' these revealed maps in a 7-tuple I, and then count the racks consistent with I. The main term in the argument comes from considering maps in $S_{\leqslant\Delta} \setminus T$ acting within components of G_T; we can control the action of these maps by first revealing some extra information corresponding to the neighbours of T in G_R^0, which can themselves be controlled as T consists of low degree vertices.

In Section 7.1 we formally define the information I, which requires some straightforward graph theory; in Section 7.2 we show that the number of possibilities for I is at most $2^{o(n^2)}$. In Section 7.3 we will show the desired upper bound on $|\mathcal{R}_n|$, and in Section 7.4 we show that this bound is only achieved if almost every component of R has size 2.

7.1 Important information in a rack

7.1.1 Degrees in graphical representations of racks

We will first need to consider the degrees of vertices in the graphical representation of a rack. Let $R = ([n], \triangleright)$ be a rack and $T \subseteq [n]$; for each $v \in [n]$, define

$$N_T^+(v) = \{j \in T \mid (v)f_j \neq v\}$$

and

$$\Gamma_T^+(v) = \{(v)f_j \mid j \in N_T^+(v)\},$$

so $N_T^+(v)$ is the set of colours of out-edges of G_T incident to v, while $\Gamma_T^+(v)$ is the set of vertices w such that $\overrightarrow{vw} \in \overrightarrow{E}(G_T)$. If $V \subseteq [n]$, we define $N_T^+(V) := \bigcup_{v \in V} N_T^+(v)$ and $\Gamma_T^+(V) := \bigcup_{v \in V} \Gamma_T^+(v)$.

Definition 7.1.1. With notation as above, define the *out-edge-degree* of v in G_T to be $b_T^+(v) = |N_T^+(v)|$ and the *out-vertex-degree* of v in G_T to be $d_T^+(v) = |\Gamma_T^+(v)|$. Thus $b_T^+(v)$ is the out-degree of v in the multigraph G_T while $d_T^+(v)$ is the out-degree of v in the simple graph G_T^0.

We can of course define the in-degrees $b_T^-(v)$ and $d_T^-(v)$ similarly. Each map f_i consists of disjoint directed cycles, so for any $v, i \in [n]$ there is a directed edge $\overrightarrow{vw} \in \overrightarrow{E}(G_T)$ of colour i for some w if and only if there is a directed edge $\overrightarrow{uv} \in \overrightarrow{E}(G_T)$ of colour i for some u, and thus $b_T^+(v) = b_T^-(v)$. Note also that $d_T^+(v) \leqslant b_T^+(v)$, as every out-neighbour of v is the endpoint of at least one out-edge from v.

We now show that when S is a subrack, all components of G_S are out-regular in both a vertex and an edge sense.

Lemma 7.1.2. *Let $([n], \triangleright)$ be a rack and $(S, \triangleright)$ be a subrack, and let C span a component of G_S, and hence also of G_S^0. Then for any $u, v \in C$, $b_S^+(u) = b_S^+(v)$ and $d_S^+(u) = d_S^+(v)$, i.e. both $G_S[C]$ and $G_S^0[C]$ are out-regular.*

Proof. If $|C| = 1$ then the result is trivial, so we may take $|C| \geqslant 2$. Let $u, v \in C$ be distinct.

First suppose that v is an out-neighbour of u, so that there is a directed edge $\overrightarrow{uv} \in \overrightarrow{E}(G_S)$ of some colour $i \in S$, i.e. $(u)f_i = v$. Note that $i \in N_S^+(u)$ and thus $N_S^+(u) \neq \emptyset$; take an arbitrary $j \in N_S^+(u)$ and note that as S is a subrack, $k := (j)f_i \in S$. Suppose for a contradiction that $(v)f_k = v$; then

$$(u)f_i = v = (v)f_k = (v)f_i^{-1} f_j f_i = (u)f_j f_i \tag{7.1.1}$$

and thus (as f_i is a bijection) $u = (u)f_j$, contradicting the fact that $j \in N_S^+(u)$. Hence $k \in N_S^+(v)$ and $(N_S^+(u))f_i \subseteq N_S^+(v)$; as f_i is a bijection we conclude that $b_S^+(u) = |N_S^+(u)| \leqslant |N_S^+(v)| = b_S^+(v)$.

Now let $w \in \Gamma_S^+(u)$, so there exists a $j \in N_S^+(u)$ such that $w = (u)f_j$. As before, let $k = (j)f_i \in N_S^+(v)$, so $(v)f_k \in \Gamma_S^+(v)$; from the last two equalities in (7.1.1), $(w)f_i = (v)f_k \in \Gamma_S^+(v)$ and thus $(\Gamma_S^+(u))f_i \subseteq \Gamma_S^+(v)$. As f_i is a bijection, $d_S^+(u) = |\Gamma_S^+(u)| \leqslant |\Gamma_S^+(v)| = d_S^+(v)$.

Now let $u, v \in C$ be arbitrary; from Lemma 4.3.7, there is a directed path $u = u_0, \ldots, u_r = v$ in $G_S[C]$. From above, we have that $b_S^+(u) \leqslant b_S^+(u_1) \leqslant \cdots \leqslant b_s^+(v)$ and similarly $d_S^+(u) \leqslant d_S^+(v)$. By instead considering a directed vu-path we have that $b_S^+(v) \leqslant b_S^+(u)$ and $d_S^+(v) \leqslant d_S^+(u)$, and thus the result follows. □

7.1.2 Some multigraph theory

The construction of the information I requires some straightforward graph theory. We will write $\mathrm{cp}(G)$ for the number of components of a multigraph G.

In this subsection we will consider only undirected multigraphs for clarity. As stated in the Introduction, 'paths' and 'components' of a directed multigraph are defined by the underlying undirected multigraph, so all of these results remain true for directed multigraphs.

Let G be a multigraph. Then for any distinct $u, v \in V(G)$ the following is true: if there exists a uv-path in G then there is a (possibly degenerate) cycle in $G + uv$ containing uv and $\mathrm{cp}(G + uv) = \mathrm{cp}(G)$, while if there is no uv-path in G then $\mathrm{cp}(G + uv) = \mathrm{cp}(G) - 1$ and there is no cycle in $G + uv$ containing uv. In the latter case, adding uv merges the components of G containing u and v, and thus coarsens the partition of $V(G)$ into vertex sets of components; it follows easily that we always have $\text{cp}(G+E) \leqslant \text{cp}(G)$. These facts allow us to prove the following standard lemma.

Lemma 7.1.3. *Let G be a multigraph and E_1 and E_2 be multisets of unordered pairs of elements of $V(G)$. Then* $\text{cp}(G) - \text{cp}(G+E_2) \geqslant \text{cp}(G+E_1) - \text{cp}(G+E_1+E_2)$.

Proof. Write $E_2 = \{e_1, \ldots, e_l\}$, and let $a = \text{cp}(G+E_1) - \text{cp}(G+E_1+E_2)$. Consider adding the edges of E_2 in the given order to $G+E_1$; as $\text{cp}(G+E_1+E_2) = \text{cp}(G+E_1)-a$ we have from above that there are exactly a edges $e_{i_1}, \ldots, e_{i_a}$ such that $\text{cp}(G+E_1+\{e_1, \ldots, e_{i_j}\}) = \text{cp}(G+E_1+\{e_1, \ldots, e_{i_j-1}\}) - 1$ for each j. Now write $e_{i_j} = u_j v_j$ for each j; we have shown that there is no $u_j v_j$ path in $G+E_1+\{e_1, \ldots, e_{i_j-1}\}$ for any j.

Now add the edges $(e_{i_1}, \ldots, e_{i_a})$ in the same order to G; if for any j we have that $\text{cp}(G+\{e_{i_1}, \ldots, e_{i_j}\}) = \text{cp}(G+\{e_{i_1}, \ldots, e_{i_{j-1}}\})$ then there must be a $u_j v_j$-path in $G+\{e_{i_1}, \ldots, e_{i_{j-1}}\} \subseteq G+E_1+\{e_1, \ldots, e_{i_j-1}\}$, a contradiction. Hence $\text{cp}(G+\{e_{i_1}, \ldots, e_{i_a}\}) = \text{cp}(G) - a$ and

$$\text{cp}(G)-\text{cp}(G+E_2) \geqslant \text{cp}(G)-\text{cp}(G+\{e_{i_1}, \ldots, e_{i_a}\}) = a = \text{cp}(G+E_1)-\text{cp}(G+E_1+E_2).$$

□

Definition 7.1.4. Let G be a multigraph and E be a multiset of unordered pairs of elements of $V(G)$. Let $C \subseteq V(G)$ span a component of G; we say that C is *merged by E* if there exist $u \in C$, $v \in V(G) \setminus C$ such that uv is an edge from E. We denote by $M(G, E)$ the set of (vertex sets of) components of G merged by E.

Note that by considering the components merged by multisets E and F separately, we have that $M(G, E \uplus F) = M(G, E) \cup M(G, F)$. As a single edge can merge at most two components, $|M(G, \{e\})| \leqslant 2$ for any unordered pair e.

Lemma 7.1.5. *Let G be a multigraph and E and $\{e\}$ be multisets of unordered pairs of elements of $V(G)$. Then if* $\text{cp}(G+E+\{e\}) = \text{cp}(G+E)$, $M(G, E \uplus \{e\}) = M(G, E)$.

Proof. Write $e = uv$; if u and v are in the same component of G then e is not a merging edge and $M(G, \{e\}) = \emptyset$, so suppose that u and v are in different components of G. Write C for the vertex set of the component containing u and D for the vertex set of the component containing v, so that $M(G, \{e\}) = \{C, D\}$.

Now as $\text{cp}(G+E+e) = \text{cp}(G+E)$ there is a uv-path in $G+E$; let $u = u_0, \ldots, u_r = v$ be such a path and suppose that u_t is the last vertex on the path belonging to C. Then $t \leqslant r-1$ and $u_{t+1} \notin C$, so by considering the edge $u_t u_{t+1}$ we see that $C \in M(G, E)$. By the symmetry in u and v, $D \in M(G, E)$ as well. We have shown in all cases that $M(G, \{e\}) \subseteq M(G, E)$ and thus that $M(G, E \uplus \{e\}) = M(G, E)$. □

Corollary 7.1.6. *Let G be a multigraph and E be a multiset of unordered pairs of elements of $V(G)$. Then $|M(G,E)| \leqslant 2(\mathrm{cp}(G) - \mathrm{cp}(G+E))$.*

Proof. As in the proof of Lemma 7.1.3, order E as $\{e_1, \ldots, e_l\}$ and write $a = \mathrm{cp}(G) - \mathrm{cp}(G+E)$; then there are precisely a edges $e_{i_1}, \ldots, e_{i_a}$ such that $\mathrm{cp}(G+\{e_1, \ldots, e_{i_j}\}) = \mathrm{cp}(G+\{e_1, \ldots, e_{i_j - 1}\}) - 1$ for each j. Write $E_k = \{e_1, \ldots, e_k\}$ for each k, so we always have $M(G, E_k) = M(G, E_{k-1} \uplus \{e_k\}) = M(G, E_{k-1}) \cup M(G, \{e_k\})$. Now consider adding the edges of E in the order given to G; if $k \neq i_j$ for any j then $\mathrm{cp}(G, E_k) = \mathrm{cp}(G, E_{k-1})$ and so from Lemma 7.1.5 we have that $M(G, E_k) = M(G, E_{k-1})$, while if $k = i_j$ for some j we have that $|M(G, E_k)| \leqslant |M(G, E_{k-1})| + 2$. As there are only a such edges it follows that $|M(G,E)| \leqslant 2a = 2(\mathrm{cp}(G) - \mathrm{cp}(G+E))$. □

7.1.3 The information $\mathcal{I}(R)$

We will introduce the following terminology. For any rack $R = ([n], \triangleright)$ and any Δ with $1 \leqslant \Delta \leqslant n-1$, let

$$S_{\leqslant \Delta}(R) := \{v \in [n] \mid d_R^+(v) \leqslant \Delta\}$$

denote the set of all elements of $[n]$ with out-vertex-degree at most Δ. Write $S_{>\Delta}(R) = [n] \setminus S_{\leqslant \Delta}(R)$ for the set of all elements with out-vertex-degree strictly greater than Δ. We now show that this partition is actually a partition into subracks.

Lemma 7.1.7. *Let $R = ([n], \triangleright)$ be a rack and $1 \leqslant \Delta \leqslant n-1$. Then $(S_{>\Delta}(R), \triangleright)$ and $(S_{\leqslant \Delta}(R), \triangleright)$ are both subracks of R.*

Proof. Suppose for a contradiction that there exist $u \in S_{>\Delta}(R)$ and $v \in S_{\leqslant \Delta}(R)$ such that u and v are adjacent in G_R. Then u and v are in the same component of G_R and thus by Lemma 7.1.2 (with $S = R$) they have the same out-vertex-degree, a clear contradiction to the definitions of $S_{>\Delta}(R)$ and $S_{\leqslant \Delta}(R)$. Hence $S_{>\Delta}(R)$ and $S_{\leqslant \Delta (R)}$ are separated in G_R and thus both $(S_{>\Delta}(R), \triangleright)$ and $(S_{\leqslant \Delta}(R), \triangleright)$ are subracks. □

Now fix

$$\Delta := (\log_2 n)^3,$$

so $\Delta \leqslant n-1$ for sufficiently large n. For a given rack R, we will construct a specific set $T(R) \subseteq S_{\leqslant \Delta}(R)$, with $|T(R)| \leqslant (\log_2 n)^2$, by the following procedure.

Construction 7.1.8. If $S_{\leqslant \Delta}(R) = \emptyset$ then $T(R) = \emptyset$. Otherwise, order the vertices of $S_{\leqslant \Delta}(R)$ as follows: first choose u_1 so that $\mathrm{cp}(G_{u_1}) \leqslant \mathrm{cp}(G_v)$ for any $v \in$

$S_{\leqslant\Delta}(R)$, then given a partial ordering $u_1, \ldots, u_k$, choose the next vertex u_{k+1} such that $\operatorname{cp}(G_{\{u_1,\ldots,u_k,u_{k+1}\}}) \leqslant \operatorname{cp}(G_{\{u_1\ldots,u_k,v\}})$ for any $v \in S_{\leqslant\Delta}(R) \setminus \{u_1, \ldots, u_k\}$. Now take

$$L := \lfloor (\log_2 n)^2 \rfloor$$

and define

$$T(R) := \begin{cases} \{u_1, \ldots, u_L\} & \text{if } |S_{\leqslant\Delta}(R)| > (\log_2 n)^2 \\ S_{\leqslant\Delta}(R) & \text{otherwise.} \end{cases}$$

We now introduce some more notation. For any $j \in [n]$, write $\overrightarrow{E}_j = \overrightarrow{E}(G_j)$ for the set of edges of G_R of colour j and $M_j = M(G_{T(R)}, \overrightarrow{E}_j)$ for the set of (vertex sets of) components of $G_{T(R)}$ merged by the edges of colour j. Note that if $j \in T(R)$ then $\overrightarrow{E}_j \subseteq \overrightarrow{E}(G_{T(R)})$ and so $M_j = \emptyset$.

The key property of the set $T(R)$, in terms of the above concepts, is given in this next lemma.

Lemma 7.1.9. *Let $R = ([n], \triangleright)$ be a rack. Then for any $j \in S_{\leqslant\Delta}(R) \setminus T(R)$,*

$$|M_j| \leqslant \frac{2n}{(\log_2 n)^2}.$$

Proof. Note that if $|S_{\leqslant\Delta}(R)| \leqslant L = \lfloor (\log_2 n)^2 \rfloor$ then $T(R) = S_{\leqslant\Delta}(R)$ and the statement is trivial. We will thus assume that $s = |S_{\leqslant\Delta}(R)| \geqslant L + 1$ and so $|T(R)| = L$.

For $1 \leqslant i \leqslant s$, write $H_i = G_{\{u_1,\ldots,u_i\}}$ and $x_i = \operatorname{cp}(H_{i-1}) - \operatorname{cp}(H_i)$, where $H_0 = G_\emptyset$ is the empty graph on $[n]$. Note that $\operatorname{cp}(H_i) = \operatorname{cp}(H_{i-1} + \overrightarrow{E}_{u_i}) \leqslant \operatorname{cp}(H_{i-1})$, and so $x_i \geqslant 0$ for each i; we also have that

$$\sum_{i=1}^{s} x_i = \sum_{i=i}^{s} \big(\operatorname{cp}(H_{i-1}) - \operatorname{cp}(H_i)\big) = \operatorname{cp}(H_0) - \operatorname{cp}(H_s) \leqslant \operatorname{cp}(H_0) = n.$$

Now fix an i and suppose that $x_i < x_{i+1}$; then

$$\begin{aligned} \operatorname{cp}(H_{i-1}) - \operatorname{cp}(H_i) &< \operatorname{cp}(H_i) - \operatorname{cp}(H_{i+1}) \\ &= \operatorname{cp}(H_{i-1} + \overrightarrow{E}_{u_i}) - \operatorname{cp}(H_{i-1} + \overrightarrow{E}_{u_i} + \overrightarrow{E}_{u_{i+1}}) \\ &\leqslant \operatorname{cp}(H_{i-1}) - \operatorname{cp}(H_{i-1} + \overrightarrow{E}_{u_{i+1}}), \end{aligned}$$

from Lemma 7.1.3. But then $\operatorname{cp}(G_{\{u_1,\ldots u_{i-1},u_i\}}) > \operatorname{cp}(G_{\{u_1,\ldots u_{i-1},u_{i+1}\}})$, contradicting our ordering of the vertices. Hence $x_i \geqslant x_{i+1}$ and as i was arbitrary we conclude that $(x_i)_{i=1}^s$ is a decreasing sequence.

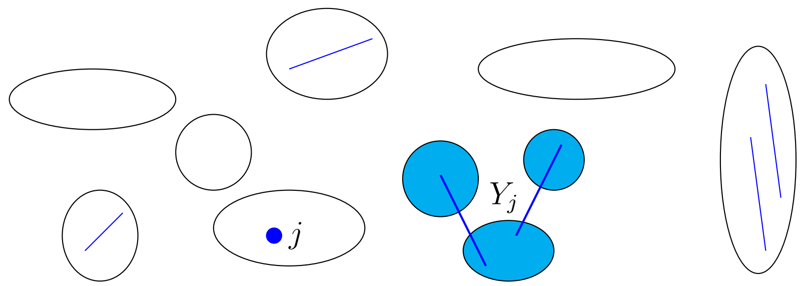

Figure 7.1: A representation of the components of $G_{T(R)}$, with the edges of colour j in blue. Here, precisely three components (shaded light blue) are merged by the edges of colour j, so $|M_j| = 3$; Y_j is the set of vertices in these shaded components. If $j \in S_{\leqslant\Delta}(R) \setminus T(R)$ then only the restriction $f_j|_{Y_j}$ is included in the information $\mathcal{I}(R)$.

Now suppose $x_{L+1} > n/(L+1)$. Then for $1 \leqslant i \leqslant L+1$, $x_i \geqslant x_{L+1} > n/(L+1)$, and as each x_i is non-negative we have

$$n \geqslant \sum_{i=1}^{L+1} x_i > \frac{(L+1)n}{L+1} = n,$$

a clear contradiction. Hence $x_{L+1} \leqslant n/(L+1) \leqslant n(\log_2 n)^{-2}$.

Now take $j \in S_{\leqslant\Delta}(R) \setminus T(R)$, so in our ordering of $S_{\leqslant\Delta}(R)$ we have $j = u_k$ for some $k > L$. By our construction, $\mathrm{cp}(G_{\{u_1,\ldots,u_L,j\}}) \geqslant \mathrm{cp}(G_{\{u_1,\ldots,u_L,u_{L+1}\}})$; noting that $G_{T(R)} = G_{\{u_1,\ldots,u_L\}} = H_L$, we may rewrite this as $\mathrm{cp}(G_{T(R)} + \overrightarrow{E}_j) \geqslant \mathrm{cp}(H_{L+1})$, and thus

$$\mathrm{cp}(G_{T(R)}) - \mathrm{cp}(G_{T(R)} + \overrightarrow{E}_j) \leqslant \mathrm{cp}(H_L) - \mathrm{cp}(H_{L+1}) = x_{L+1} \leqslant \frac{n}{(\log_2 n)^2}.$$

We can combine this with Corollary 7.1.6 to see that

$$|M_j| = |M(G_{T(R)}, \overrightarrow{E}_j)| \leqslant 2(\mathrm{cp}(G_{T(R)}) - \mathrm{cp}(G_{T(R)} + \overrightarrow{E}_j)) \leqslant \frac{2n}{(\log_2 n)^2},$$

showing the result. □

Before formally defining the information $\mathcal{I}(R)$ associated with a rack R, we will need some more notation. Firstly, write $T^+(R) = T(R) \cup \Gamma_R^+(T(R))$. For any $j \in S_{\leqslant\Delta}(R) \setminus T(R)$, define $Y_j = \bigcup_{C \in M_j} C$ to be the set of *vertices* in components of $G_{T(R)}$ merged by $\overrightarrow{E}_j$, and write $Z_j = [n] \setminus Y_j$ (see Figure 7.1). Now let $\mathbf{M}$ be the $(|S_{\leqslant\Delta}(R)| - |T(R)|)$-tuple $(M_j)_{j \in S_{\leqslant\Delta}(R) \setminus T(R)}$, where we order the vertices in some arbitrary way, and let $Y = \bigcup_{j \in S_{\leqslant\Delta}(R) \setminus T(R)} (Y_j \times \{j\}) \subseteq [n]^2$. The following lemma gives the key property of the set Z_j, in a slightly more general setting.

