Combinatorial methods in group theory and group-theoretic methods in combinatorics

G2T2 Summer School, Hong Kong,

22-25 July 2026

Marston Conder

University of Auckland and University of Primorska

m.conder@auckland.ac.nz

Outline of topics

1. Basic applications of counting
2. Methods for generating random elements of a group
3. Cayley graphs
4. Schreier coset graphs and their applications
5. Back-track search to find small index subgroups
6. Double-coset graphs and some applications
7. Möbius inversion on lattices and applications

Copies of slides can be made available by email or USB stick.

But before I start ... **where am I from?**

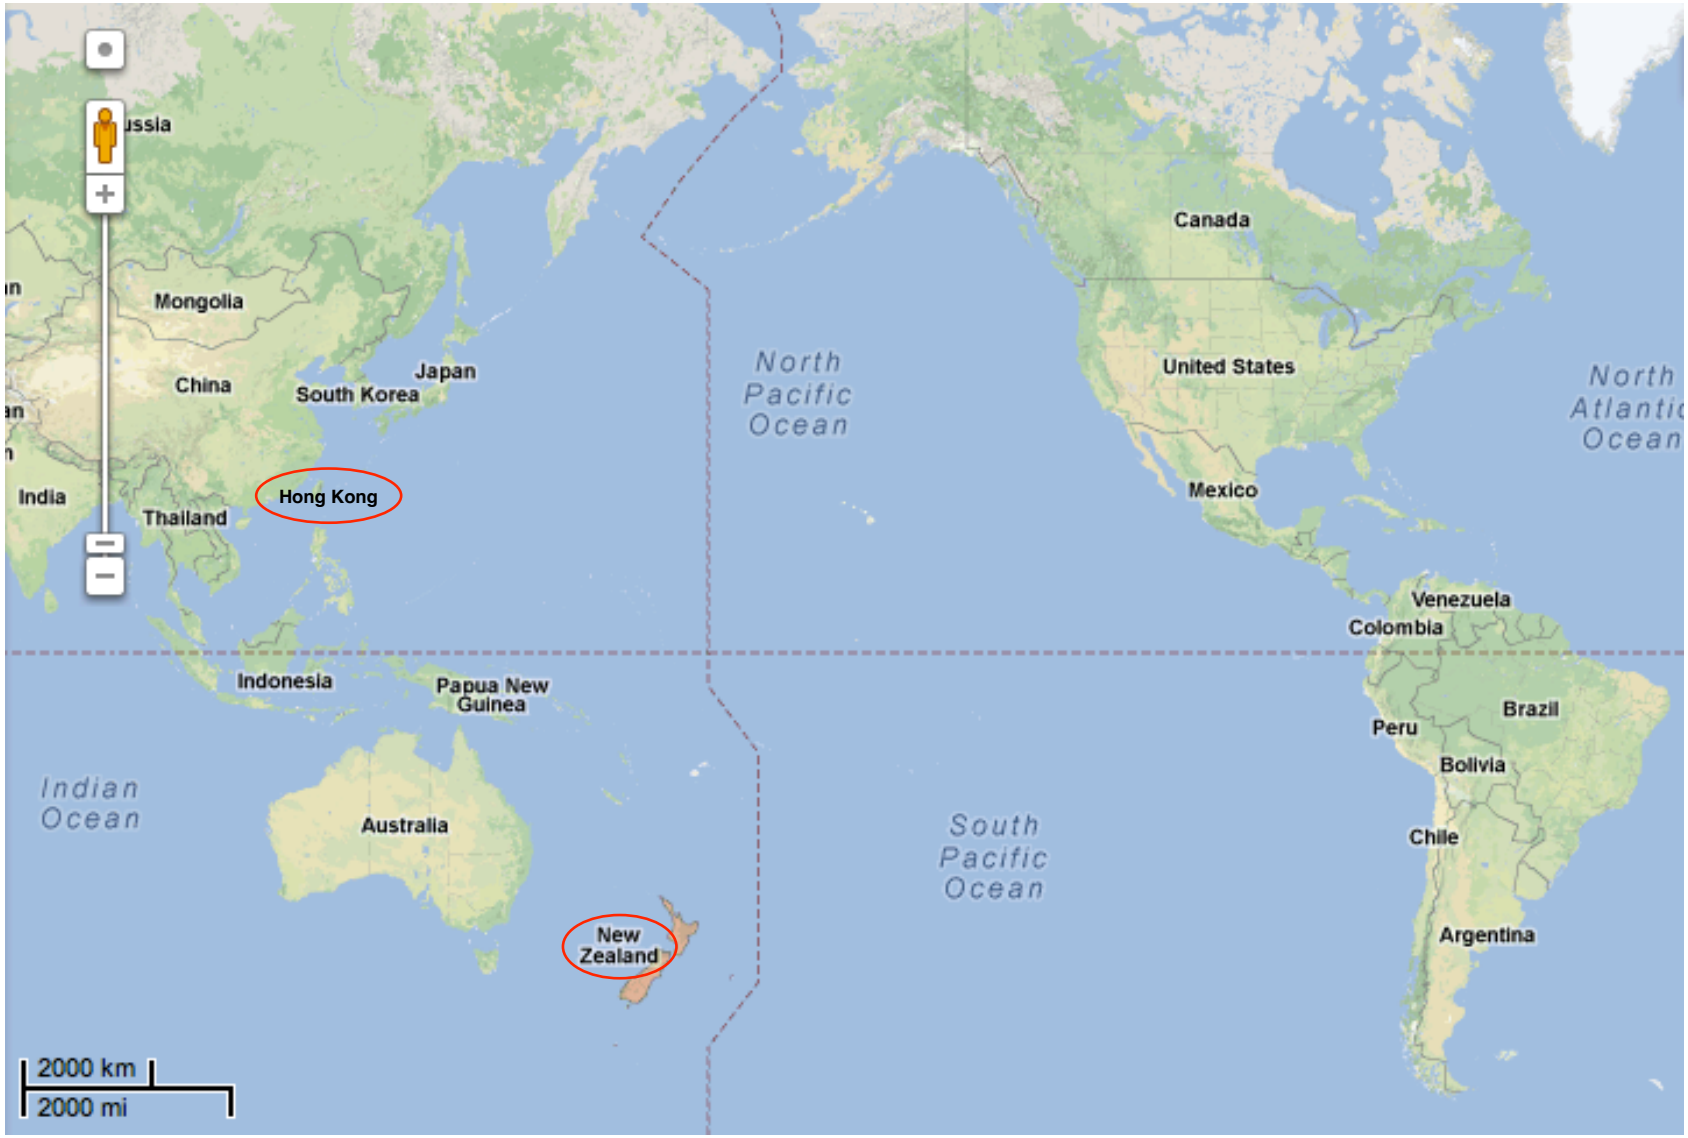
Canada
Mongolia
China
Japan
South Korea
North Pacific Ocean
United States
North Atlantic Ocean
India
Hong Kong
Thailand
Mexico
Venezuela
Colombia
Indonesia
Papua New Guinea
Brazil
Peru
Bolivia
Indian Ocean
Australia
South Pacific Ocean
Chile
New Zealand
Argentina
2000 km
2000 mi