Lemma 7.1.10. *Let $R = ([n], \triangleright)$ be a rack, $W \subseteq [n]$ and $j \in [n]$. Let $C \subseteq [n]$ be a W-orbit (so C spans a component of G_W), with $C \notin M(G_W, \overrightarrow{E}_j)$. Then C is j-invariant.*

Proof. Take a set C as described and let $x \in C$ be arbitrary. If $(x)f_j = y \neq x$ then $\overrightarrow{xy}$ is an edge of colour j; as C is *not* merged by $\overrightarrow{E}_j$, we must have $(x)f_j = y \in C$ and (as x was arbitrary and f_j is a bijection) it follows that $(C)f_j = C$. □

Now apply this lemma with $W = T(R)$; for any $C \subseteq [n]$ spanning a component of $G_{T(R)}$ with $C \notin M_j$, $(C)f_j = C$. As Z_j consists of all vertices in components of $G_{T(R)}$ not merged by $\overrightarrow{E}_j$, we have that $(Z_j)f_j = Z_j$, and thus that $(Y_j)f_j = Y_j$.

We will now formally define the information associated with a rack R.

Definition 7.1.11. Let $R = ([n], \triangleright)$ be a rack and let $\Delta = (\log_2 n)^3$. Then with notation as above define

$$\mathcal{I}(R) := \left(S_{\leqslant\Delta}(R),\ \triangleright|_{[n]\times S_{>\Delta}},\ T(R),\ \triangleright|_{T(R)\times[n]},\ \triangleright|_{[n]\times T^{+}(R)},\ \mathbf{M},\ \triangleright|_{Y}\right).$$

As $i \triangleright j = (i)f_j$, the second entry of this 7-tuple is equivalent to the set of maps $(f_j)_{j\in S_{>\Delta}(R)}$ or alternatively the graph $G_{S_{>\Delta}(R)}$. The fourth entry is equivalent to the set of maps $(f_j|_{T(R)})_{j\in[n]}$; knowing this set determines $N^{+}_{R}(T(R))$ and thus $\Gamma^{+}_{R}(T(R))$, which is necessary for the fifth entry. Note also that the fifth entry is equivalent to the set of maps $(f_i)|_{i\in T^{+}(R)}$ and thus the graph $G_{T^{+}(R)}$, while the seventh entry is equivalent to $(f_j|_{Y_j})_{j\in S_{\leqslant\Delta}(R)\setminus T(R)}$.

We will think of $\mathcal{I}$ as a map from $\mathcal{R}_n$ to a set $\mathfrak{I}_n$ of 7-tuples; the form of this set $\mathfrak{I}_n$ will be considered in more detail in Section 7.3. In the next section, we show that the image $\mathcal{I}(\mathcal{R}_n) \subseteq \mathfrak{I}_n$ is small. We will do this by first considering the map $\mathcal{I}'$, where $\mathcal{I}'(R) = (\mathcal{I}_1(R), \mathcal{I}_2(R))$ and then $\mathcal{I}''$, where $\mathcal{I}''(R) = (\mathcal{I}_3(R), \ldots, \mathcal{I}_7(R))$.

7.2 Determining the information $\mathcal{I}(R)$

7.2.1 Random subsets

The part of the argument relating to the vertices of high degree requires some probabilistic tools. In particular, we will require the following result of Hoeffding [15]; these results are known generally as the *Chernoff bounds*.

Theorem 7.2.1. *Let $X_1, \dots, X_n$ be independent random variables, each taking values in the range $[0,1]$. Let $X = \sum_{i=1}^n X_i$ and let $p = \mathbb{E}[X]/n$. Then for $x \geqslant p$,*

$$\mathbb{P}(X \geqslant nx) \leqslant \left(\left(\frac{p}{x}\right)^x \left(\frac{1-p}{1-x}\right)^{1-x} \right)^n,$$

and for $x \leqslant p$,

$$\mathbb{P}(X \leqslant nx) \leqslant \left(\left(\frac{p}{x}\right)^x \left(\frac{1-p}{1-x}\right)^{1-x} \right)^n.$$

□

We will use the following, more workable corollary (see for example Theorems 2.1 and 2.8 and Corollary 2.3 of [20]).

Corollary 7.2.2. *Let $X_1, \dots, X_n$ be independent random variables, each taking values in the range $[0,1]$. Let $X = \sum_{i=1}^n X_i$. Then for any $\epsilon \in [0,1]$,*

$$\mathbb{P}(X \geqslant (1+\epsilon)\mathbb{E}[X]) \leqslant e^{-\epsilon^2 \mathbb{E}[X]/3}$$

and

$$\mathbb{P}(X \leqslant (1-\epsilon)\mathbb{E}[X]) \leqslant e^{-\epsilon^2 \mathbb{E}[X]/2}.$$

□

Let $R = ([n], \triangleright)$ be a rack. For the rest of this chapter, we will consider only the out-*vertex*-degree $d_T^+(v)$ of any given vertex; if $T = [n]$, we will write $d_v^+ = d_R^+(v)$.

Lemma 7.2.3. *Let $R = ([n], \triangleright)$ be a rack and let $p, \epsilon \in (0,1)$. Construct a random subset X of $[n]$ by retaining each element with probability p, independently for all elements. Then*

1. *$\mathbb{P}(|X| \geqslant (1+\epsilon)np) \leqslant e^{-\epsilon^2 np/3}$;*

2. *For any $v \in [n]$ and $0 < \delta \leqslant d_v^+ p$, $\mathbb{P}\big(d_X^+(v) \leqslant (1-\epsilon)\delta\big) \leqslant e^{-\epsilon^2 \delta/2}$.*

Proof. For each j, let $X_j = \mathbf{1}_{\{j \in X\}}$, so that each $X_j \sim \mathrm{Ber}(p)$ and the variables $(X_j)_{j=1}^n$ are independent. Then $|X| = \sum_{j=1}^n X_j \sim \mathrm{Bi}(n,p)$ and $\mathbb{E}[X] = np$, so we can apply Theorem 7.2.2 to $|X|$, showing the first statement.

For the second statement, take a vertex $v \in [n]$ and let $v_1, \dots, v_{d_v^+}$ be the out-neighbours of v in G_R^0; for each $1 \leqslant i \leqslant d_v^+$, choose an element $j_{vi} \in [n]$ such that $(v)f_{j_{vi}} = v_i$ and put

$$J_v := \{j_{vi} \mid i = 1, \dots, d_v^+\}.$$

The elements $j_{v1}, \ldots, j_{vd_v^+}$ are clearly distinct and so $|J_v| = d_v^+$. Now

$$|J_v \cap X| = \sum_{i=1}^{d_v^+} X_{j_{v_i}} \sim \mathrm{Bi}(d_v^+, p),$$

so $\mathbb{E}[|J_v \cap X|] = d_v^+ p$ and we can apply Theorem 7.2.2 to see that for any $0 < \delta \leqslant d_v^+ p$,

$$\mathbb{P}\big(|J_v \cap X| \leqslant (1-\epsilon)\delta\big) \leqslant \mathbb{P}\big(|J_v \cap X| \leqslant (1-\epsilon)d_v^+ p\big) \leqslant e^{-\epsilon^2 d_v^+ p/2} \leqslant e^{-\epsilon^2 \delta/2}.$$

Since $d_X^+(v) \geqslant |J_v \cap X|$, the second result follows. □

7.2.2 The high degree part

We will need the following crucial lemma, which will also be used later.

Lemma 7.2.4. *Let $R = ([n], \triangleright)$ be a rack and $T \subseteq [n]$. Let C be a T-orbit, so C spans a component of G_T, and let $v \in C$. Let $A \subseteq [n]$ be $[n]$-invariant. Then knowledge of the maps $(f_i|_A)_{i \in T}$ and $f_v|_A$ is sufficient to determine the maps $(f_u|_A)_{u \in C}$, and further, the maps $(f_u|_A)_{u \in C}$ are conjugate in* $\mathrm{Sym}(A)$.

Proof. Let $u \in C$, so from Lemma 4.3.7 there is a directed vu-path in G_T. Let the colours of the edges of this path be $i_1, \ldots, i_l$, so $(v)f_{i_1} \cdots f_{i_l} = u$ and thus $f_u = f_{i_l}^{-1} \cdots f_{i_1}^{-1} f_v f_{i_1} \cdots f_{i_l}$. As A is $[n]$-invariant, $(A)f_{i_j} = A$ for any j, and thus

$$f_u|_A = f_{i_l}^{-1}|_A \cdots f_{i_1}^{-1}|_A f_v|_A f_{i_1}|_A \cdots f_{i_l}|_A.$$

But as each $i_j \in T$, all of these maps are determined and thus $f_u|_A$ is determined. Also note that $f_u|_A$ is conjugate to $f_v|_A$ by the map $(f_{i_1} \cdots f_{i_r})|_A$, proving the result. □

We will now show the first main result of this section.

Proposition 7.2.5. *Let $R = ([n], \triangleright)$ be a rack. Then for n sufficiently large there exists a set $W \subseteq [n]$, with*

$$|W| \leqslant w_0(n) := \frac{7}{2} \frac{n}{(\log_2 n)^{3/2}},$$

such that $\mathcal{I}'(R)$ is determined by the sets $S_{\leqslant \Delta}(R)$ and W and the maps $(f_i)_{i \in W}$.

Proof. We will construct the set W by a mixture of probabilistic and deterministic arguments. Let $p = (\log_2 n)^{-3/2}$ and consider a random subset X of $[n]$ as described in Lemma 7.2.3. For each n, let E_n be the event that $|X| \leqslant 3np/2$; as $np \to \infty$ as $n \to \infty$, we have from item 1 of the lemma (with $\epsilon = 1/2$) that

$$\mathbb{P}(E_n^{\mathrm{c}}) \leqslant \mathbb{P}(|X| \geqslant 3np/2) \leqslant e^{-np/12} \to 0$$

as $n \to \infty$.

Now for each $v \in [n]$ and $U \subseteq [n]$, call v *U-bad* if $d_U^+(v) \leqslant (\log_2 n)^{3/2}/2$, and let B_v be the event that v is X-bad, so that $\mathbb{P}(B_v) = \mathbb{P}\big(d_X^+(v) \leqslant (\log_2 n)^{3/2}/2\big)$. Let $N = N(X)$ denote the number of X-bad vertices in $S_{>\Delta} := S_{>\Delta}(R)$, so that $N = \sum_{v \in S_{>\Delta}} \mathbf{1}_{B_v}$. Now $d_v^+ p > \Delta p = (\log_2 n)^{3/2}$ for each $v \in S_{>\Delta}$; hence from item 2 of Lemma 7.2.3 (with $\epsilon = 1/2$ and $d = \Delta p$),

$$\mathbb{E}[N] = \sum_{v \in S_{>\Delta}} \mathbb{P}(B_v) = \sum_{v \in S_{>\Delta}} \mathbb{P}\big(d_X^+(v) \leqslant (\log_2 n)^{3/2}/2\big) \leqslant e^{-(\log_2 n)^{3/2}/8} |S_{>\Delta}|.$$

Now for each n let F_n be the event that $N = 0$, i.e. that every vertex in $S_{>\Delta}$ is X-good. Then from Markov's Inequality

$$\mathbb{P}(F_n^{\mathrm{c}}) = \mathbb{P}(N \geqslant 1) \leqslant \mathbb{E}[N] \leqslant |S_{>\Delta}| e^{-(\log_2 n)^{3/2}/8} \leqslant n e^{-(\log_2 n)^{3/2}/8} \to 0$$

as $n \to \infty$, noting that $(\log_2 n)^{3/2} = \omega(\log n)$; hence $\mathbb{P}(F_n) \to 1$ as $n \to \infty$.

We have shown that $\mathbb{P}(E_n \cap F_n) \to 1$ as $n \to \infty$, and thus in particular $\mathbb{P}(E_n \cap F_n) > 0$ for sufficiently large n. Hence we can find a set $U \subseteq [n]$ such that $|U| \leqslant 3np/2$ and each vertex in $S_{>\Delta}$ is U-good; this means that $d_U^+(v) > (\log_2 n)^{3/2}/2$ whenever $v \in S_{>\Delta}$, or in graphical terms, that each vertex $v \in S_{>\Delta}$ is adjacent to at least $(\log_2 n)^{3/2}/2$ vertices in G_U^0.

Now from Lemma 7.1.7 there are no edges from $S_{>\Delta}$ to $S_{\leqslant\Delta} = [n] \setminus S_{>\Delta}$, so $S_{>\Delta}$ is a disjoint union of vertex sets of components of G_U^0. We have shown that each component of G_U^0 contained within $S_{>\Delta}$ has size at least $(\log_2 n)^{3/2}/2 + 1$ and hence there are at most $2|S_{>\Delta}|(\log_2 n)^{-3/2}$ such components. Write the vertex sets of these components as $\{C_1, \ldots, C_r\}$ and take a set of vertices $V = \{v_1, \ldots, v_r\}$ such that $v_k \in C_k$ for each k; we have shown that

$$|V| = r \leqslant \frac{2n}{(\log_2 n)^{3/2}}.$$

Now from Lemma 7.2.4 (with $T = U$ and $A = [n]$), knowledge of the maps $(f_i)|_{i \in U}$ and f_{v_k} is sufficient to determine the maps $(f_u)_{u \in C_k}$; applying this to each component in turn shows that knowledge of the maps $(f_i)_{i \in U}$ and $(f_v)_{v \in V}$ are sufficient to determine $(f_u)_{u \in S_{>\Delta}}$, i.e. the second entry of $\mathcal{I}'(R)$. So put $W = U \cup V$; as $|U| \leqslant 3n(\log_2 n)^{-3/2}/2$ we have the result. $\square$

Corollary 7.2.6. *Let $\epsilon > 0$. There exists some positive integer n_1 such that for all $n \geqslant n_1$, $|\mathcal{I}'(\mathcal{R}_n)| \leqslant 2^{\epsilon n^2}$.*

Proof. Take n sufficiently large for the previous result to hold; then $\mathcal{I}'(R)$ is defined by $S_{\leqslant\Delta}(R)$ (or equivalently $S_{>\Delta}(R)$), W and the maps $(f_i)_{i\in W}$ for a suitable set W depending on R, with $|W| \leqslant w_0(n)$. Now fix such an n; it follows that $|\mathcal{I}'(\mathcal{R}_n)|$ is at most the number of distinct triples $(S_{\leqslant\Delta}(R), W, (f_i)_{i\in W})$ arising from all racks in $\mathcal{R}_n$. There are clearly at most 2^n possibilities for each of $S_{\leqslant\Delta}(R)$ and W; as there are $n! \leqslant n^n$ choices for any map f_i, and $|W| \leqslant w_0(n)$, there are at most $n^{nw_0(n)}$ possibilities for the maps $(f_i)_{i\in W}$. Hence $|\mathcal{I}'(\mathcal{R}_n)| \leqslant 2^{2n} n^{nw_0(n)}$, and from the previous result we have that

$$\begin{aligned}
\log_2\left(2^{2n} n^{nw_0(n)}\right) &= 2n + n(\log_2 n)w_0(n)\\
&= 2n + n(\log_2 n)\frac{7}{2}\frac{n}{(\log_2 n)^{3/2}}\\
&= 2n + \frac{7}{2}\frac{n^2}{(\log_2 n)^{1/2}}\\
&= o(n^2).
\end{aligned}$$

Hence for any $\epsilon > 0$, there exists a positive integer n_1 such that for $n \geqslant n_1$, $|\mathcal{I}'(\mathcal{R}_n)| \leqslant 2^{2n} n^{nw_0} \leqslant 2^{\epsilon n^2}$, proving the result. $\square$

7.2.3 Components of the graph $G_{T(R)}$

In order to prove that there are few choices for $\mathcal{I}''(R)$, we will need the following lemma; the ideas used in this argument are similar to those in Lemmas 5.1.1 and Lemma 7.2.4.

Lemma 7.2.7. *Let $R = ([n], \triangleright)$ be a rack. Let C be a $T(R)$-orbit, so C spans a component of $G_{T(R)}$, and let $v \in C$. Then for any $j \in [n]$, knowledge of the maps $(f_l|_{T(R)})_{l\in[n]}$ and $(f_i)|_{i\in T^+(R)}$ and the vertex $(v)f_j$ is sufficient to determine the map $f_j|_C$.*

As in Definition 7.1.11, the maps $(f_j|_{T(R)})_{j\in[n]}$ determine the set $\Gamma_R^+(T(R))$ and thus the set $T^+(R)$.

Proof. Let $u \in C$; from Lemma 4.3.7 there is a directed vu-path in $G_{T(R)}$ (note that this graph is determined from the maps $(f_i)_{i\in T(R)}$, knowledge of which is assumed). Let $\overrightarrow{d}(v,u)$ denote the length of the shortest directed vu-path; we will show that $(u)f_j$ is determined by induction on the graph distance $d = \overrightarrow{d}(v,u)$. The base case $d = 0$ is true by assumption, so take $d > 0$ and suppose the result is true for smaller d.

Take a shortest directed path (of length d) from v to u and let w be the penultimate vertex on the path; then $\overrightarrow{wu}$ is an edge in $G_{T(R)}$, so there exists some $i \in T(R)$ such that $(w)f_i = u$. As we have knowledge of $f_j|_{T(R)}$, $k := (i)f_j$ is determined; as $k \in T(R) \cup \Gamma_R^+(T(R)) = T^+(R)$, the map f_k is also determined. Further, $\overrightarrow{d}(v, w) = d-1$, so by the inductive hypothesis the vertex $(w)f_j$ is determined. Now $f_k = f_{(i)f_j} = f_j^{-1} f_i f_j$, and so $f_j = f_i^{-1} f_j f_k$; hence

$$(u)f_j = (u)f_i^{-1} f_j f_k = (w)f_j f_k,$$

which is determined as we know the vertex $(w)f_j$ and the map f_k. The result follows by induction. □

We can now show the second main result of this section.

Proposition 7.2.8. *Let $\epsilon > 0$. Then there exists a positive integer n_2 such that for any $n \geqslant n_2$, $|\mathcal{I}''(\mathcal{R}_n)| \leqslant 2^{\epsilon n^2}$.*

Proof. As in Corollary 7.2.6, $|\mathcal{I}''(\mathcal{R}_n)|$ is equal to the number of distinct values of the 5-tuple $\mathcal{I}''(R)$ as R ranges over all racks of order n; we will produce bounds for each of these entries. There are clearly at most 2^n choices for the set $T(R)$; recall that by construction $|T(R)| \leqslant L = \lfloor (\log_2 n)^2 \rfloor$. There are at most n possibilities for the vertex $(u)f_i$ for any $i, u \in [n]$, so there are at most $(n^L)^n$ possibilities for the maps $(f_j|_{T(R)})_{j \in [n]}$ and at most $(n^n)^L$ possibilities for the maps $(f_i)_{i \in T(R)}$. Now recall that by construction $T(R) \subseteq S_{\leqslant \Delta}(R)$[1], so by definition $d_v^+ = |\Gamma_R^+(v)| \leqslant \Delta$ for any $v \in S_{\leqslant \Delta}(R)$; it follows that $|\Gamma_R^+(T(R))| \leqslant \sum_{v \in T(R)} d_v^+ \leqslant L\Delta$, and so there are at most $(n^n)^{L\Delta}$ possibilities for the maps $(f_i)_{i \in \Gamma_R^+(T(R))}$. Hence there are at most $2^n n^{(2+\Delta)Ln}$ possible values for the first three entries of $\mathcal{I}''(R)$ (as R ranges over all racks of order n).

Now consider a rack $R \in \mathcal{R}_n$ and suppose that the first three entries of $\mathcal{I}''(R)$ have been determined. Fix a $j \in S_{\leqslant \Delta}(R) \setminus T(R)$; we must consider the possibilities for the set M_j of components of $G_{T(R)}$ merged by $\overrightarrow{E}_j$, and the restricted map $f_j|_C$ for each $C \in M_j$. If $|M_j| = a_j$ then there are (crudely) at most n^{a_j} possibilities for M_j as there are at most n components of $G_{T(R)}$; from Lemma 7.1.9, $a_j = |M_j| \leqslant 2n(\log_2 n)^{-2}$, so the number of possibilities for M_j is at most

$$\sum_{a_j=0}^{\lfloor 2n(\log_2 n)^{-2} \rfloor} n^{a_j} \leqslant \left(1 + \frac{2n}{(\log_2 n)^2}\right) n^{2n(\log_2 n)^{-2}} \leqslant 3n \cdot n^{2n(\log_2 n)^{-2}},$$

[1]We need not fix this set for this proof, as we will consider every vertex outside of $T(R)$.

for sufficiently large n. Now take a $C \in M_j$ and choose an arbitrary $v \in C$; as the maps $(f_j|_{T(R)})_{j\in[n]}$ and $(f_i)_{i\in T^+(R)}$ have been determined already, we have from Lemma 7.2.7 that the restriction $f_j|_C$ is determined entirely by $(v)f_j$. There are at most n possibilities for this vertex, and so considering the $|M_j|$ components making up Y_j, there are at most $n^{|M_j|} \leqslant n^{2n(\log_2 n)^{-2}}$ possibilities for the restriction $f_j|_{Y_j}$.

Now note that $\mathbf{M}$ and $\triangleright|_Y$ are determined by M_j and $f_j|_{Y_j}$ for each $j \in S_{\leqslant\Delta}(R) \setminus T(R)$; as there are at most n such elements regardless of the set $S_{\leqslant\Delta}(R)$, there are at most $(3n)^n n^{2n^2(\log_2 n)^{-2}}$ possibilities for $\mathbf{M}$ and at most $n^{2n^2(\log_2 n)^{-2}}$ possibilities for $\triangleright|_Y$. Combining these numbers, there are at most $2^n n^{(2+\Delta)Ln}(3n)^n n^{4n^2(\log_2 n)^{-2}} = (6n)^n n^{(2+\Delta)Ln+4n^2(\log_2 n)^{-2}}$ possibilities for $\mathcal{I}''(R)$ (as R ranges over all racks of order n, for n sufficiently large) and thus $|\mathcal{I}''(\mathcal{R}_n)| \leqslant (6n)^n n^{(2+\Delta)Ln+4n^2(\log_2 n)^{-2}}$.

We have that $\Delta = (\log_2 n)^3$, so $2 + \Delta \leqslant 2(\log_2 n)^3$ for sufficiently large n. Then for n sufficiently large

$$\begin{aligned}
\log_2\left((6n)^n n^{(2+\Delta)Ln+4n^2(\log_2 n)^{-2}}\right) &= n\log_2(6n) + (\log_2 n)\left((2+\Delta)Ln + \frac{4n^2}{(\log_2 n)^2}\right) \\
&\leqslant n\log_2(6n) + 2n(\log_2 n)^6 + \frac{4n^2}{\log_2 n} \\
&= o(n^2).
\end{aligned}$$

Hence for any $\epsilon > 0$, there exists a positive integer n_2 such that for $n \geqslant n_2$, $|\mathcal{I}''(\mathcal{R}_n)| \leqslant (6n)^n n^{(2+\Delta)Ln+4n^2(\log_2 n)^{-2}} \leqslant 2^{\epsilon n^2}$, proving the result. □

7.3 Maps acting within components of $G_{T(R)}$

7.3.1 Some preparatory results

We will need the following easy claims. The first is standard calculus and will not be proved.

Claim 7.3.1. *Let $\alpha > 0$ and define $g_\alpha : (0,\infty) \to \mathbb{R}$ by $x \mapsto (\log_2 x)/x^\alpha$. Then $g_\alpha(x)$ is increasing for $x < e^{1/\alpha}$, obtains its unique maximum at $x = e^{1/\alpha}$, and is decreasing for $x > e^{1/\alpha}$.* □

Claim 7.3.2. *For real numbers $x, y \geqslant 0$, $xy/3 + x^2/9 \leqslant (x+y)^2/8$.*

Proof. Put $z = x + y$ and consider z fixed; differentiating the function $h : x \mapsto x(z-x)/3 + x^2/9$ gives a single critical point occurring at the value $x = 3z/4$. Then $h(0) = 0$, $h(3z/4) = z^2/8$ and $h(z) = z^2/9$, giving the result. □

The above claims are used to prove this important technical lemma. The notation is chosen to match the quantities defined in the next subsection.

Lemma 7.3.3. *Let n be a positive integer and $(\eta_q)_{q=1}^n$ be a sequence of non-negative integers such that $\sum_{q=1}^n \eta_q = n$. Now put*

$$\zeta = \left(\sum_{p=1}^n \frac{\eta_p}{p}\right)\left(\sum_{q=1}^n \frac{\log_2 q}{q}\eta_q\right).$$

Then $\zeta \leqslant n^2/4$.

Proof. By expanding the sums, we have that

$$\zeta = \sum_{q=1}^n \frac{\log_2 q}{q^2}\eta_q^2 + \sum_{p=1}^{n-1}\sum_{q=p+1}^n \frac{\log_2 p + \log_2 q}{pq}\eta_p\eta_q \tag{7.3.1}$$

and

$$\frac{n^2}{4} = \frac{(\eta_1 + \cdots + \eta_n)^2}{4} = \sum_{q=1}^n \frac{\eta_q^2}{4} + \sum_{p=1}^{n-1}\sum_{q=p+1}^n \frac{\eta_p\eta_q}{2},$$

so if we set $\nu = n^2/4 - \zeta$ we have

$$\begin{aligned}\nu &= \sum_{q=1}^n \left(\frac{1}{4} - \frac{\log_2 q}{q^2}\right)\eta_q^2 + \sum_{p=1}^{n-1}\sum_{q=p+1}^n \left(\frac{1}{2} - \frac{\log_2(pq)}{pq}\right)\eta_p\eta_q \\ &:= \sum_{q=1}^n d_q\eta_q^2 + \sum_{p=1}^{n-1}\sum_{q=p+1}^n c_{p,q}\eta_p\eta_q.\end{aligned}$$

From Claim 7.3.1 with $\alpha = 1$, $g_\alpha(x) = (\log_2 x)/x$ is increasing for $x < e$ and decreasing for $x > e$. We have $g_1(1) = 0$, $g_1(2) = g_1(4) = 1/2$ and $g_1(3) = (\log_2 3)/3 \approx 0.528$, so for all positive integers $u \neq 2, 3, 4$, $(\log_2 u)/u \leqslant (\log_2 5)/5$. Now define $\gamma := 1/4 - (\log_2 5)/10 \approx 0.0178$; we have shown that for $pq = 1$ and $pq \geqslant 5$, the off-diagonal coefficient $c_{p,q} = 1/2 - (\log_2(pq))/pq \geqslant 1/2 - (\log_2 5)/5 = 2\gamma$. As $p < q$, the only remaining off-diagonal coefficients are $c_{1,2}$, $c_{1,3}$ and $c_{1,4}$; note that $c_{1,2} = c_{1,4} = 0$ while $c_{1,3} < 0$.

Similarly, if we set $\alpha = 2$ we have that $g_\alpha(x) = (\log_2 x)/x^2$ is increasing for $x < e^{1/2} \approx 1.649$ and decreasing for $x > e^{1/2}$. As $g_2(1) = 0 < 1/4 = g_2(2)$, we see that for all positive integers $u \neq 2, 3$, $(\log_2 u)/u^2 \leqslant (\log_2 4)/16 = 1/8$. This means that for $q \geqslant 4$, the diagonal coefficient $d_q = 1/4 - (\log_2 q)/q^2 \geqslant 1/8 > \gamma$; of the remaining diagonal coefficients, note that $d_2 = 0$. Hence

$$\begin{aligned}
\nu &= \sum_{q=1}^{3} d_q \eta_q^2 + \sum_{q=2}^{4} c_{1,q}\eta_1\eta_q + \sum_{q=4}^{n} d_q \eta_q^2 + \sum_{q=5}^{n} c_{1,q}\eta_1\eta_q + \sum_{p=2}^{n-1}\sum_{q=p+1}^{n} c_{p,q}\eta_p\eta_q \\
&\geqslant d_1\eta_1^2 + d_3\eta_3^2 + c_{1,3}\eta_1\eta_3 + \gamma\sum_{q=4}^{n}\eta_q^2 + 2\gamma\eta_1\sum_{q=5}^{n}\eta_q + 2\gamma\sum_{p=2}^{n-1}\sum_{q=p+1}^{n}\eta_p\eta_q. \qquad (7.3.2)
\end{aligned}$$

Now note that $d_1 = 1/4$, $d_3 = 1/4 - (\log_2 3)/9 > 0$ and $c_{1,3} - 1/2 - (\log_2 3)/3 < 0$. We can bound the first three terms of (7.3.2) from below by using Claim 7.3.2 with $x = \eta_3$ and $y = \eta_1$; we have that

$$\begin{aligned}
d_1\eta_1^2 + d_3\eta_3^2 + c_{1,3}\eta_1\eta_3 &= \frac{\eta_1^2}{4} + \left(\frac{1}{4} - \frac{\log_2 3}{9}\right)\eta_3^2 + \left(\frac{1}{2} - \frac{\log_2 3}{3}\right)\eta_1\eta_3 \\
&= \frac{\eta_1^2}{4} + \frac{\eta_3^2}{4} + \frac{\eta_1\eta_3}{2} - (\log_2 3)\left(\frac{\eta_1\eta_3}{3} + \frac{\eta_3^2}{9}\right) \\
&\geqslant \frac{(\eta_1+\eta_3)^2}{4} - (\log_2 3)\frac{(\eta_1+\eta_3)^2}{8} \\
&> \gamma(\eta_1+\eta_3)^2,
\end{aligned}$$

as $1/4 - (\log_2 3)/8 \approx 0.0519$. We now have from (7.3.2) that

$$\begin{aligned}
\nu &\geqslant \gamma\eta_1^2 + \gamma\eta_3^2 + 2\gamma\eta_1\eta_3 + \gamma\sum_{q=4}^{n}\eta_q^2 + 2\gamma\eta_1\sum_{q=5}^{n}\eta_q + 2\gamma\eta_2\sum_{q=3}^{n}\eta_q + 2\gamma\sum_{p=3}^{n-1}\sum_{q=p+1}^{n}\eta_p\eta_q \\
&= \gamma\eta_1^2 + 2\gamma\eta_1\left(\eta_3 + \sum_{q=5}^{n}\eta_q\right) + 2\gamma\eta_2\sum_{q=3}^{n}\eta_q + \gamma\sum_{q=3}^{n}\eta_q^2 + 2\gamma\sum_{p=3}^{n-1}\sum_{q=p+1}^{n}\eta_p\eta_q \\
&\geqslant \gamma\eta_1^2 + 2\gamma\eta_1\sum_{q=3}^{n}\eta_q - 2\gamma\eta_1\eta_4 + \gamma\left(\sum_{q=3}^{n}\eta_q\right)^2 \\
&= \gamma\left(\eta_1 + \sum_{q=3}^{n}\eta_q\right)^2 - 2\gamma\eta_1\eta_4 \\
&= \gamma(n-\eta_2)^2 - 2\gamma\eta_1\eta_4.
\end{aligned}$$

Now by the AM–GM inequality, $\eta_1\eta_4 \leqslant (\eta_1+\eta_4)^2/4 \leqslant (n-\eta_2)^2/4$, and so

$$\frac{n^2}{4} - \zeta = \nu \geqslant \gamma(n-\eta_2)^2 - \gamma\frac{(n-\eta_2)^2}{2} = \frac{\gamma}{2}(n-\eta_2)^2 \geqslant 0, \qquad (7.3.3)$$

proving the result. □

The form of the proof also gives a corresponding stability result.

Corollary 7.3.4. *Let $(\eta_q)_{q=1}^{\infty}$ be a sequence of positive, integer-valued sequences such that $\sum_{q=1}^{n} \eta_q(n) = n$ for all n. Now define a new sequence $\zeta : \mathbb{N} \to \mathbb{R}$ by*

$$\zeta(n) = \left(\sum_{p=1}^{n} \frac{\eta_p(n)}{p} \right) \left(\sum_{q=1}^{n} \frac{\log_2 q}{q} \eta_q(n) \right).$$

Then $\zeta(n) \leqslant n^2/4$, and further, $\zeta(n) \sim n^2/4$ if and only if $\eta_2(n) \sim n$.

Proof. We need only prove the stability statement. First assume that $\eta_2(n) \sim n$; then from (7.3.1) and (7.3.3), $n^2/4 \sim \eta_2(n)^2/4 \leqslant \zeta(n) \leqslant n^2/4$, and so $\zeta(n) \sim n^2/4$. Now suppose that $\zeta(n) \sim n^2/4$; from (7.3.3),

$$0 \leqslant 2\gamma \left(\frac{n - \eta_2(n)}{n} \right)^2 \leqslant 1 - \frac{4\zeta(n)}{n^2} \to 0$$

as $n \to \infty$, and thus $1 - \eta_2(n)/n \to 0$ as $n \to \infty$, i.e. $\eta_2(n) \sim n$. $\square$

7.3.2 A tight upper bound

In this section, we will consider all racks $R \in \mathcal{R}_n$ such that $\mathcal{I}(R) = I$ for some fixed $I \in \mathfrak{I}_n$; in Section 7.2 we showed that the number of possible 7-tuples I is at most $2^{o(n^2)}$.

We will first consider the set $\mathfrak{I}_n$ in more detail. From Definition 7.1.11, $\mathfrak{I}_n$ is a set of 7-tuples, where each $I = (I_1, \ldots, I_7) \in \mathfrak{I}_n$ has the form described below.

1. I_1 is a set $S^I \subseteq [n]$;
2. I_2 is an $(n - |S^I|)$-tuple (σ_i^I) of elements of $\mathfrak{S}_n$, indexed by $\tilde{S}^I := [n] \setminus S^I$ in some arbitrary order;
3. I_3 is a subset T^I of $[n]$ such that $T^I \subseteq S^I$;
4. I_4 is an n-tuple $(\tau_i^I)_{i=1}^n$ of injective maps from T^I to $[n]$;
5. I_5 is a sequence (σ_i^I) of elements of $\mathfrak{S}_n$, indexed by the set $(T^I)^+ := T^I \cup (T^I)\tau_1^I \cup \cdots \cup (T^I)\tau_n^I$ in some arbitrary order.

$\mathcal{I}_6(R)$ and $\mathcal{I}_7(R)$ are graphical in nature; to relate them to an abstract $I \in \mathfrak{I}_n$ we will formally define the graph associated with such a 7-tuple I.

Definition 7.3.5. Let $I \in \mathfrak{I}_n$. Write $\Sigma^I = \{\sigma_i^I \mid i \in T^I\}$ and define $G_I := G_{\Sigma^I}$, in the sense of Definition 4.3.5, so G_I is a $|T^I|$-edge-coloured multigraph on $[n]$. We write c^I for $\mathrm{cp}(G_I)$ and $\mathcal{C}^I$ for the set of vertex sets of components of G_I, so $|\mathcal{C}^I| = c^I$.

We can now describe the form of I_6 and I_7, namely:

6. I_6 is a sequence (M_j^I) of subsets of $\mathcal{C}^I$, indexed by $S^I \setminus T^I$ in some arbitrary order;

7. I_7 is a sequence (ψ_j^I) indexed by $S^I \setminus T^I$, where $Y_j^I = \bigcup_{C\in M_j^I} C$ and $\psi_j^I \in \mathrm{Sym}(Y_j^I)$ for all j.

To avoid later inconvenience, we will extend the definition of M_j^I and Y_j^I to all $j \in [n]$ as follows. For each $j \in \tilde{S}^I \cup T^I$, define a set of ordered pairs from $[n]$ (which we can think of as edges of colour j) by setting

$$\overrightarrow{E}_j^I := \{\overrightarrow{xy} \mid (x)\sigma_j^I = y,\ x \neq y\}. \tag{7.3.4}$$

Then we define $M_j^I = M(G_I, \overrightarrow{E}_j^I)$ and $Y_j^I = \bigcup_{C\in M_j^I} C$.

Now fix an $I \in \mathfrak{I}_n$ and suppose $R \in \mathcal{R}_n$ is a rack such that $\mathcal{I}(R) = I$. Recall that $\Delta = (\log_2 n)^3$ and

$$\mathcal{I}(R) = \left(S_{\leqslant\Delta}(R),\ \triangleright|_{[n]\times S_{>\Delta}},\ T(R),\ \triangleright|_{T(R)\times[n]},\ \triangleright|_{[n]\times T^+(R)},\ \mathbf{M},\ \triangleright|_Y\right),$$

where each entry defined by $\triangleright$ can also be determined in terms of the maps $f_1, \dots, f_n$ corresponding to R. Comparing $\mathcal{I}(R)$ with I we see that $S_{\leqslant\Delta}(R) = S^I$, $S_{>\Delta}(R) = \tilde{S}^I$ and $T(R) = T^I$. We also have that $f_j = \sigma_j^I$ for $j \in \tilde{S}^I \cup (T^I)^+$, and thus $G_{T(R)} = G_I$. Finally, $M_j^I = M_j$ and $Y_j^I = Y_j$ for all $j \in [n]$, and $f_j|_{Y_j} = \psi_j^I$ for $j \in S^I \setminus T^I$.

This means that if the information $I = \mathcal{I}(R)$ is known, we need only determine the maps $(f_j|_{Z_j})_{j\in S_{\leqslant\Delta}(R)\setminus T(R)}$ to determine the entire rack R (as described above, the various sets are determined by I; $S_{\leqslant\Delta}(R) = S^I$, $T(R) = T^I$ and $Z_j = [n]\setminus Y_j = [n]\setminus Y_j^I$ for each $j \in [n]$). It follows that an upper bound on the number of possibilities for these maps is also an upper bound for the number of racks R such that $\mathcal{I}(R) = I$.

We can further reduce the number of maps left to determine by considering components of the graph G_I.

Lemma 7.3.6. *Let $I \in \mathfrak{I}_n$ be fixed and let $R \in \mathcal{R}_n$ be a rack such that $\mathcal{I}(R) = I$. Let $\mathcal{C}^I = \{C_1, \dots, C_{c^I}\}$ and let $V = \{v_1, \dots, v_{c^I}\}$ be a set of vertices such that $v_j \in C_j$ for each j. Then R is determined by I and the maps $(f_v|_{Z_v})_{v\in V}$.*

Proof. As $\mathcal{I}(R) = I$, we have that $G_{T(R)} = G_I$ and so $\mathcal{C}(T(R)) = \mathcal{C}^I$. Take any $j = 1, \dots, c^I$; the knowledge of I determines the maps $(f_i)_{i\in T(R)} = (\sigma_i^I)_{i\in T'}$ and thus from Lemma 7.2.4 (with $A = [n]$) knowledge of the map f_{v_j} is sufficient to determine the maps $(f_u)_{u\in C_j}$. Applying this over all the components of $G_{T(R)}$ shows that we can determine the entire rack R by determining the set of maps $(f_v)_{v\in V}$. As the restrictions $(f_v|_{Y_v})_{v\in V}$ are determined by I ($f_v|_{Y_v} = f_v|_{Y_v^I}$ is equal to ψ_v^I if $v \in S^I \setminus T^I$ and $\sigma_v^I|_{Y_v^I}$ otherwise), we need only determine the restrictions $(f_v|_{Z_v})_{v\in V}$. □

Before proving an upper bound we will need some more notation. For $I \in \mathfrak{I}_n$ and $1 \leqslant q \leqslant n$, let η_q^I denote the number of vertices in components of G_I of size exactly q. Then there are η_q^I/q components of size exactly q and thus

$$n = \sum_{q=1}^{n} \eta_q^I \quad \text{and} \quad c^I = \sum_{q=1}^{n} \frac{\eta_q^I}{q}.$$

Proposition 7.3.7. *Let $I \in \mathfrak{I}_n$ and define*

$$\zeta^I = \left(\sum_{p=1}^{n} \frac{\eta_p^I}{p}\right)\left(\sum_{q=1}^{n} \frac{\log_2 q}{q}\eta_q^I\right).$$

Then there are at most 2^{ζ^I} racks $R \in \mathcal{R}_n$ such that $\mathcal{I}(R) = I$.

Proof. Write $\mathcal{C}^I = \{C_1, \dots, C_{c^I}\}$ and let $V = \{v_1, \dots, v_{c^I}\}$ be a set of vertices such that $v_j \in C_j$ for each j. Let $R \in \mathcal{R}_n$ be a rack such that $\mathcal{I}(R) = I$; from Lemma 7.3.6, R is determined by I and the maps $(f_v|_{Z_v})_{v\in V}$. It follows that an upper bound on the number of possibilities for the maps $(f_v|_{Z_v})_{v\in V}$ is also an upper bound for the number of racks R such that $\mathcal{I}(R) = I$.

Now fix an element $v \in V$ (so f_v is determined by I on all merged components). Let $C \subseteq [n]$ span a component of $G_{T(R)} = G_I$ with $C \notin M_v$, and let $x \in C$; from Lemma 7.1.10 (with $W = T(R)$), $(C)f_v = C$ and so in particular there are at most $|C|$ possibilities for $(x)f_v$. Now $(f_j|_{T(R)})_{j\in[n]} = (\tau_j^I)_{j\in[n]}$ and $(f_i)_{i\in T^+(R)} = (\sigma_i^I)_{i\in(T^I)^+}$ are determined by $I = \mathcal{I}(R)$, so from Lemma 7.2.7 $f_v|_C$ is determined by $(x)f_v$. Thus there are at most $|C|$ possibilities for $f_v|_C$.

As there are η_q^I/q components of $G_{T(R)} = G_I$ of size q, there are at most $\Pi := \prod_{q=1}^{n} q^{\eta_q^I/q}$ possibilities for the map $f_v|_{Z_v}$. Considering all of the c^I components together, there are at most Π^{c^I} possibilities for the maps $(f_v|_{Z_v})_{v\in V}$; it follows that there are at most Π^{c^I} racks R such that $\mathcal{I}(R) = I$. Now

$$\begin{aligned}
\log_2\left(\Pi^{c^I}\right) &= c^I \sum_{q=1}^{n} \frac{\eta_q^I}{q} \log_2 q \\
&= \left(\sum_{p=1}^{n} \frac{\eta_p^I}{p}\right)\left(\sum_{q=1}^{n} \frac{\eta_q^I}{q} \log_2 q\right) \\
&= \zeta^I,
\end{aligned}$$

proving the result. □

This proposition gives us the final result.

Theorem 7.3.8. *Let $\epsilon > 0$ and let $\mathcal{R}'_n$ denote the set of all isomorphism classes of racks of order n. Then for all sufficiently large integers n,*

$$2^{(1/4-\epsilon)n^2} \leqslant |\mathcal{R}'_n| \leqslant 2^{(1/4+\epsilon)n^2}.$$

Proof. The lower bound is exactly Corollary 5.2.6; as $|\mathcal{R}'_n| \leqslant |\mathcal{R}_n|$ it suffices to find an upper bound on $|\mathcal{R}_n|$. Let $\epsilon > 0$; we have from Corollary 7.2.6 and Proposition 7.2.8 that for n sufficiently large $|\mathcal{I}'(\mathcal{R}_n)| \leqslant 2^{\epsilon n^2/2}$ and $|\mathcal{I}''(\mathcal{R}_n)| \leqslant 2^{\epsilon n^2/2}$. Now take such a sufficiently large n; as $\mathcal{I}(R) = (\mathcal{I}'(R), \mathcal{I}''(R))$ for any $R \in \mathcal{R}_n$, $|\mathcal{I}(\mathcal{R}_n)| \leqslant |\mathcal{I}'(\mathcal{R}_n)||\mathcal{I}''(\mathcal{R}_n)| \leqslant 2^{\epsilon n^2}$, and so there are at most $2^{\epsilon n^2}$ possibilities for a 7-tuple $I \in \mathcal{I}(\mathcal{R}_n)$. From Proposition 7.3.7 and Lemma 7.3.3, there are at most $2^{\zeta^I} \leqslant 2^{n^2/4}$ racks R such that $\mathcal{I}(R) = I$, and thus $|\mathcal{R}_n| \leqslant 2^{\epsilon n^2} 2^{n^2/4}$. $\square$

7.4 Racks with given component sizes

Recall that if $\mathcal{P}$ is a partition of $[n]$ then $\mathcal{R}_{\mathcal{P}}$ denotes the set of all racks such that the partition of $[n]$ into orbits of the operator group of R is exactly $\mathcal{P}$ (or alternatively such that $\mathcal{C}([n]) = \mathcal{P}$). Proposition 5.2.3 and the following discussion show that for any $\epsilon > 0$, there exists a positive integer n_0 such that for $n \geqslant n_0$ the following is true: if $\mathcal{P}$ is a partition of $[n]$ into pairs, then $|\mathcal{R}_{\mathcal{P}}| \geqslant 2^{(1/4-\epsilon)n^2}$. In this section we show that if $\mathcal{P}$ is significantly different to a partition of $[n]$ into pairs then $|\mathcal{R}_{\mathcal{P}}|$ is much smaller than $2^{n^2/4}$.

7.4.1 Another graph associated with I

Recall that for any $I \in \mathfrak{I}_n$ we define a corresponding graph $G_I = G_{\Sigma^I}$, where $\Sigma^I = \{\sigma_i^I \mid i \in T^I\}$; G_I is equal to $G_{T(R)}$ if R is a rack on $[n]$ such that $\mathcal{I}(R) = I$. By using some more of the information I we can construct another edge-coloured multigraph on $[n]$.

First define $\tilde{\Sigma}^I = \{\sigma_j \mid j \in \tilde{S}^I \cup T^I\}$ and put $G'_I = G_{\tilde{\Sigma}^I}$ (in the sense of Definition 4.3.5); note from (7.3.4) that

$$G'_I = G_I + \biguplus_{j \in \tilde{S}^I} \overrightarrow{E}_j.$$

For each $j \in S^I \setminus T^I$, define a set of ordered pairs from Y_j^I (which we can think of as directed edges of colour j) by setting

$$\overrightarrow{F}_j^I := \{\overrightarrow{xy} \mid (x)\psi_j^I = y,\ x \neq y\},$$

analogously to Definition 4.3.5. Now define

$$H_I = G'_I + \biguplus_{j \in S^I \setminus T^I} \overrightarrow{F}^I_j.$$

Now suppose R is a rack on $[n]$ such that $\mathcal{I}(R) = I$; then $G_{S_{>\Delta}(R) \cup T(R)} = G'_I$. Further ,for each $j \in [n]$, let $\overrightarrow{E}'_j$ be the subset of $\overrightarrow{E}_j$ consisting only of edges of colour j with both endpoints in Y_j, so that $\overrightarrow{F}^I_j = \overrightarrow{E}'_j$. Then we have

$$H_I = G'_I + \biguplus_{j \in S^I \setminus T^I} \overrightarrow{F}^I_j \;=\; G_{S > \Delta(R) \cup T(R)} + \biguplus_{j \in S_{\leqslant \Delta} \setminus T} \overrightarrow{E}'_j \;\subseteq\; G_R, \tag{7.4.1}$$

where we have suppressed the argument R for improved readability. Hence for each $j \in [n]$, every edge in H_I of colour j corresponds to an edge of colour j in G_R; this correspondence is not one-to-one, but we now show that the graphs G_R and H_I have the same components.

Lemma 7.4.1. *Let $I \in \mathfrak{I}_n$ and $R = ([n], \triangleright)$ be a rack such that $I(R) = I$. Let $C \subseteq [n]$. Then C spans a component of G_R if and only if C spans a component of H_I.*

Proof. From (7.4.1), $H_I \subseteq G_R$; it follows that if $u, v \in [n]$ are in the same component of H_I then they are also in the same component of G_R.

Now write

$$\mathcal{E} = \bigcup_{j \in S^I \setminus T^I} \overrightarrow{E}_j \setminus \overrightarrow{F}^I_j = \bigcup_{j \in S^I \setminus T^I} \overrightarrow{E}_j \setminus \overrightarrow{E}'_j;$$

we have from (7.4.1) again that $G_R = H_I + \mathcal{E}$. Let $\overrightarrow{xy}$ be an edge from $\mathcal{E}$ of colour j; by definition, $\overrightarrow{xy} \notin \overrightarrow{E}'_j$ and thus $x, y \in Z_j$. As Z_j consists of vertices in components of $G_{T(R)}$ *not* merged by $\overrightarrow{E}_j$, $y = (x)f_j$ is in the same component of $G_{T(R)}$ as x, and thus from Lemma 4.3.7 there is a directed xy-path in $G_{T(R)} = G_I \subseteq H_I$. As j was arbitrary, every edge in $\mathcal{E}$ can be replaced by a path from H_I, and thus every two vertices in the same $[n]$-orbit are joined by a path in H_I, proving the result. □

Fix some $I \in \mathfrak{I}_n$; we will consider the graphs G_I and H_I together. Let $B \subseteq [n]$ span a component of H_I; as $G_I \subseteq H_I$ the component $H_I[B]$ is a disjoint union of components of G_I. Suppose that $B = C_1 \cup \cdots \cup C_r$, where each C_i spans a component of G_I. As B spans a component of H_I we can reorder the C_is such that for each $1 < k \leqslant r$ there exists a $1 \leqslant j < r$ such that C_j and C_k are joined by an edge of H_I of some colour and direction. We will use this fact to formally fix an ordering of $\mathcal{C}^I$.

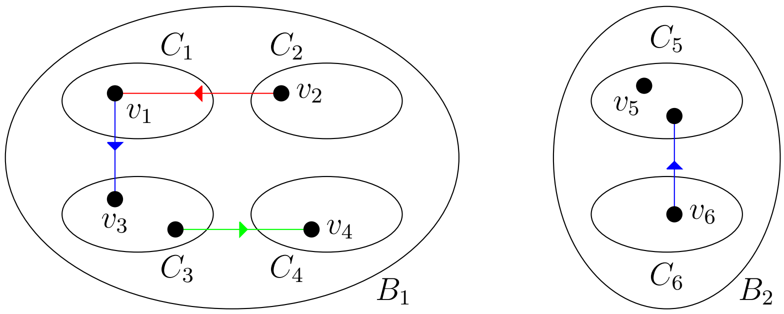

Figure 7.2: An example where $c^I = 6$ and $m^I = 2$. The first set B_1 contains the first four components of $\mathcal{C}^I$, and the second set B_2 the next two. Here $r_1 = 4$ and $r_2 = 6$, so $J_1 = \{1, 5\}$. The coloured edges are edges from H_I.

Write $m^I = \operatorname{cp}(H_I)$ and let $\mathcal{B}^I = \{B_1, \ldots, B_{m^I}\}$ be the set of vertex sets of components of H_I. Write $\mathcal{C}^I = \{C_1, \ldots, C_{c^I}\}$; as described above, each B_i is a disjoint union of sets from $\mathcal{C}^I$. We can reorder $\mathcal{C}^I$ such that the sets from $\mathcal{C}^I$ making up B^I are *consecutive*, i.e. that

$$B_i = \bigcup_{j=1}^{r_i - r_{i-1}} C_{r_{i-1}+j},$$

for some strictly increasing sequence $r_0 = 0, r_1, \ldots, r_{m^I} = c^I$ of indices, where we have the additional property that for any i and any $r_{i-1} < k \leqslant r_i$ there exists some $r_{i-1} \leqslant j < r_i$ such that there is an edge between C_j and C_k of some colour and direction; let v_k be the endpoint of this edge lying in C_k. Now for $i = 0, \ldots, m^I - 1$, choose v_{r_i+1} to be an arbitrary vertex of C_{r_i+1}; we have thus determined a set of vertices $V = \{v_1, \ldots, v_{c^I}\}$ with $v_l \in C_l$ for all l. We will also define

$$J_1 := \{r_i + 1 \mid 0 \leqslant i < m^I\},$$

so J_1 indexes the first component of G_I in each component of H_I (see Figure 7.2). Now let $I \in \mathfrak{I}_n$ be fixed and take a rack $R \in \mathcal{R}_n$ such that $\mathcal{I}(R) = I$; recall that $G_{T(R)} = G_I$ and that $M_l = M_l^I$ and $Y_l = Y_l^I$ for all $l \in [n]$. For $1 \leqslant j \leqslant c^I$, partition the set Z_{v_j} into

$$P_j := \bigcup_{\substack{C \in \mathcal{C}(T(R)) \setminus M_{v_j} \\ |C|=2}} C \quad \text{and} \quad Q_j := \bigcup_{\substack{C \in \mathcal{C}(T(R)) \setminus M_{v_j} \\ |C| \neq 2}} C.$$

As both P_j and Q_j are disjoint unions of f_{v_j}-invariant sets, both P_j and Q_j are f_{v_j}-invariant. We will bound the number of possibilities for the maps $(f_{v_j}|_{P_j})_{j=1}^{c^I}$ by considering maps f_{v_j} and f_{v_k} together if v_j and v_k lie in the same component of the larger graph H_I.

Lemma 7.4.2. *Let $I \in \mathfrak{I}_n$ be fixed and take a rack $R \in \mathcal{R}_n$ such that $\mathcal{I}(R) = I$. Then with notation as above, the maps $(f_{v_j}|_{Z_{v_j}})_{j=1}^{c^I}$ are determined by I and the maps $(f_{v_j}|_{P_j})_{j \in J_1}$ and $(f_{v_j}|_{Q_j})_{j=1}^{c^I}$.*

Proof. Fix some i, $0 \leqslant i < m^I$; we will show that the maps $(f_{v_k}|_{Z_{v_k}})_{k=r_i+1}^{r_{i+1}}$ are determined by I and the maps $f_{v_{r_i+1}}|_{P_{r_i+1}}$ and $(f_{v_k}|_{Q_k})_{k=r_i+1}^{r_{i+1}}$, and the full result then follows by considering all values of i. Next fix a k with $r_i < k \leqslant r_{i+1}$; we will show that the maps $(f_{v_l}|_{P_l})_{l=r_i+1}^{k-1}$, along with the assumed information, determine $f_{v_k}|_{P_k}$. The result then follows inductively.

By our construction of V, v_k is the endpoint of an edge from H_I whose other endpoint, say u, lies in C_j for some j such that $r_i \leqslant j < r_{i+1}$. Let l denote the colour of this edge, so $\overrightarrow{uv_k}$ or $\overrightarrow{v_k u}$ is an edge of H_I of colour l; as $\mathcal{I}(R) = I$, $v_k = (u)f_l^\epsilon$ for some $\epsilon \in \{-1, 1\}$. Now u and v_k are in different components of $G_{T(R)} = G_I$, so $l \notin T(R)$. We also have that $u \in C_j$, so from Lemma 4.3.7, there is a directed $v_j u$-path in $G_{T(R)}$; let the colours of the edges on this path be $i_1, \dots, i_p$, so $u = (v_j)f_{i_1} \cdots f_{i_p}$ and

$$v_k = (u)f_l^\epsilon = (v_j)f_{i_1} \cdots f_{i_p} f_l^\epsilon = (v_j)g,$$

say. Then from Lemma 4.1.10 we have that $f_{v_k} = g^{-1} f_{v_j} g$. As $i_q \in T(R)$ for $1 \leqslant q \leqslant p$, each map $f_{i_q} = \sigma_{i_q}^I$ is determined by I; if $l \in S_{>\Delta(R)} = \tilde{S}^I$ then $f_{v_l} = \sigma_{v_l}^I$ is also determined by I, so g is determined entirely. In this case $f_{v_k} = g^{-1} f_{v_j} g$ is determined by I, $f_{v_j}|_{P_j}$ and $f_{v_j}|_{Q_j}$, so we may suppose that $l \in S_{\leqslant \Delta(R)} \setminus T(R) = S^I \setminus T^I$. In this case the map f_l is *not* fully determined by I, but we do know the set $M_l = M_l^I$ and the restriction $f_{v_l}|_{Y_l} = \psi_l^I$.

Now take $C \in P_k$, so by definition C is a $T(R)$-orbit of size 2 such that $(C)f_{v_k} = C$; as C is a $T(R)$-orbit we have that C is i_q-invariant for $1 \leqslant q \leqslant p$. First assume that $C \in \mathcal{C}(T(R)) \setminus M_l$; from Lemma 7.1.10, C is l-invariant, and so $(C)f_l^\epsilon = C$ for $\epsilon \in \{-1, 1\}$. It follows that g is a composition of maps preserving C and thus $(C)g^\epsilon = C$ for $\epsilon \in \{-1, 1\}$. But then $(C)f_{v_j} = (C)g f_{v_k} g^{-1} = C$, and so C is also f_{v_j}-invariant.

We have shown that $g|_C, f_{v_j}|_C \in \mathrm{Sym}(C)$; as $|C| = 2$ the group $\mathrm{Sym}(C) \cong C_2$, which is abelian. It follows that

$$f_{v_k}|_C = g^{-1}|_C \circ f_{v_j}|_C \circ g|_C = f_{v_j}|_C,$$

and thus $f_{v_k}|_C$ is determined by $f_{v_j}|_{P_j}$.

Now suppose that $C \in M_l$; as $G_{T(R)} \subseteq G_{T(R) \cup \{l\}}$ there exists a set $D \supseteq C$ such that D is a disjoint union of components of $G_{T(R)}$ and D spans a component of

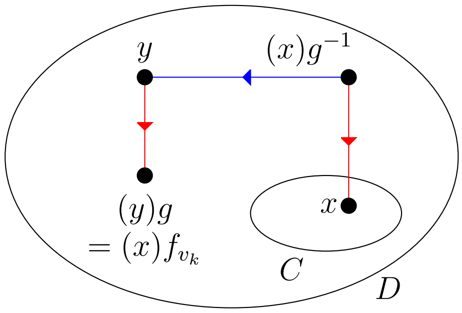

Figure 7.3: The red lines represent g, which is not determined everywhere outside of D. We must ensure that the blue edge (of colour v_j) maps $(x)g^{-1}$ to some other vertex in D.

$G_{T(R)\cup\{l\}}$. Let $C' \subseteq D \setminus C$ span such a component of $G_{T(R)}$, and let $x \in C$ and $y \in C'$ be arbitrary; from Lemma 4.3.7 there is a directed xy-path $x = x_0, \ldots, x_m = y$ in $G_{T(R)\cup\{l\}}[D]$. Now let t be minimal such that $x_t \in C'$; then $t > 0$ and $x_{t-1} \notin C'$, so the edge $\overrightarrow{x_{t-1}x_t}$ is of colour l and $C' \in M(G_{T(R)}, \overrightarrow{E}_l) = M_l$. As C' was arbitrary we conclude that each $T(R)$-orbit contained in D lies in M_l, and so $D \subseteq Y_l$; it follows that $f_l|_D = \psi_l^I|_D$ (which is determined by I), and as D spans a component of $G_{T(R)\cup\{l\}}$, $(D)(\psi_l^I)^\epsilon = D$ for $\epsilon \in \{-1, 1\}$. As each $T(R)$-orbit is i_q-invariant for $1 \leqslant q \leqslant p$, D is also i_q-invariant for each q, and so $g^\epsilon|_D$ is determined by I for $\epsilon \in \{-1, 1\}$, with $(D)g^\epsilon = D$.

Let $x \in C$ be arbitrary; we will now show that the vertex $(x)f_{v_k}$ is determined by I and f_{v_j} (see Figure 7.3). We have shown that $(x)g^{-1}$ is determined by I, with $(x)g^{-1} \in D$. Suppose that $y := (x)g^{-1}f_{v_j} \notin D$; then $(x)f_{v_k}g^{-1} \notin D$, which is a contradiction as $(x)f_{v_k} \in (C)f_{v_k} = C \subseteq D$ and $(D)g^{-1} = D$. Hence $y \in D$, and so the vertex $(y)g$ is determined by I and lies in D. As $y = (x)g^{-1}f_{v_j}$ is determined by I and f_{v_j}, we get the desired result, and as x was arbitrary it follows that $(C)f_{v_j}$ is determined by I and f_{v_j}. This proves the lemma as we have determined f_{v_k} on every component of $G_{T(R)}$ lying within P_k. $\square$

We can now give a bound on the number of racks R such that $\mathcal{I}(R) = I$ in terms of m^I and η_q^I for $q \geqslant 2$; Lemma 7.4.1 then allows us to relate this bound back to the given component structure of G_R.

Proposition 7.4.3. *Let $I \in \mathfrak{I}_n$, and define*

$$\zeta_2^I = m^I \frac{\eta_2^I}{2} + n \sum_{q=3}^{n} \frac{\log_2 q}{q} \eta_q^I.$$

Then there are at most $2^{\zeta_2^I}$ racks $R \in \mathcal{R}_n$ such that $\mathcal{I}(R) = I$.

Proof. We will use the same notation as the previous lemma. Let $R \in \mathcal{R}_n$ be a rack such that $\mathcal{I}(R) = I$; from Lemma 7.3.6, R is determined by the maps $(f_{v_j}|_{Z_{v_j}})_{j=1}^{c^I}$. From Lemma 7.4.2 these maps are determined by I and and the maps $(f_{v_j}|_{P_j})_{j\in J_1}$ and $(f_{v_j}|_{Q_j})_{j=1}^{c^I}$, so it suffices to find an upper bound on the number of possibilities for the maps $(f_{v_j}|_{P_j})_{j\in J_1}$ and $(f_{v_j}|_{Q_j})_{j=1}^{c^I}$.

Now fix a j such that $1 \leqslant j \leqslant c^I$; we will consider the possibilities for the map $f_{v_j}|_{Q_j}$ much as before. Take a $C \in \mathcal{C}(T(R)) \setminus M_{v_j}$ such that $|C| \neq 2$; for an arbitrary $x \in C$ there are at most $|C|$ possibilities for $(x)f_{v_j}$. We argue as in the proof of Proposition 7.3.7; as the maps $(f_j|_{T(R)})_{j\in[n]} = (\tau_j^I)_{j\in[n]}$ and $(f_i)_{i\in\Gamma_R^+(T)} = (\sigma_i^I)_{i\in(T^I)^+}$ are determined by I, we have from Lemma 7.2.7 that there are at most $|C|$ possibilities for $f_{v_j}|_C$. As Q_j is a union of $T(R)$-orbits of size not equal to 2 and there are η_q^I/q $T(R)$-orbits of size q for each q, the number of possibilities for the maps $(f_{v_j}|_{Q_j})_{j=1}^{c^I}$ is at most $(\Pi_2)^{c^I}$, where $\Pi_2 = \prod_{q=3}^n q^{\eta_q^I/q}$.

Now fix a $j \in J_1$; there are at most 2 options for the restriction of the map f_{v_j} to any component of size 2. P_j is a disjoint union of $T(R)$-orbits of size 2, and there are $\eta_2^I/2$ such orbits, so there are at most $2^{\eta_2^I/2}$ possibilities for the map $f_{v_j}|_{P_j}$. Now $|J_1| = m^I$ and thus there are at most $2^{m^I\eta_2^I/2}$ possibilities for the maps $(f_{v_j}|_{P_j})_{j\in J_1}$. Hence there are at most $2^{m^I\eta_2^I/2}\Pi_2^{c^I}$ racks $R \in \mathcal{R}_n$ such that $\mathcal{I}(R) = I$; note that

$$\begin{aligned}
\log_2\left(2^{m^I\eta_2^I/2}\Pi_2^{c^I}\right) &= m^I\frac{\eta_2^I}{2} + c^I\sum_{q=3}^{n}\frac{\log_2 q}{q}\eta_q^I \\
&\leqslant m^I\frac{\eta_2^I}{2} + n\sum_{q=3}^{n}\frac{\log_2 q}{q}\eta_q^I \\
&= \zeta_2^I,
\end{aligned}$$

showing the result. □

7.4.2 An extremal result

We will now show that unless $\mathcal{P}$ is 'almost' a partition of $[n]$ into pairs, the size of $(\log_2|\mathcal{R}_\mathcal{P}|/n^2)$ is bounded away from 1/4. Formally, we need to consider a *sequence* of partitions $(\mathcal{P}_n)_{n=1}^\infty$, where $\mathcal{P}_n$ is a partition of $[n]$ for each n. For each positive integer q, let $m_q(n)$ denote the number of parts of $\mathcal{P}_n$ of size exactly q and let $m(n)$ denote the total number of parts, so

$$m(n) = \sum_{q=1}^{n} m_q(n) \quad \text{and} \quad n = \sum_{q=1}^{n} q m_q(n)$$

for all n.

Proposition 7.4.4. *Let $(\mathcal{P}_n)_{n=1}^{\infty}$ be a sequence of partitions, where $\mathcal{P}_n$ is a partition of $[n]$ for all n, and suppose that*

$$\frac{\log_2 |\mathcal{R}_{\mathcal{P}_n}|}{n^2} \to \frac{1}{4}$$

as $n \to \infty$. Then $m_2(n) \sim n/2$.

Proof. Let $\gamma = 1/4 - (\log_2 5)/10$ and let

$$A = \frac{2}{\gamma}\left(1 - \frac{\log_2 3}{3}\right).$$

Now let $\epsilon \in (0, A)$; this range is chosen such that $\gamma\epsilon^2/2 + \epsilon(\log_2 3)/3 < \epsilon$. Put $\delta = \gamma\epsilon^2/8$; from Corollary 7.2.6 and Proposition 7.2.8 there exists some positive integer n_0 such that for all $n \geqslant n_0$, $|\mathcal{I}(\mathcal{R}_n)| \leqslant 2^{\delta n^2}$.

Now for any $n \geqslant n_0$, partition the set $\mathfrak{I}_n$ into two sets as follows: let

$$\mathfrak{I}_n^{(1)} := \{I \in \mathfrak{I}_n \mid \eta_2^I \leqslant (1-\epsilon)n\} \quad \text{and} \quad \mathfrak{I}_n^{(2)} := \{I \in \mathfrak{I}_n \mid \eta_2^I > (1-\epsilon)n\}.$$

Now partition $\mathcal{R}_{\mathcal{P}_n}$ accordingly; let

$$\mathcal{R}_{\mathcal{P}_n}^{(1)} := \{R \in \mathcal{R}_{\mathcal{P}_n} \mid \mathcal{I}(R) \in \mathfrak{I}_n^{(1)}\} \quad \text{and} \quad \mathcal{R}_{\mathcal{P}_n}^{(2)} := \{R \in \mathcal{R}_{\mathcal{P}_n} \mid \mathcal{I}(R) \in \mathfrak{I}_n^{(2)}\},$$

so $|\mathcal{R}_{\mathcal{P}_n}| = |\mathcal{R}_{\mathcal{P}_n}^{(1)}| + |\mathcal{R}_{\mathcal{P}_n}^{(2)}|$. We will find an upper bound on $|\mathcal{R}_{\mathcal{P}_n}^{(1)}|$ by using Proposition 7.3.7; for any $I \in \mathfrak{I}_n^{(1)}$ there are at most 2^{ζ^I} racks R such that $\mathcal{I}(R) = I$. Now for such I, $\eta_2^I \leqslant (1-\epsilon)n$, so from (7.3.3),

$$\zeta^I \leqslant \frac{n^2}{4} - \frac{\gamma}{2}(n - \eta_2^I)^2 \leqslant \frac{n^2}{4} - \frac{\gamma(\epsilon n)^2}{2} = \left(\frac{1}{4} - \frac{\gamma\epsilon^2}{2}\right)n^2.$$

As $|\mathcal{I}(\mathcal{R}_{\mathcal{P}_n}^{(1)})| \leqslant |\mathcal{I}(\mathcal{R}_{\mathcal{P}_n})| \leqslant 2^{\delta n^2}$, we have that $|\mathcal{R}_{\mathcal{P}_n}^{(1)}| \leqslant 2^{\alpha_1 n^2}$, where

$$\alpha_1 = \delta + \frac{1}{4} - \frac{\gamma\epsilon^2}{2} = \frac{1}{4} - \frac{3\gamma\epsilon^2}{8}.$$

Now let $R \in \mathcal{R}_{\mathcal{P}_n}$ and put $\mathcal{I}(R) = I$. By definition, η_2^I is the number of vertices in components of $G_I = G_{T(R)}$ of size exactly 2; as $G_{T(R)} \subseteq G_R$, every such vertex must be in a component of G_R of size at least 2. There are precisely $m_1(n)$ vertices of G_R in components of size 1, so we have that $\eta_2^I \leqslant n - m_1(n)$. Hence if $m_1(n) \geqslant \epsilon n$ then $\eta_2^I \leqslant n - \epsilon n = (1-\epsilon)n$, and thus $I \in \mathfrak{I}_n^{(1)}$; it follows in this case that $\mathcal{R}_{\mathcal{P}_n}^{(2)} = \emptyset$ and $|\mathcal{R}_{\mathcal{P}_n}| = |\mathcal{R}_{\mathcal{P}_n}^{(1)}| \leqslant 2^{\alpha_1 n^2}$. Then

$$\frac{\log_2 |\mathcal{R}_{\mathcal{P}_n}|}{n^2} \leqslant \alpha_1 = \frac{1}{4} - \frac{3\gamma\epsilon^2}{8}. \tag{7.4.2}$$

Now if there are infinitely many n such that $m_1(n) \geqslant n$ then (7.4.2) holds for infinitely many n, contradicting the fact that $(\log_2 |\mathcal{R}_{\mathcal{P}_n}|)/n^2 \to 1/4$ as $n \to \infty$. Hence there exists some positive integer $n_0' > n_0$ such that for $n \geqslant n_0'$, $m_1(n) < \epsilon n$; this suffices to show that $m_1(n) = o(n)$.

Now suppose that $n \geqslant n_0'$, so $m_1(n) < \epsilon n$; we will find an upper bound on $|\mathcal{R}^{(2)}_{\mathcal{P}_n}|$ by using Proposition 7.4.3, and thus we need to consider ζ_2^I. Let $I \in \mathfrak{I}_n^{(2)}$ and let $R \in \mathcal{R}_{\mathcal{P}_n^{(2)}}$ be such that $\mathcal{I}(R) = I$; from Lemma 7.4.1, $m^I = \mathrm{cp}(H_I) = \mathrm{cp}(G_R) = m(n)$ as $R \in \mathcal{R}_{\mathcal{P}_n}$. Now from Claim 7.3.1 (and the proof of Lemma 7.3.3) $(\log_2 q)/q \leqslant (\log_2 3)/3$ for any positive integer q; hence

$$\zeta_2^I = m^I \frac{\eta_2^I}{2} + n \sum_{q=3}^{n} \frac{\log_2 q}{q} \eta_q^I \leqslant m(n) \frac{\eta_2^I}{2} + \frac{\log_2 3}{3} n \sum_{q=3}^{n} \eta_q^I.$$

As $I \in \mathfrak{I}_n^{(2)}$ we have that $\eta_2^I > (1-\epsilon)n$; as $n = \sum_{q=1}^n \eta_q^I$, it follows that $\sum_{q=3}^n \eta_q^I = n - \eta_2^I - \eta_1^I < \epsilon n$. Thus, approximating η_2^I by n,

$$\zeta_2^I \leqslant \frac{nm(n)}{2} + \frac{\epsilon \log_2 3}{3} n^2$$

for any $I \in \mathfrak{I}_n^{(2)}$. Now for any such I we have from Proposition 7.4.3 that there are at most $2^{\zeta_2^I}$ racks R such that $\mathcal{I}(R) = I$; as $|\mathcal{I}(\mathcal{R}^{(2)}_{\mathcal{P}_n})| \leqslant |\mathcal{I}(\mathcal{R}_{\mathcal{P}_n})| \leqslant 2^{\delta n^2}$, we have that $|\mathcal{R}^{(2)}_{\mathcal{P}_n}| \leqslant 2^{\alpha_2 n^2}$, where

$$\alpha_2 = \frac{\delta n^2 + \zeta_2^I}{n^2} \leqslant \frac{m(n)}{2n} + \frac{\gamma \epsilon^2}{8} + \frac{\epsilon \log_2 3}{3}.$$

Now suppose that $m(n)/2n < 1/4 - (\gamma\epsilon^2)/2 - \epsilon(\log_2 3)/3$; then $\alpha_2 < \alpha_1$. It follows that $|\mathcal{R}_{\mathcal{P}_n}| = |\mathcal{R}^{(1)}_{\mathcal{P}_n}| + |\mathcal{R}^{(2)}_{\mathcal{P}_n}| \leqslant 2^{\alpha_1 n^2} + 2^{\alpha_2 n^2} < 2^{\alpha_1 n^2 + 1}$, and so

$$\frac{\log_2 |\mathcal{R}_{\mathcal{P}_n}|}{n^2} < \alpha_1 + \frac{1}{n^2} = \frac{1}{4} - \frac{3\gamma\epsilon^2}{8} + \frac{1}{n^2} \leqslant \frac{1}{4} - \frac{\gamma\epsilon^2}{4} \tag{7.4.3}$$

for sufficiently large n. If there are infinitely many n such that $m(n)/2n < 1/4 - (\gamma\epsilon^2)/2 - \epsilon(\log_2 3)/3$ then (7.4.3) holds for infinitely many n, contradicting the fact that $(\log_2 |\mathcal{R}_{\mathcal{P}_n}|)/n^2 \to 1/4$ as $n \to \infty$. Hence there exists some positive integer $n_0'' > n_0'$ such that for $n \geqslant n_0''$, $m(n)/2n \geqslant 1/4 - (\gamma\epsilon^2)/2 - \epsilon(\log_2 3)/3$. As ϵ was chosen such that $(\gamma\epsilon^2)/2 + \epsilon(\log_2 3)/3 < \epsilon$, $m(n)/2n > 1/4 - \epsilon$ for $n \geqslant n_0''$.

Now take $n \geqslant n_0''$, so $m_1(n) < \epsilon n$ and $m(n)/2n > 1/4 - \epsilon$, so $m(n) > n/2 - 2\epsilon n$. Recall that $m(n) = \sum_{q=1}^n m_q(n)$ and $n = \sum_{q=1}^n q m_q(n)$; we thus have that

$$2\epsilon n > \frac{n}{2} - m(n) = \sum_{q=1}^{n} \left(\frac{q}{2} - 1\right) m_q(n) > \frac{-\epsilon n}{2} + \sum_{q=3}^{n} \left(\frac{q}{2} - 1\right) m_q(n). \tag{7.4.4}$$

Now $6(q/2-1) \geqslant q$ for $q \geqslant 3$, so multiplying (7.4.4) by $6/n$ gives

$$0 \leqslant \sum_{q=3}^{n} \frac{q m_q(n)}{n} \leqslant 6 \sum_{q=3}^{n} \left(\frac{q}{2} - 1\right) \frac{m_q(n)}{n} \leqslant 15\epsilon,$$

which suffices to show that $\sum_{q=3}^{n} q m_q(n) = o(n)$. But then $n = m_1(n) + 2m_2(n) + \sum_{q=3}^{n} q m_q(n) = 2m_2(n) + o(n)$, and thus $2m_2(n)/n = 1 + o(1)$, i.e. $m_2(n) \sim n/2$. $\square$

Chapter 8

Connected kei

There has been a high level of interest in connected racks in recent years; in [11] Grana classified all racks of order p^2 (for a prime number p), while McCarron showed in [25] that for a prime number $p \neq 2, 3, 5$ there is no connected quandle of order $2p$. In this chapter we exploit an interesting property of the graphical representation of a kei to show that the number of connected kei of order n is at most $2^{o(n^2)}$; McCarron's result shows that there is no non-trivial lower bound.

Because we are working exclusively with *kei* on $[n]$, we will omit the inverse notation for the corresponding maps $f_1, \ldots, f_n$. Recall that if K is a kei then the edges of G_K are undirected, so we may adapt Definition 7.1.1 and write $b(v) = b^+(v) = b^-(v)$ and $d(v) = d^+(v) = d^-(v)$ for any $v \in [n]$.

8.1 Diameter of graphical representations

Let $([n], \triangleright)$ be a *kei* and $f_1, \ldots, f_n$ be the associated maps. For any $S \subseteq [n]$ and $u, v \in [n]$, denote by $d_S(u, v)$ the graph distance between u and v in the graph G_S. Recall that the reduced graph G_S^0 is the simple graph on $[n]$ formed by ignoring all colours and multiple edges; $d_S(u, v)$ is clearly identical to the graph distance between u and v in the reduced graph G_S^0. We will show that any component with diameter d must have at least 2^d vertices, and also at least 2^{d-1} edges of the same colour. We begin by showing that there are many shortest paths between pairs of vertices.

Lemma 8.1.1. *Let $S \subseteq [n]$ be a subkei, and let $u, v \in [n]$ such that $d_S(u, v) = d \in [2, \infty)$; hence there exists a shortest path $u = v_0, v_1, \ldots, v_d = v$ of length d in G_S between u and v. For each i, let c_i be the colour of the $v_{i-1}v_i$ edge on the path, so $(v_i)f_{c_i} = v_{i-1}$. Now fix some $i \geqslant 2$ and define $w_k = (v_k)f_{c_i}$ for $k = 0, \ldots, i-2$. Then*

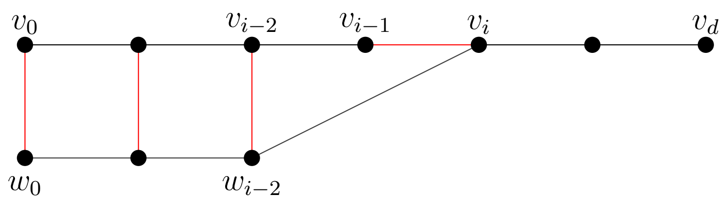

Figure 8.1: An example with $i = 4$. The colour c_i is represented by red, with the colours on the lower path being the colours c', c'' etc. found as we progress from right to left along the upper path.

the vertices $w_0, \ldots, w_{i-2}, v_0, \ldots, v_d$ are all distinct and $v_0, w_0, \ldots, w_{i-2}, v_i, \ldots, v_d$ is a (shortest) uv-path in G_S.

Proof. (See Figure 8.1 throughout). For any $j \leqslant i-2$, denote by $P(j)$ the statement that the vertices $w_j, \ldots, w_{i-2}, v_0, \ldots, v_d$ are all distinct, and that the path $v_0, \ldots, v_j, w_j, \ldots, w_{i-2}, v_i, \ldots, v_d$ is a (shortest) uv-path in G_S. Then the lemma is the statement $P(0)$; we will prove that $P(j)$ holds for all j by reverse induction.

Firstly, as $w_{i-2} = (v_{i-2})f_{c_i}$ and $v_i = (v_{i-1})f_{c_i}$ we see that $w_{i-2} \notin \{v_{i-1}, v_i\}$ (or there would be a vertex with two incident edges of the same colour). If we had $w_{i-2} = v_l$ for some $l > i$ then $v_0, \ldots, v_{i-2}, v_l, \ldots, v_d$ would be a uv-path in G_S of length $d-(l-i+1) < d$, a contradiction. Now observe that

$$w_{i-2} := (v_{i-2})f_{c_i} = (v_{i-1})f_{c_{i-1}}f_{c_i} = (v_i)f_{c_i}f_{c_{i-1}}f_{c_i} = (v_i)f_{c'},$$

where $c' = (c_{i-1})f_{c_i}$; note that $c' \in S$ as S is a subkei and so there is an edge in G_S between v_i and w_{i-2}. So now suppose $w_{i-2} = v_l$ for $l < i-1$; then $v_0, \ldots, v_l = w_{i-2}, v_i, \ldots, v_d$ is a uv-path in G_S of length $d-(i-l-1)$, again a contradiction. So $w_{i-2} \neq v_l$ for any l, and we can use the $v_{i-2}w_{i-2}$ and $w_{i-2}v_i$ edges, of colours c_i and c' respectively, to construct a uv-path $v_0, \ldots, v_{i-2}w_{i-2}v_i, \ldots, v_d$ of length d in G_S. This establishes $P(i-2)$.

So now take $j < i-2$ and suppose $P(j+1)$ holds. As $v_{j+1}w_{j+1}, \ldots, v_{i-2}w_{i-2}$ and $v_{i-1}v_i$ are all edges of colour c_i, $w_j \neq v_l$ for $j < l \leqslant i$ and it is also not equal to any other w_k. As before, if $w_j = v_l$ for $l > i$ we can construct a uv-path in G_S of length $d-(l-i+1)$, a contradiction. Now observe that

$$w_j := (v_j)f_{c_i} = (v_{j+1})f_{c_{j+1}}f_{c_i} = (w_{j+1})f_{c_i}f_{c_{j+1}}f_{c_i} = (w_{j+1})f_{c''},$$

where $c'' = (c_{j+1})f_{c_i} \in S$, so there is an edge in G_S between w_j and w_{j+1}. Now suppose $w_j = v_l$ for some $l \leqslant j$; then $v_0, \ldots, v_l = w_j, w_{j+1}, \ldots, w_{i-2}v_i, \ldots, v_d$ is a

uv-path in G_S of length $d-(j-l+1)$, a contradiction (note the existence of the $w_{j+1}v_i$-path follows from the assumed $P(j+1)$). Hence $w_j \neq v_l$ for any l, and we can use the v_jw_j and w_jw_{j+1} edges, of colours c_i and c'' respectively, to construct a uv-path $v_0,\ldots,v_j,w_j,\ldots,w_{i-2},v_i,\ldots,v_d$ of length d in G_S. Thus $P(j)$ holds, and so by reverse induction $P(0)$ holds, proving the lemma. □

Corollary 8.1.2. *Let P be any shortest uv-path in G_S. Then no two edges of P have the same colour.*

Proof. Suppose not, so there is a path $u=v_0,\ldots,v_d=v$ and a colour $c \in S$ such that $(v_{j-1})f_c = v_j$ and $(v_{i-1})f_c = v_i$ for some $j < i-1$ (we can't have $j = i-1$ as then v_j is incident to two edges of colour c). From the lemma, with $c = c_i$, $w_j = (v_j)f_{c_i} \neq v_l$ for any l, but here $w_j = (v_j)f_c = v_{j-1}$, a contradiction. □

Suppose we are considering a shortest path $P = v_0,\ldots,v_d$ in G_S. In the light of Corollary 8.1.2, and to ease notation, we will assume without loss of generality that $[d] \subseteq S$ and that the edge between v_{i-1} and v_i in P is of colour i. Now let $\mathbf{s} = (a_1,\ldots,a_r)$ be a strictly increasing sequence of elements from $[d]$ and define

$$u_{\mathbf{s}} = (u)f_{a_1}\cdots f_{a_r},$$

where we note $u_{\mathbf{s}} \in S$ as S is a subkei. Now define $U_0 = \{u\}$ and for $r = 1,\ldots,d$, define

$$U_r = \{u_{\mathbf{s}} \mid \mathbf{s} \text{ is strictly increasing}, |\mathbf{s}| = r\},$$

so in particular $U_d = \{v\}$. Now let $\mathbf{e}_i = (1,\ldots,i)$ for all i; then $v_i = (u)f_1\cdots f_i = u_{\mathbf{e}_i}$ and so $v_i \in U_i$ for all i. We also have the property that as $\mathbf{s}$ is increasing, $a_i \geqslant i$ for any i, with equality if and only if the subsequence consisting of the first i terms of $\mathbf{s}$ is the canonical sequence $\mathbf{e}_i$.

We will show that any strictly increasing sequence $\mathbf{s}$ can appear at the start of a shortest uv-path without affecting too many other edges, and that any such path passes sequentially through each of $U_0,U_1,\ldots,U_d$.

Lemma 8.1.3. *Suppose we are in the same situation as Lemma 8.1.1 with $c_i = i$ for $1 \leqslant i \leqslant d$, and let $\mathbf{s} = (a_1,\ldots,a_r)$ be a strictly increasing sequence of elements from $[d]$. Then there is a shortest path $P_{\mathbf{s}} = (u = x_0, x_1, \ldots, x_{a_r-1}, x_{a_r} = v_{a_r}, \ldots, v_d)$ such that $x_k \in U_k$ for $1 \leqslant k < a_r$ and the $x_{k-1}x_k$ edge in $P_{\mathbf{s}}$ is of colour a_k for $1 \leqslant k \leqslant r$ (which means in particular that $x_r = u_{\mathbf{s}}$).*

Proof. We shall prove the following stronger statement by induction on r: there exists a shortest path $P_{\mathbf{s}} = (u = x_0, x_1, \ldots, x_{a_r-1}, x_{a_r} = v_{a_r}, \ldots, v_d)$ such that for $1 \leqslant k < a_r$,

$$x_k = u_{\mathbf{t}} \tag{8.1.1}$$

for a (strictly increasing) sequence $\mathbf{t}$ of length k whose largest element is at most a_r, and that for $1 \leqslant k \leqslant r$ the $x_{k-1}x_k$ edge in $P_{\mathbf{s}}$ is of colour a_k. These statements clearly imply the result.

First consider the base case $r = 1$; if $\mathbf{s} = \mathbf{e}_1 = (1)$ then the original path P suffices as the first edge is of colour 1, so $v_1 = (u)f_1 = u_{(1)}$. Hence we may assume $a_1 > 1$. Then applying Lemma 8.1.1 to P with $i = a_1$, there exists a shortest path $P' = (u = v_0, w_0, \ldots, w_{a_1-2}, v_{a_1}, \ldots, v_d)$ in G_S, with an edge of colour a_1 between v_k and w_k for $0 \leqslant k \leqslant a_1 - 2$; in particular, the first edge of P' is of colour a_1. Now for $k \leqslant a_1 - 2$, $w_k = (v_k)f_{a_1} = (u_{\mathbf{e}_k})f_{a_1}$, so $w_k = u_{(\mathbf{e}_k, a_1)} \in U_{k+1}$. So we obtain (8.1.1) by setting $x_k = w_{k-1}$ for $1 \leqslant k < a_1$ and putting $P_{\mathbf{s}} = P'$.

So now take $r > 1$ and assume the result for smaller r. If $a_r = r$ we have $\mathbf{s} = \mathbf{e}_r$, and in this situation the original path P suffices as the first r edges are of colours $1, \ldots, r$ and $v_k \in U_k$ for $1 \leqslant k \leqslant r$. So we can assume that $a_r > r$.

By applying the inductive hypothesis to the sequence $\mathbf{s}' = (a_1, \ldots, a_{r-1})$ we see that there exists a path

$$P_{\mathbf{s}'} = (u = y_0, y_1, \ldots, y_{a_{r-1}-1}, y_{a_{r-1}} = v_{a_{r-1}}, \ldots, v_d)$$

such that, for $1 \leqslant k < a_{r-1}$, $y_k = u_{\mathbf{t}}$ for a strictly increasing sequence $\mathbf{t}$ whose largest element is at most a_{r-1}. We also have that the $y_{k-1}y_k$ edge is of colour a_k for $1 \leqslant k < r$, but we don't know the colours of the edges $y_{k-1}y_k$ for $r \leqslant k \leqslant a_{r-1}$. But as $a_r > a_{r-1}$ the $v_{a_{r-1}}v_{a_r}$ edge of colour a_r is still present in $P_{\mathbf{s}'}$; hence we can apply Lemma 8.1.1 to $P_{\mathbf{s}'}$, with $i = a_r$, to obtain a shortest path $P'' = (u = v_0, w_0, \ldots, w_{a_r-2}, v_{a_r}, \ldots, v_d)$ in G_S, with an edge of colour a_r between y_k and w_k for $0 \leqslant k \leqslant \min\{a_{r-1}, a_r - 2\}$ (note that $a_r - 2 < a_{r-1}$ if and only if $a_r = a_{r-1} + 1$). We aim to show that for all $0 \leqslant k \leqslant a_r - 2$, $w_k = u_{\mathbf{t}'}$ for some sequence $\mathbf{t}'$ of length $k+1$ whose largest element is at most a_r. The proof will vary slightly depending on whether $a_r = a_{r-1} + 1$ (see Figure 8.2) or $a_r > a_{r-1} + 1$ (see Figure 8.3).

Fix a k such that $0 \leqslant k \leqslant a_{r-1} - 1 \leqslant a_r - 2$; by the inductive hypothesis (specifically (8.1.1)) $y_k = u_{\mathbf{t}}$ for some sequence $\mathbf{t}$ whose kth and largest element is at most a_{r-1}. But $w_k = (y_k)f_{a_r} = (u_{\mathbf{t}})f_{a_r}$, so $w_k = u_{(\mathbf{t}, a_r)}$; as $a_r > a_{r-1}$ we have $w_k \in U_{k+1}$, where $w_k = u_{\mathbf{t}'}$ for some sequence t' whose largest element is at most a_r. If $a_r = a_{r-1} + 1$ then we have considered all the vertices $w_1, \ldots, w_{a_r-2}$.

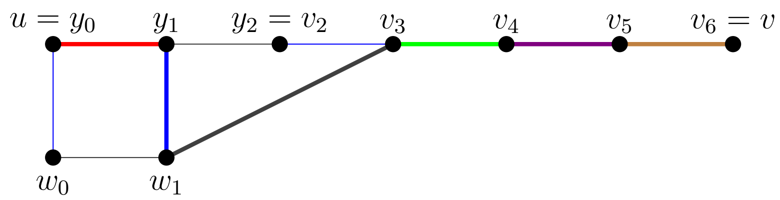

Figure 8.2: Suppose $\mathbf{s} = (2, 3)$ *and that edges of colour* 2 *are red while those of colour* 3 *are blue. The top path* $P_{\mathbf{s}'}$ *corresponds to the sequence* (2)*, and the bottom path* P'' *is a result of applying Lemma 8.1.1. The new path* P''' *is shown in bold.*

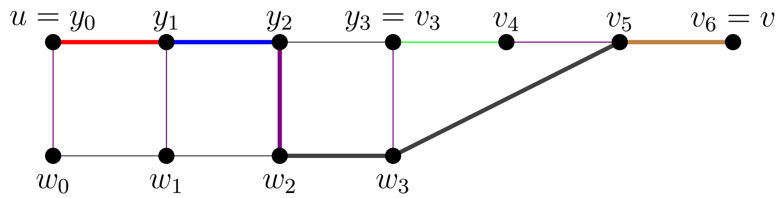

Figure 8.3: Suppose we now have the sequence $\mathbf{s} = (2, 3, 5)$*, with edges of colour 5 being violet. The top path* $P_{s'}$ *is the (relabelled) bold path from Figure 8.2, while the bottom path is again a result of applying Lemma 8.1.1; note the additional edges. The new path* P''' *is shown in bold.*

Now suppose that $a_r > a_{r-1} + 1$; in this case we also have an edge of colour a_r between v_k and w_k for $a_{r-1} \leqslant k \leqslant a_r - 2$. But as before, $w_k = (v_k)f_{a_r} = (u_{\mathbf{e}_k})f_{a_r} = u_{(\mathbf{e}_k, a_r)}$ for any such k, so $w_k \in U_{k+1}$ and $w_k = u_{\mathbf{t}'}$ for some sequence t' whose largest element is at most a_r. Thus we have the desired result on w_k for the entire range $0 \leqslant k \leqslant a_r - 2$.

Now as $a_r > r$ and $a_{r-1} \geqslant r - 1$ we have $\min\{a_{r-1}, a_r - 2\} \geqslant r - 1$, so there is always an edge of colour a_r between y_{r-1} and w_{r-1}. Consider the path P''' obtained by using the first $r - 1$ edges of $P_{\mathbf{s}'}$, the edge $y_{r-1}w_{r-1}$ and then the remaining edges of P''; we obtain

$$P''' = (u, y_1, \ldots, y_{r-1}, w_{r-1}, \ldots, w_{a_r-2}, v_{a_r}, \ldots, v_d),$$

where the first r edges are of colours $a_1, \ldots, a_r$. We get the result by putting $x_k = y_k \in U_k$ for $1 \leqslant k < r$ and $x_k = w_{k-1} \in U_k$ for $r \leqslant k < a_r$, and setting $P_{\mathbf{s}} = P'''$. $\square$

We can now use this result to show that different sequences correspond to distinct vertices.

Corollary 8.1.4. *Let* $\mathbf{s}$ *and* $\mathbf{t}$ *be distinct, strictly increasing sequences from* $[d]$*. Then* $u_{\mathbf{s}} \neq u_{\mathbf{t}}$*.*

Proof. Let $|\mathbf{s}| = q$ and $|\mathbf{t}| = r$, where we may assume $q \leqslant r$, and suppose that $u_\mathbf{s} = u_\mathbf{t}$. Write $\mathbf{s} = (a_1, \ldots, a_q)$ and $\mathbf{t} = (b_1, \ldots, b_r)$ and consider the paths $P_\mathbf{s} : (ux_1 \cdots x_{a_q} = v_{a_q} \cdots v_d)$ and $P_\mathbf{t} : (uy_1 \cdots y_{b_r} = v_{b_r} \cdots v_d)$ as described in Lemma 8.1.3; note that $x_q = u_\mathbf{s} = u_\mathbf{t} = y_r$. Thus we may replace the uy_r-segment of $P_\mathbf{t}$ with the ux_q-segment of $P_\mathbf{s}$ to obtain a uv-walk in G_S of length $d - (r - q)$, a contradiction if $q < r$. This shows that $U_q \cap U_r = \emptyset$ for $q \neq r$.

So suppose $q = r$ (so $x_r = y_r$) and for now that $a_r \neq b_r$; we may assume that $a_r < b_r$. Now replace the ux_r-segment of $P_\mathbf{s}$ with the uy_r-segment of $P_\mathbf{t}$ to obtain a new uv-path

$$P' : (uy_1 \cdots y_r = x_r \cdots x_{a_r - 1} x_{a_r} = v_{a_r} \cdots v_{b_r - 1} v_{b_r} \cdots v_d),$$

where we note this is a path as both $P_\mathbf{s}$ and $P_\mathbf{t}$ pass sequentially through the pairwise disjoint sets $U_0, U_1, \ldots, U_d$. But both the $y_{r-1}y_r$ and $v_{b_r-1}v_{b_r}$ edges on P' are of colour b_r, contradicting Corollary 8.1.2.

Now note that in any case $x_{r-1} = (x_r)f_{a_r}$ and $y_{r-1} = (y_r)f_{b_r}$, so if $a_r = b_r$ we have $x_{r-1} = y_{r-1}$. An easy inductive argument (and the fact that $\mathbf{s} \neq \mathbf{t}$) shows that there exists some $p \leqslant r$ such that $x_i = y_i$ for $p \leqslant i \leqslant r$ but $a_p \neq b_p$, so we may apply the above argument to the sequences $\mathbf{s}' = (a_1, \ldots, a_p)$ and $\mathbf{t}' = (b_1, \ldots, b_p)$ to get a contradiction, thus proving the result. □

We can now state the main result of this section.

Theorem 8.1.5. *Let $([n], \triangleright)$ be a kei, and let $(S, \triangleright)$ be a subkei. Let $C \in \mathcal{C}(S)$ and suppose that $G_S[C]$ has diameter d. Then $|C| \geqslant 2^d$, and further there exists some $k \in S$ such that there are at least 2^{d-1} edges of colour k in $G_S[C]$.*

Proof. If $d = 0$ or 1 then the result is trivial, so suppose $d \geqslant 2$. $G_S[C]$ has diameter d, so there exist $u, v \in C$ such that $d_S(u, v) = d$; as before, assume $[d] \subseteq S$ and that there is a shortest uv-path using edges of colours $1, \ldots, d$ sequentially. There are 2^d strictly increasing sequences of elements from $[d]$ (including the empty sequence) and from Corollary 8.1.4 these correspond to 2^d distinct vertices (including u). Hence $|C| \geqslant 2^d$.

Note that there are 2^{d-1} strictly increasing sequences from $[d-1]$; adding d to the end of such a sequence $\mathbf{s}'$ gives a strictly increasing sequence $\mathbf{s}$ from $[d]$. But then $(u_{\mathbf{s}'})f_d = u_\mathbf{s}$ so the edge $u_{\mathbf{s}'}u_\mathbf{s}$ is of colour d for each such sequence $\mathbf{s}'$; as all vertices are distinct this gives us 2^{d-1} edges of colour d. □

Note that this result is tight for an infinite class of subkei. Consider Construction 5.2.2 with $k_i = d$ for all i and some $m \geqslant d+1$; P_1 is an $[n]$-orbit and so a subkei. Without loss of generality we can take $[d] \subseteq P_1$; now choose the maps $\phi_1^{(P_2)}, \ldots, \phi_d^{(P_2)}$ to form the canonical framework of P_2. Thus P_2 is a component in the graph G_{P_1} isomorphic to the cube Q_d, which has diameter d. But $|P_2| = 2^d$, and each $\phi_i^{(P_2)}$ is an involution on P_2 with no fixed points, so acts as 2^{d-1} transpositions; hence there are 2^{d-1} edges of each colour i.

8.2 Degree and diameter

8.2.1 A bound on the diameter

We showed in Lemma 7.1.2 that if $([n], \triangleright)$ is a rack and $(S, \triangleright)$ is a subrack then each component of G_S^0 is out-regular; if $([n], \triangleright)$ is a kei then all edges are undirected and thus each component of G_S^0 is regular. This result means that whenever C spans a component of G_S^0 (or equivalently G_S), we can refer to the degree of the *graph* $G_S^0[C]$, rather than the degree of any individual vertex.

The following straightforward lemma shows how the diameter and (vertex) degree of a component of G_S are related.

Lemma 8.2.1. *Let $([n], \triangleright)$ be a kei, $(S, \triangleright)$ be a subkei and C span a component of G_S, and hence also of G_S^0. Then the diameter of $G_S^0[C]$ is at most the degree of $G_S^0[C]$.*

Proof. Suppose that $G_S^0[C]$ has diameter d and degree Δ. If $d = 0$ or 1 then the result is trivial, so we may assume that $d \geqslant 2$. By definition there exist $u, v \in C$ such that $d_S(u, v) = d$; as before, assume $[d] \subseteq S$ and that there is a shortest uv-path using edges of colours $1, \ldots, d$ sequentially. Now apply Corollary 8.1.4 to the d singleton sequences $(1), \ldots, (d)$; we have that $u_{(1)}, \ldots, u_{(d)}$ are distinct vertices not equal to u (as this corresponds to the empty sequence). But for each i, $u_{(i)} = (u)f_i$ is a neighbour of u, and so u has at least d neighbours. As $G_S^0[C]$ is regular with degree Δ, we have that $d \leqslant \Delta$. $\square$

8.2.2 Transposition numbers

Let $K = ([n], \triangleright)$ be a connected kei. As all vertices are in the same component, it follows from Lemma 7.2.4 (with $T = A = [n]$) that *all* the maps $f_1, \ldots, f_n$ are conjugate; suppose that each map acts as a product of t disjoint transpositions (we may call t the *transposition number* of the kei). We also have from Lemma 7.1.2 (with

$S = [n]$) that both G_K and G_K^0 are regular; let $\Delta_e(K)$ and $\Delta(K)$ respectively denote the common degrees of these graphs, i.e. the common edge- and vertex-degrees of G_K. From the discussion following Definition 7.1.1, $\Delta(K) \leqslant \Delta_e(K)$.

We will use the result of the previous section to show that given a transposition number t, there is an upper bound on the number of vertices in a connected kei with transposition number t. We will need a small result from graph theory; given a simple graph G with degree at most Δ and diameter at most d, it is straightforward to show that

$$|G| \leqslant 1 + \Delta \sum_{i=0}^{d-1} (\Delta - 1)^i =: M_{\Delta,d}$$

for $\Delta > 2$. The expression $M_{\Delta,d}$ is known as the *Moore bound* (see [16]); if $\Delta = 2$ then clearly $M_{\Delta,d} = 1 + 2d$, while if $\Delta > 2$ we have

$$M_{\Delta,d} = 1 + \frac{\Delta}{\Delta - 2}((\Delta - 1)^d - 1).$$

We can easily find a cruder bound for $M_{\Delta,d}$; we have that

$$M_{\Delta,d} = 1 + \Delta \sum_{i=0}^{d-1} (\Delta - 1)^i \leqslant 1 + \sum_{i=0}^{d-1} \Delta^{i+1} = \frac{\Delta^{d+1} - 1}{\Delta - 1} < \Delta^{d+1}, \tag{8.2.1}$$

valid for $\Delta > 1$.

We can now obtain the following upper bound for the size of a connected kei with transposition number t; the lower bound is straightforward.

Proposition 8.2.2. *Let $K = ([n], \triangleright)$ be a connected kei with transposition number t. Then*

$$2t + 1 \leqslant n \leqslant 1 + \frac{t}{t-1}\left((2t-1)^{1+\lfloor \log_2 t \rfloor} - 1\right), \tag{8.2.2}$$

for $t > 1$.

Proof. For the lower bound, observe that for any $x \in [n]$, f_x acts as t transpositions and fixes the element x, so there are at least $2t$ elements of $[n] \setminus \{x\}$ not fixed by f_x. Hence $n \geqslant 2t + 1$.

For the upper bound, let $d = \operatorname{diam}(G_K) = \operatorname{diam}(G_K^0)$. From Theorem 8.1.5, there are at least 2^{d-1} edges of some colour $i \in [n]$, so we must have $2^{d-1} \leqslant t$, i.e. $d \leqslant 1 + \log_2 t$. Now $\Delta_e(K)$ is the common degree of G_K; from the Handshaking Lemma, $n\Delta_e(K) = 2e(G_K)$. But as each of the n maps acts as t transpositions, there are t edges of each colour, and so $e(G_K) = nt$. Hence $\Delta_e(K) = 2t$, and so

$\Delta(K) \leqslant \Delta_e(K) = 2t$. But the number of vertices in a graph of maximum degree $2t$ and maximum diameter $1 + \lfloor \log_2 t \rfloor$ is bounded above by the Moore bound

$$M_{2t,1+\lfloor \log_2 t \rfloor} = 1 + \frac{2t}{2t-2}\left((2t-1)^{1+\lfloor \log_2 t \rfloor} - 1\right),$$

and the result follows. □

If $t = 1$ the upper bound as given is not valid; as $\Delta = 2t = 2$ we instead conclude that $n \leqslant 1 + 2d = 3$, and it is straightforward to see that $K = D_3$. In the Appendix, we will classify all kei such that each map acts as at most one transposition, connected or otherwise.

8.3 Counting connected kei

We will denote by $\mathcal{CK}_n$ the set of all connected kei on $[n]$. Let $K \in \mathcal{CK}_n$; we can rearrange (8.2.2) to see that there is a non-trivial lower bound on its transposition number t, of the form $2^{C\sqrt{\log_2 n}}$ for some constant C. From the proof of Proposition 8.2.2, $\Delta_e(K) = 2t$, so we obtain a lower bound on the edge degree of a connected kei; we can use Lemma 8.2.1 to give a lower bound on the vertex degree.

Lemma 8.3.1. *Let $K = ([n], \triangleright)$ be a connected kei. Then for sufficiently large n, $\Delta(K) > \log_2 \log_2 n$.*

Proof. Put $\lambda = \log_2 \log_2 n$ and suppose for a contradiction that $\Delta(K) \leqslant \lambda$. Then from Lemma 8.2.1 (with $S = [n]$), $\operatorname{diam}(G_K^0) \leqslant \Delta(K) \leqslant \lambda$, and from (8.2.1), $n = |G_K^0| \leqslant \lambda^{\lambda+1}$. But then

$$\log_2 n \leqslant (\lambda + 1)\log_2 \lambda \leqslant (\lambda+1)^2 = (\log_2 \log_2 n + 1)^2 < \log_2 n$$

for sufficiently large n, a clear contradiction. Hence $\Delta(K) > \log_2 \log_2 n$. □

We will now count connected kei on $[n]$ by first fixing the degree $\Delta(K)$ of the underlying graph G_K^0; let $\mathcal{CK}_{n,\Delta}$ denote the set of all connected kei K on $[n]$ with $\Delta(K) = \Delta$. We have shown that for n sufficiently large, $\mathcal{CK}_{n,\Delta} = \emptyset$ for $\Delta \leqslant \log_2 \log_2 n$.

We now consider a larger range of Δ. If $\Delta(K) > \Delta_0 := (\log_2 n)^3$ then we have, in the notation of Chapter 7,

$$S_{>\Delta_0}(K) = \{v \in [n] \mid d_K^+(v) > \Delta_0\} = \{v \in [n] \mid d_K(v) > (\log_2 n)^3\} = [n],$$

noting that G_K^0 is undirected and regular. We can now appeal to Proposition 7.2.5[1]; for n sufficiently large, K is determined by a set W and the maps $(f_i)_{i\in W}$ for some set $W \subseteq [n]$ such that $|W| \leqslant w_0(n)$. It follows that for n sufficiently large and $\Delta > (\log_2 n)^3$ there are at most $2^n n^{w_0(n)}$ connected kei K on $[n]$ with $\Delta(K) = \Delta$; as in the proof of Corollary 7.2.6,

$$\log_2\left(2^n n^{w_0(n)}\right) \leqslant n + \frac{7}{2}\frac{n^2}{(\log_2 n)^{1/2}} = o(n^2).$$

Hence for any $\epsilon > 0$ and $\Delta > (\log_2 n)^3$, there exists a positive integer n_1 such that for $n \geqslant n_1$, $|\mathcal{CK}_{n,\Delta}| \leqslant 2^{\epsilon n^2}$.

For the remaining values of Δ, we use a similar probabilistic argument to Proposition 7.2.5. We will prove the statement for racks as well as kei.

Proposition 8.3.2. *Let $\log_2 \log_2 n < \Delta \leqslant (\log_2 n)^3$, and let $R = ([n], \triangleright)$ be a connected rack with $d_R^+(v) = \Delta$ for all $v \in [n]$. Then for n sufficiently large there exists a set $W' \subseteq [n]$, with*

$$|W'| \leqslant w_1(n) := \frac{3n}{2(\log_2 \Delta)^{3/2}} + \frac{2n(\log_2 \Delta)^{3/2}}{\Delta} + 2\exp\left(-\frac{\Delta}{8(\log_2 \Delta)^{-3/2}}\right) n,$$

such that R is determined by W' and the maps $(f_i)_{i\in W'}$.

Proof. We will construct W' by a mixture of probabilistic and deterministic methods. Let $p := (\log_2 \Delta)^{-3/2}$ and consider a random subset X of $[n]$ as described in Lemma 7.2.3. For each n, let E_n be the event that $|X| \leqslant 3np/2$; as $np \geqslant n(\log_2 n)^{-3/2} \to \infty$ as $n \to \infty$, we have from item 1 of the lemma (with $\epsilon = 1/2$) that $\mathbb{P}(E_n) \to 1$ as $n \to \infty$.

We will now redefine the concept of a 'bad' vertex, giving a different threshold; for $v \in [n]$ and $U \subseteq [n]$, we now say that v is *U-bad* if $d_U^+(v) \leqslant \Delta p/2$. Let B_v be the event that that v is X-bad, and let $N = N(X)$ denote the number of X-bad vertices in $[n]$, so $N = \sum_{v=1}^n \mathbf{1}_{B_v}$. Now $d_v^+ = d_R^+(v) = \Delta$ for all v, so from item 2 of Lemma 7.2.3 (again with $\epsilon = 1/2$ and $\delta = \Delta p$),

$$\mathbb{E}[N] = \sum_{v=1}^n \mathbb{P}(B_v) = \sum_{v=1}^n \mathbb{P}\big(d_X^+(v) \leqslant \Delta p/2\big) \leqslant e^{-\Delta p/8} n.$$

Now for each n let F_n be the event that $N \leqslant 2e^{-\Delta p/8} n$. From Markov's Inequality,

$$\mathbb{P}(F_n^{\mathrm{c}}) \leqslant \mathbb{P}(N \geqslant 2e^{-\Delta p/8} n) \leqslant \frac{\mathbb{E}[N]}{2e^{-\Delta p/8} n} \leqslant \frac{1}{2},$$

[1]The quantity Δ in this result is defined in Chapter 7 to be $(\log_2 n)^3 = \Delta_0$.

so $\mathbb{P}(F_n) \geqslant 1/2$. Now $\mathbb{P}(E_n \cup F_n) \geqslant \mathbb{P}(E_n) \to 1$ as $n \to \infty$, so $\mathbb{P}(E_n) - \mathbb{P}(E_n \cup F_n) \to 0$ as $n \to \infty$. Then $\mathbb{P}(E_n \cap F_n) = \mathbb{P}(F_n) + \mathbb{P}(E_n) - \mathbb{P}(E_n \cup F_n) > 0$ for sufficiently large n, so we can find a set $U' \subseteq [n]$ with

$$|U'| \leqslant \frac{3np}{2} = \frac{3n}{2(\log_2 \Delta)^{3/2}} =: w_1'(n), \tag{8.3.1}$$

such that the number of U'-bad vertices is at most αn, where $\alpha = \alpha(n) = 2e^{-\Delta p/8}$.

Now let $\mathcal{C}$ be the set of components of $G^0_{U'}$, $\mathcal{C}_g \subseteq \mathcal{C}$ be the subset of components containing at least one U'-good vertex and $\mathcal{C}_b$ be the subset of components consisting entirely of U'-bad vertices; then $\mathcal{C}_g \cap \mathcal{C}_b = \emptyset$ and $\mathcal{C}_g \cup \mathcal{C}_b = \mathcal{C}$. Note that any component in $\mathcal{C}_g$ contains an U'-good vertex, i.e. a vertex with at least $\lceil \Delta p/2 \rceil$ neighbours; hence any such component has size at least $\lceil \Delta p/2 \rceil + 1$ and $|\mathcal{C}_g| \leqslant 2n/\Delta p$. As there are at most αn bad vertices, $|\mathcal{C}_b| \leqslant \alpha n$. Now let $\mathcal{C} = \{C_1, \dots, C_r\}$ and let $V' = \{v_1, \dots, v_r\}$ be a set of vertices such that $v_k \in C_k$ for each k; we have shown that

$$|V'| = |\mathcal{C}| \leqslant \frac{2n}{\Delta p} + \alpha n = \frac{2n(\log_2 \Delta)^{3/2}}{\Delta} + 2\exp\left(-\frac{\Delta}{8(\log_2 \Delta)^{-3/2}}\right) n =: w_1''(n). \tag{8.3.2}$$

Now from Lemma 7.2.4 (with $T = U'$ and $A = [n]$), knowledge of the maps $(f_i)_{i \in U'}$ and f_{v_k} are sufficient to determine the maps $(f_u)_{u \in C_k}$; applying this to each component in turn shows that knowledge of the maps $(f_i)_{i \in U'}$ and $(f_v)_{v \in V'}$ are sufficient to determine R. So put $W' = U' \cup V'$; from (8.3.1) and (8.3.2), $|W'| \leqslant |U'| + |V'| \leqslant w_1'(n) + w_1''(n) = w_1(n)$, proving the result. □

Corollary 8.3.3. *Let $\log_2 \log_2 n < \Delta \leqslant (\log_2 n)^3$, and let $\epsilon > 0$. Then there exists a positive integer n_2 such that for all $n \geqslant n_2$, $|\mathcal{CK}_{n,\Delta}| \leqslant 2^{\epsilon n^2}$.*

Proof. Let $K \in \mathcal{CK}_{n,\Delta}$; we have from the previous result that for n sufficiently large, K is determined by a set W' with $|W'| \leqslant w_1(n)$ and the maps $(f_i)_{i \in W'}$. We can lower the number of possibilities for these maps by first determining the graph G^0_K.

By assumption, the simple graph G^0_K is Δ-regular and thus has $\Delta n/2$ edges; there are (crudely) at most $\binom{n}{2}^{\Delta n/2} < n^{\Delta n} \leqslant n^{n(\log_2 n)^3}$ such labelled graphs on the vertex set $[n]$. There are clearly at most 2^n possibilities for the set W', and the knowledge of G^0_K limits the number of possibilities for the maps $(f_i)_{i \in W'}$. Let $i, v \in [n]$; as G^0_K is determined, $(v)f_i \in \{v\} \cup \Gamma_K(v)$, and so there are at most $1 + d_V = 1 + \Delta$ possibilities for each $(v)f_i$ and at most $(\Delta+1)^n$ possibilities for each map f_i. Putting this together gives that $|\mathcal{CK}_{n,\Delta}| \leqslant n^{n(\log_2 n)^3} 2^n (\Delta+1)^{n w_1(n)} = n^{n(\log_2 n)^3} 2^n (\Delta+1)^{n w_1'(n)} (\Delta+1)^{n w_1''(n)}$, as $|W| \leqslant w_1(n) = w_1'(n) + w_1''(n)$. Then

$$\log_2 \left(n^{n(\log_2 n)^3} 2^n\right) = n(\log_2 n)^4 + n = o(n^2),$$

while for n sufficiently large, $\log_2(\Delta+1) \leqslant 2\log_2\Delta$ and thus

$$\begin{aligned}
\log_2\left((\Delta+1)^{nw_1'(n)}\right) &\leqslant 2(\log_2\Delta)nw_1'(n)\\
&= 2(\log_2\Delta)n\frac{3n}{2(\log_2\Delta)^{3/2}}\\
&= \frac{3n^2}{(\log_2\Delta)^{1/2}}\\
&= o(n^2),
\end{aligned}$$

noting that $\Delta > \log_2\log_2 n \to \infty$ as $n \to \infty$. We also have for sufficiently large n that

$$\begin{aligned}
\log_2\left((\Delta+1)^{nw_1''(n)}\right) &\leqslant 2(\log_2\Delta)nw_1''(n)\\
&= 2(\log_2\Delta)n\left(\frac{2n(\log_2\Delta)^{3/2}}{\Delta} + 2\exp\left(-\frac{\Delta}{8(\log_2\Delta)^{-3/2}}\right)n\right)\\
&= \left(\frac{4n^2(\log_2\Delta)^{5/2}}{\Delta} + 4(\log_2\Delta)\exp\left(-\frac{\Delta}{8(\log_2\Delta)^{-3/2}}\right)n^2\right)\\
&= o(n^2),
\end{aligned}$$

noting additionally that $\Delta(\log_2\Delta)^{-3/2} = \omega(\log_2\log_2\Delta)$. Hence for any $\epsilon > 0$, there exists a positive integer n_2 such that for all $n \geqslant n_2$, $|\mathcal{CK}_{n,\Delta}| \leqslant 2^{\epsilon n^2}$. □

We can now combine these results to prove the main theorem.

Theorem 8.3.4. *Let $\epsilon > 0$. Then there exists a positive integer n_0 such that for $n \geqslant n_0$, $|\mathcal{CK}_n| \leqslant 2^{\epsilon n^2}$.*

Proof. We have from Lemma 8.3.1 that $|\mathcal{CK}_{n,\Delta}| = 0$ if $\Delta \leqslant \log_2\log_2 n$. Let $\epsilon > 0$; we can combine the discussion following Lemma 8.3.1 and Corollary 8.3.3 to see that there exists a positive integer n' such that for $n \geqslant n'$ and $\Delta \geqslant \log_2\log_2 n$, $|\mathcal{CK}_{n,\Delta}| \leqslant 2^{\epsilon n^2/2}$. There is another positive integer n'' such that for $n \geqslant n''$, $n \leqslant 2^{\epsilon n^2/2}$; put $n_0 = \max\{n', n''\}$. Then for $n \geqslant n_0$,

$$|\mathcal{CK}_n| = \sum_{\Delta=\lceil\log_2\log_2 n\rceil}^{n-1} |\mathcal{CK}_{n,\Delta}| \leqslant n2^{\epsilon n^2/2} \leqslant 2^{\epsilon n^2},$$

proving the result. □

We do not know if the above result extends to racks or quandles. Although Proposition 8.3.2 can be generalised to racks, the crucial Lemma 8.2.1, used to prove a lower bound on the vertex degree of a connected kei, does not. For example, we can take a rack R on $[n]$ by setting $f_1 = (2\ \cdots\ n)$ and $f_i = \iota$ for $i > 1$; then G_R^0 consists of a directed cycle on $n-1$ vertices and one isolated vertex. The in-degree and out-degree of the non-trivial component are both 1, whilst the diameter is $\lfloor (n-1)/2 \rfloor$, giving a counterexample to Lemma 8.2.1. It may be possible however that these counterexamples are restricted to disconnected racks.

Appendix A

Miscellaneous results

This Appendix contains smaller miscellaneous results on kei.

A.1 Kei with maps acting as at most one transposition

Let $K = ([n], \triangleright)$ be a kei such that for each $i \in [n]$, the associated map f_i is either the identity or a single transposition. Let $C \subseteq [n]$ be an $[n]$-orbit, so $G_K[C]$ is a component of G_K. Suppose that $G_K[C]$ has diameter d; the from Theorem 8.1.5, there are at least 2^{d-1} edges of some colour i in $G_K[C]$, and thus $d \leqslant 1$. It follows that each component of the reduced graph G^0_K is a (possibly trivial) clique. Suppose one of these cliques were of size at least 4; then we could find a path $wxyz$ with edges of colour j, i and k successively (see Figure A.1; these colours must all be distinct as each colour appears at most once).

Now as in the proof of Lemma 7.1.2, $(j)f_i \in N(y)$; as f_j acts as a single transposition, $j \notin N(y)$ and thus f_i acts non-trivially on j. Similarly, $(k)f_i \in N(x)$ and f_i acts non-trivially on k. Note that f_i and f_k both fix w; then $(w)f_i f_j f_i = y \neq (w)f_k$, and so $f_{(j)f_i} \neq f_k$ and $(j)f_i \neq k$. Hence f_i must act as at least two transpositions (namely $(j\ (j)f_i)$ and $(k\ (k)f_i)$), contradicting our original assumption. Thus each

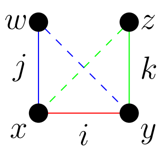

Figure A.1: A putative clique of size 4. The blue dashed edge is of colour $(j)f_i$ while the green dashed edge is of colour $(k)f_i$.

component of G_K^0 is isomorphic to a complete graph on at most three vertices.

Proposition A.1.1. *Let $K = ([n], \triangleright)$ be a kei such that for each i, the associated map f_i acts as either the identity or a single transposition. Then there exists a partition*

$$\mathcal{P} = \{X_1, \ldots, X_r, Y_1, \ldots, Y_s, Z_1, \ldots, Z_t\}$$

of $[n]$, with $|X_i| = 3$ for all i, $|Y_j| = 2$ for all j and $|Z_k| = 1$ for all k, such that K has the properties described below.

1. *For each i, $(X_i, \triangleright)$ is a kei isomorphic to D_3, and for each $p \in X_i$, f_p fixes every element in $[n] \setminus X_i$;*

2. *If we label some $Y_j = \{p, q\}$, then either $f_p = f_q = \iota$ or there exists some $Y_{j'} = \{p', q'\}$, $j' \neq j$, such that $f_p = f_q = (p'\ q')$;*

3. *If we label some $Z_k = \{p\}$, then either $f_p = \iota$ or there exists some $Y_{j'} = \{p', q'\}$ such that $f_p = (p'\ q')$.*

Proof. From the remarks above, no component of G_K can have size four or greater; partition $[n]$ into the sets $(X_i)_{i=1}^r$, $(Y_j)_{j=1}^s$ and $(Z_k)_{k=1}^t$, representing the components (or $[n]$-orbits) of size 3, 2 and 1 respectively. We will use the same language as in Section 6.3, in terms of 'blocks' acting on other blocks. Let T be any block (i.e. any X_i, Y_j or Z_k); from Lemma 6.2.4, any two components of G_T lying entirely within any X_i must be of the same size, and thus there must be 3 components of size 1 (in which case T doesn't act on X_i) or a single component of size three. We will show that X_i must act on itself, with no other blocks acting on X_i.

First suppose $T = Z_k$ for some k, and write $T = \{p\}$. G_T is at most a single edge (representing the map f_p), so each component has size 1 or 2; hence T doesn't act on X_i. Now suppose $T = Y_j$ for some j and write $T = \{p, q\}$; by the orbit structure we have $(q)f_p \in \{p, q\}$. But f_p is a bijection fixing p, so $(q)f_p = q$ and the maps f_p and f_q commute. Then from Proposition 5.1.4, any component of G_T has size a power of 2, and thus T doesn't act on X_i.

Now as in Lemma 6.3.1, each $(X_i, \triangleright)$ is either trivial or isomorphic to D_3, and in the former case X_i acts on no other X_j. Now suppose there is some X_i that doesn't act on itself; from above, only the blocks X_j, $j \neq i$, can act on other blocks of size three. But as each map acts as at most one transposition, each such X_j can act on at most one other block X_k; say without loss of generality that X_r is not acted on by any of $X_1, \ldots, X_r$. But then no block acts on X_r and so X_r is not a component in

G_K; a contradiction. Thus each X_i acts on itself (such that $K_i := (X_i, \triangleright) \cong D_3$) and on no other block. This shows item 1.

Now take an arbitrary j and write $Y_j = \{p, q\}$; we have already shown that $f_p = f_q$. As each map acts as at most one transposition we have shown items 2 and 3. $\square$

Note that we can write $K = K_1 \oplus \cdots \oplus K_r \oplus L$, where each $K_I \cong D_3$ and the kei L is an output of Construction 5.2.2, with the appropriate restrictions.

A.2 Kei with no subkei

Let $R = (X, \triangleright)$ be a rack. Recall that a subset $Y \subseteq X$ forms a subrack $(Y, \triangleright)$ if and only if $(Y)f_y = Y$ for all $y \in Y$. The following lemma shows that for any subset $T \subseteq R$, any disjoint union of components of G_T forms a subrack.

Lemma A.2.1. *Let $R = (X, \triangleright)$ be a rack and let $T \subseteq X$. Let $x_1, \ldots, x_m \in T$, and for each i let C_i be the element of $\mathcal{C}(T)$ containing x_i (so $G_T[C_i]$ is the component of G_T containing x_i). Then $(\bigcup_{i=1}^m C_i, \triangleright)$ is a subrack of R.*

Proof. Let $C = \bigcup_{i=1}^m C_i$ and suppose $y \in C$; then $y \in C_i$ for some i. As $G_T[C_i]$ contains x_i, we have from Lemma 4.3.7 that there is a directed $x_i y$-path in G_T; label the colours of the edges on this path $t_1, \ldots, t_k$. Then $y = (x_i)f_{t_1} \cdots f_{t_k}$ and $f_y = f_{t_k}^{-1} \cdots f_{t_1}^{-1} f_{x_i} f_{t_1} \cdots f_{t_k}$. Now $x_i, t_1, \ldots, t_k \in T$ and so their associated maps all preserve the component structure of G_T; hence $(C_j)f_y = C_j$ for each j and so $(C)f_y = C$. As y was arbitrary we have the result. $\square$

We now return specifically to kei. We say that a subkei $(Y, \triangleright)$ is *trivial* if $|Y| = 1$ and *proper* if $|Y| < |X|$ (or equivalently if $Y \subset X$). The following result classifies all kei with no non-trivial, proper subkei.

Proposition A.2.2. *Let $n \geqslant 3$ and let $K = ([n], \triangleright)$ be a kei with no non-trivial, proper subkei. Then $n = p$ is prime and $K \cong D_p$.*

Proof. Let $x, y \in [n]$ be distinct and consider the subgraph G_{xy}. As each vertex is adjacent to at most one edge of each colour, each vertex of G_{xy} has edge degree at most 2, and thus G_{xy} consists of a disjoint union of paths and (possibly degenerate) cycles. Let P_x be the component of G_{xy} containing x and P_y be the component containing y; as $(x)f_x = x$, x is not adjacent to an edge of colour x and thus has edge degree at most 1. Hence P_x is a path and similarly P_y is a path. Now we can apply Lemma A.2.1 with $T = \{x, y\}$; as $n \geqslant 3$ we have that $(V(P_x) \cup V(P_y), \triangleright)$

is a non-trivial subkei of K, and thus from our assumptions $V(P_x) \cup V(P_y) = [n]$. Without loss of generality P_x has at least two vertices; applying Lemma A.2.1 again gives that $(V(P_x), \triangleright)$ is a non-trivial subkei of K; from our assumptions, we must have $V(P_x) = [n]$. Then $P_x = G_{xy}$ is connected and thus must consist of a single xy path.

Now consider D_n; denote the binary operation on $[n]$ by $\triangleright'$ and the associated maps by $g_1, \ldots, g_n$. From Proposition 6.2.9, the subgraph G_{12} of G_{D_n} is a spanning path $1 = z_1, \ldots, z_n = 2$ with edges alternating in colour between 2 and 1 (the first edge is of colour 2). Relabel the vertices along P_x as $x = w_1, \ldots, w_n = y$, and define a bijection $\phi \in \mathfrak{S}_n$ by $w_r \mapsto z_r$ for each r, so in particular $x \mapsto 1$ and $y \mapsto 2$. Now note that

$$f_x = (w_2\ w_3) \cdots (w_{n-1}\ w_n)$$

while

$$g_1 = (z_2\ z_3) \cdots (z_{n-1}\ z_n) = ((w_2)\phi\ (w_3)\phi) \cdots ((w_{n-1})\phi\ (w_n)\phi)$$

and similarly

$$f_y = (w_1\ w_2) \cdots (w_{n-2}\ w_{n-1})$$

while

$$g_2 = (z_1\ z_2) \cdots (z_{n-2}\ z_{n-1}) = ((w_1)\phi\ (w_2)\phi) \cdots ((w_{n-2})\phi\ (w_{n-1})\phi),$$

so $g_1 = \phi^{-1} f_x \phi$ and $g_2 = \phi^{-1} f_y \phi$. Now for $0 \leqslant r \leqslant (n-1)/2$, we have that $w_{2r+1} = (x)(f_y f_x)^r$ and $z_{2r+1} = (1)(g_2 g_1)^r$, so $f_{w_{2r+1}} = (f_x f_y)^r f_x (f_y f_x)^r$ and $g_{z_{2r+1}} = (g_1 g_2)^r g_1 (g_2 g_1)^r$. Hence, noting the cancellations,

$$g_{z_{2r+1}} = \phi^{-1} (f_x f_y)^r f_x (f_y f_x)^r \phi = \phi^{-1} \circ f_{w_{2r+1}} \circ \phi.$$

Similarly, for $0 \leqslant r \leqslant (n-3)/2$, $w_{2r+2} = (w_{2r+1}) f_z$ and $z_{2r+2} = (z_{2r+1}) g_2$, so pre- and post-multiplying all of the above by f_y or g_2 gives the result for even-indexed vertices as well. Hence $g_{(v)\phi} = \phi^{-1} f_v \phi$ for all $v \in [n]$.

We now show that ϕ induces an isomorphism from K to D_n. Let $u, v \in [n]$ be arbitrary; from the previous paragraph we have $(u) f_v \phi = (u) \phi g_{(v)\phi}$, or equivalently $(u \triangleright v)\phi = ((u)\phi) \triangleright' ((v)\phi)$.

Finally, suppose that n is composite; then $n = pm$ where $p < n$ is an odd prime. Let $T = \{im \mid 1 \leqslant i \leqslant p\}$, so $|T| = p$. Let $i, j \in [p]$ be arbitrary, so $im, jm \in T$; then $im \triangleright' jm \equiv 2jm - im \pmod{n}$, and as m divides n, we see that $im \triangleright' jm \in T$. Hence $(T, \triangleright')$ is a subkei of D_n of order p, seen to be isomorphic to D_p under the map $im \mapsto i$. But this is a contradiction as $K \cong D_n$ and K has no subkei of order p; this completes the proof. $\square$

Corollary A.2.3. *Let $n \geqslant 3$ and let $K = ([n], \triangleright)$ be a kei. Then there exists a subkei L of K such that L is either isomorphic to D_p for some odd prime p, or to one of the three exceptional subkei shown in Figure A.2.*

Proof. Put $K_0 = K$ and consider the following procedure (starting from $i = 0$):

1. If $K_i = (Y_i, \triangleright)$ has a subkei of size k for some $3 \leqslant k < |Y_i|$, choose such a subkei $K_{i+1} := (Y_{i+1}, \triangleright)$, update i to $i + 1$ and repeat;

2. Otherwise, set $t = i$ and stop.

Note that this process must eventually terminate as $(|Y_i|)_{i \geqslant 0}$ is a strictly decreasing sequence; further, as each K_{i+1} is a subkei of K_i, we have that $K_t = (Y_t, \triangleright)$ is a subkei of K.

By our construction, K_t has no proper subkei of size at least three; if K_t has no proper, non-trivial subkei at all then by Proposition A.2.2 $K_t \cong D_p$ for some odd prime p and we can put $L = K_t$. Hence we may suppose that K_t has a proper subkei of size two but no strictly larger proper subkei.

Suppose that K_t is generated by two distinct elements $x, y \in Y_t$; then the graph $G_{xy}[Y_t]$ consists of a single path spanning all the vertices. From the proof of Proposition A.2.2 we have that $K_t \cong D_{|Y_t|}$, and that K_t contains a subkei isomorphic to D_p, where p is some odd prime.[1] In this case, we can set L to be this subkei.

So we may further suppose that for any pair of distinct elements $x, y \in Y_t$, the subkei generated by x and y has size two and is thus isomorphic to the trivial kei T_2; hence x and y commute, and as these elements were arbitrary, the maps $(f_x|_{Y_t})_{x \in Y_t}$ mutually commute. Then from Proposition 5.1.4, each component of G_{K_t} has size a power of two, and is isomorphic to a multicube. If K_t is connected then in our earlier notation $G_{K_t} \cong Q_{\mathcal{F}}^{\log_2 |Y_t|}$, where Y_t is identified with the set $\{0,1\}^m$ (for some m) and the edges of the ith colour (in some arbitrary ordering) are identified with the set $\{vv_{F_i} \mid v \in \{0,1\}^m\}$. But this means the transposition number of K_t is either 0 or $|Y_t|/2$; in the first case, this implies that $K_t \cong T_{|Y_t|}$ and thus contains subkei of any order s, $1 \leqslant s \leqslant |Y_t|$, a contradiction. In the second case, we have that for any $i \in Y_t$ the map $f_i|_{Y_t}$ has no fixed points, contradicting the fact that $(i)f_i = i$ for all i.

Hence K_t is not connected. Now the vertices in any disjoint union of components of G_{K_t} form a subkei; as K_t has no proper subkei of size greater than two, each component is either an isolated vertex or a (multi)edge. Then we see that if G_{K_t}

[1] This statement contradicts the fact that K_t has a subkei of order two, but we don't need to use this information.

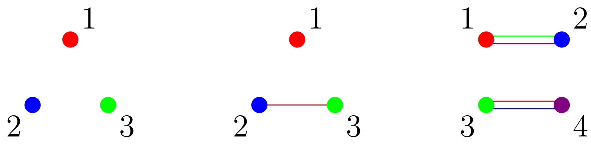

Figure A.2: Exceptional subkei.

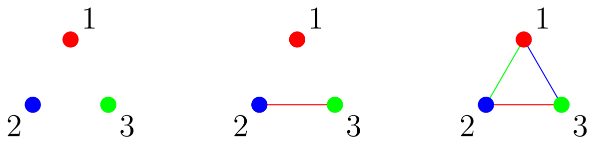

Figure A.3: Kei on $\{1, 2, 3\}$.

contains three isolated vertices, an edge and an isolated vertex, or two edges, these components must form the whole of G_{K_t}; these correspond to the three exceptional subkei and thus are the other choices for L. $\square$

A.3 Very small kei

In this last section, we show how to use some of the approaches covered to classify all kei of size at most five, up to isomorphism. These results were obtained computationally by Ho and Nelson [14]; later work by Henderson, Macedo and Nelson [13] extended these results to all quandles of size at most eight. We will restrict ourselves to the analytic tools developed in the thesis and this appendix.

A.3.1 Kei of size at most 4

Let $n \leqslant 4$ and consider a kei $([n], \triangleright)$. For any $x \in [n]$, f_x fixes x; as there are at most three other elements and f_x is an involution, f_x is the identity or a single transposition. Hence we can classify all of these kei by appealing to Proposition A.1.1. For $n = 1$ or 2 any kei is trivial; the kei on 3 or 4 elements are given (up to isomorphism) in Figures A.3 and A.4. They are obtained by considering different partitions of $[n]$ into sets of size 1, 2 and 3.

A.3.2 Kei of size 5

If $n = 5$, there is a possibility that some $x \in [n]$ acts as a double transposition. If this is not the case, then we can use Proposition A.1.1 again; see Figure A.5.

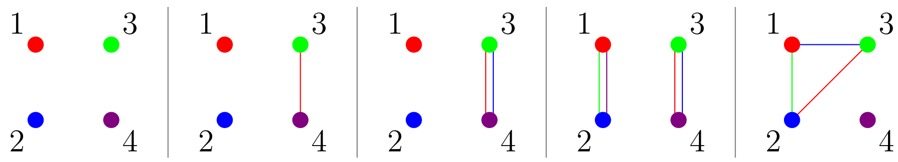

Figure A.4: Kei on $\{1, 2, 3, 4\}$*.*

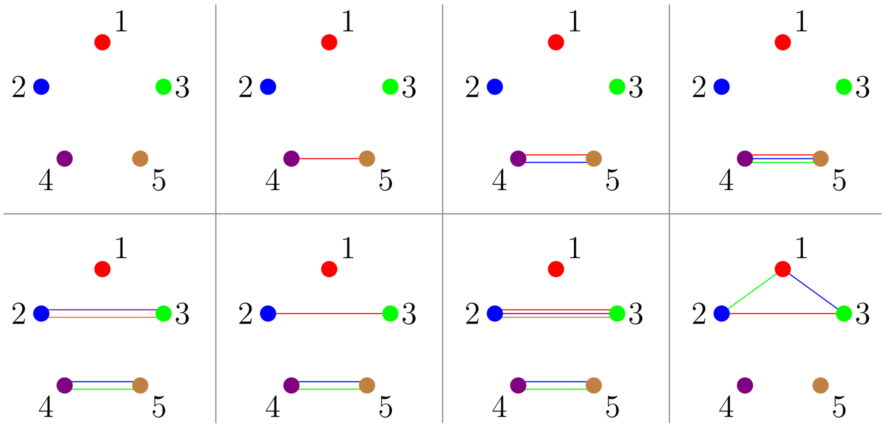

Figure A.5: Kei on $\{1, 2, 3, 4, 5\}$ *where each map is the identity or a single transposition.*

Otherwise, we may suppose without loss of generality that f_1 is a double transposition. First suppose that for each $j \geqslant 2$, f_j acts entirely within $\{2, 3, 4, 5\}$; then $(1)f_j = 1$ for each j and so f_1 commutes with all of f_2, f_3, f_4 and f_5. In particular, setting (again without loss of generality) $f_1 = (2\ 3)(4\ 5)$, we see that $f_3 = f_{(2)f_1} = f_1 f_2 f_1 = f_2$ and similarly $f_5 = f_4$. Then we can use the example at the end of Section 6.1.2 (with $\{2, 3, 4, 5\}$ instead of $\{1, 2, 3, 4\}$); the resulting kei are given in Figure A.6.

So finally, suppose there is some map not fixing 1; we can say without loss of generality that $(1)f_j = 2$ for some j. Then $f_2 = f_j f_1 f_j$, so f_2 is conjugate to f_1 and is thus also a double transposition. Arguing as in Proposition 6.2.9, the graph G_{12}

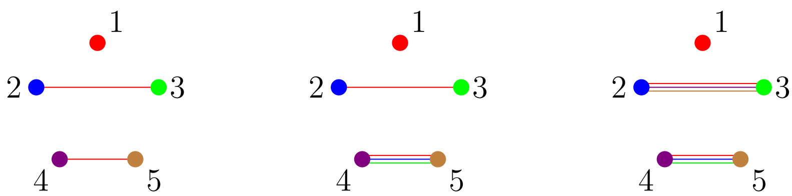

Figure A.6: More kei on $\{1, 2, 3, 4, 5\}$*; each of* f_2*,* f_3*,* f_4 *and* f_5 *was assumed to fix 1.*

Figure A.7: The last two kei on $\{1, 2, 3, 4, 5\}$.

consists of a 1-2 path and a collection of disjoint even cycles, with edges alternating in colour along the path and around any cycle. If the path is of length 4 (so it includes all the vertices), we may take the path to be (1,3,4,5,2) with the first edge of colour 2; then $f_1 = (2\ 5)(3\ 4)$ and $f_2 = (1\ 3)(4\ 5)$. We can then compute $f_3 = f_{(1)f_2} = f_2 f_1 f_2 = (1\ 5)(2\ 4)$, $f_4 = f_{(3)f_1} = f_1 f_3 f_1 = (1\ 2)(3\ 5)$ and $f_5 = f_{(4)f_2} = f_2 f_4 f_2 = (1\ 4)(2\ 3)$, and this kei is D_5.

Otherwise, the 1-2 path is of length 2, which we may take to be (1,3,2). Then 4 and 5 must form a degenerate 'cycle' of length 2 in G_{12}, so we have $f_1 = (2\ 3)(4\ 5)$, $f_2 = (1\ 3)(4\ 5)$ and $f_3 = f_{(1)f_2} = f_2 f_1 f_2 = (1\ 2)(4\ 5)$. To compute f_4, we observe that $5 = (4)f_2 = (4)f_3$, so $f_5 = f_2 f_4 f_2 = f_3 f_4 f_3$. Then $(f_3 f_2) f_4 = f_4 (f_3 f_2)$, so that f_4 commutes with $f_3 f_2 = (1\ 2\ 3)$. It is straightforward to show that the centraliser of this 3-cycle is the group $\langle (1\ 2\ 3), (4\ 5) \rangle$, and the only involution in this group fixing 4 is the identity ι. Then $f_5 = f_2 f_4 f_2 = \iota$, and we have determined the kei; see Figure A.7.

Bibliography

[1] Noga Alon and Joel H Spencer. *The Probabilistic Method*. Wiley. com, 2004.

[2] Simon R Blackburn. Enumerating finite racks, quandles and kei. *The Electronic Journal of Combinatorics*, 20(3):P43, 2013.

[3] Richard Brauer and KA Fowler. On groups of even order. *The Annals of Mathematics*, 62(3):565–583, 1955.

[4] Charles Curtis. *Linear Algebra: An Introductory Approach*. Springer Science & Business Media, 2012.

[5] Mohamed Elhamdadi, Jennifer MacQuarrie, and Ricardo Restrepo. Automorphism groups of quandles. *Journal of Algebra and Its Applications*, 11(01):1250008, 2012.

[6] Paul Erdős and Alfréd Rényi. On random graphs. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.

[7] Pavel Etingof, Alexander Soloviev, and Robert Guralnick. Indecomposable set-theoretical solutions to the quantum Yang–Baxter equation on a set with a prime number of elements. *Journal of Algebra*, 242(2):709–719, 2001.

[8] Roger Fenn and Colin Rourke. Racks and links in codimension two. *Journal of Knot Theory and Its Ramifications*, 1(04):343–406, 1992.

[9] Michael Giudici and Chris Parker. There is no upper bound for the diameter of the commuting graph of a finite group. *Journal of Combinatorial Theory, Series A*, 120(7):1600–1603, 2013.

[10] Michael Giudici and Aedan Pope. On bounding the diameter of the commuting graph of a group. *Journal of Group Theory*, 17(1):131–149, 2014.

[11] Matıas Grana. Indecomposable racks of order p^2. *Contributions to Algebra and Geometry*, 45(2):665–676, 2004.

[12] Peter Hegarty and Dmitry Zhelezov. Can connected commuting graphs of finite groups have arbitrarily large diameter? In *Geometry, Structure and Randomness in Combinatorics*, pages 141–144. Springer, 2014.

[13] Richard Henderson, Todd Macedo, and Sam Nelson. Symbolic computation with finite quandles. *Journal of Symbolic Computation*, 41(7):811–817, 2006.

[14] Benita Ho and Sam Nelson. Matrices and finite quandles. *Homology, Homotopy and Applications*, 7(1):197–208, 2005.

[15] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

[16] Alan J Hoffman and Robert R Singleton. On Moore graphs with diameters 2 and 3. *IBM Journal of Research and Development*, 4(5):497–504, 1960.

[17] A Iranmanesh and A Jafarzadeh. On the commuting graph associated with the symmetric and alternating groups. *Journal of Algebra and Its Applications*, 7(01):129–146, 2008.

[18] Svante Janson. Poisson approximation for large deviations. *Random Structures & Algorithms*, 1(2):221–229, 1990.

[19] Svante Janson. New versions of Suen's correlation inequality. *Random Structures and Algorithms*, 13(3-4):467–483, 1998.

[20] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random Graphs*. John Wiley & Sons, 2011.

[21] David Joyce. A classifying invariant of knots, the knot quandle. *Journal of Pure and Applied Algebra*, 23(1):37–65, 1982.

[22] Victor Klee and David Larman. Diameters of random graphs. *Canad. J. Math*, 33:618–640, 1981.

[23] Attila Maróti. Bounding the number of conjugacy classes of a permutation group. *Journal of Group Theory*, 8(3):273–289, 2005.

[24] Sergei Vladimirovich Matveev. Distributive groupoids in knot theory. *Sbornik: Mathematics*, 47(1):73–83, 1984.

[25] James McCarron. Connected quandles with order equal to twice an odd prime. *arXiv preprint arXiv:1210.2150*, 2012.

[26] Sam Nelson. The combinatorial revolution in knot theory. *Notices of the AMS*, 58(11):1553–1561, 2011.

[27] BH Neumann. A problem of Paul Erdős on groups. *J. Austral. Math. Soc. Ser. A*, 21(4):467–472, 1976.

[28] L Pyber. Enumerating finite groups of given order. *Annals of Mathematics*, pages 203–220, 1993.

[29] Yoav Segev and Gary M Seitz. Anisotropic groups of type A_n and the commuting graph of finite simple groups. *Pacific Journal of Mathematics*, 202(1):125–225, 2002.

[30] David Stanovskỳ. The origins of involutory quandles. *arXiv preprint arXiv: 1506.02389*, 2015.

[31] WC Suen. A correlation inequality and a Poisson limit theorem for nonoverlapping balanced subgraphs of a random graph. *Random Structures & Algorithms*, 1(2):231–242, 1990.

[32] Mituhisa Takasaki. Abstractions of symmetric functions. *Tohuko Mathematical Journal*, 49:143–207, 1943.