§0. Background on group theory

In this course, a relatively small amount of knowledge of group theory will be needed, namely the following:

- Definition and elementary properties of groups
- Some well known examples of groups – e.g. C_n (cyclic), D_n (dihedral), A_n (alternating), S_n (symmetric)
- Subgroups, cosets, conjugacy, normal subgroups, factor groups, homomorphisms, isomorphisms, automorphisms
- Permutation groups (i.e. subgroups of symmetric groups)
- Generating sets for groups
- Presentation of groups by generators and relations.

Information about most of these things is available on Wikipedia.

§1. Basic applications of counting

Many theorems in combinatorics can be proved by counting the same thing in two different ways – e.g. the 'hand-shaking lemma' in graph theory ($2|E(X)| = \sum_{v \in V(X)} \deg(v)$).

The same thing happens in aspects of group theory.

Theorem (The 'Orbit-Stabiliser Theorem')
If G is a permutation group on a set Ω, and G_α and α^G are the stabiliser $\{g \in G \mid \alpha^g = \alpha\}$ and orbit $\{\alpha^g : g \in G\}$ of a point $\alpha \in \Omega$, then $|G| = |\alpha^G|\,|G_\alpha|$ for every $\alpha \in \Omega$.

Proof. Count the number of pairs $(g, \beta) \in G \times \Omega$ such that $\alpha^g = \beta$ in two different ways. On one hand, choosing g first gives the number as $|G|$, while on the other hand, choosing β first gives the number as $|\alpha^G|\,|G_\alpha|$, because if $\beta = \alpha^h$ for some h, then $\alpha^g = \beta$ iff $g \in G_\alpha h$, and $|G_\alpha h| = |G_\alpha|$. ■

An application: Lagrange's Theorem

If H is a subgroup of the finite group G, then $|G| = |G:H|\,|H|$ – and in particular, both $|H|$ and $|G:H|$ are divisors of $|G|$.

Proof. Let G act on the right coset space $\{Hx : x \in G\}$ by right multiplication, with each element $g \in G$ inducing the permutation $\mu_g : Hx \mapsto Hxg$. Then the orbit of H is the entire coset space $(G:H)$, while the stabiliser of H is $\{g \in G \mid Hg = H\} = \{g \in G \mid g \in H\} = H$, so the Orbit-Stabiliser Theorem gives $|G| = |G:H|\,|H|$.

Exercise: Show that the 'converse' of Lagrange's Theorem is not true, by finding a finite group G that has no subgroup of order d for some divisor d of $|G|$.

Some more applications (for finite groups)

Conjugacy: Let the group G act on itself by conjugation, with each $g \in G$ inducing the permutation $\tau_g : x \mapsto g^{-1}xg$.

The orbit of x is the conjugacy class $[x] = \{g^{-1}xg : g \in G\}$, and its stabiliser is the centraliser $C_G(x) = \{g \in G \mid xg = gx\}$, so the Orbit-Stabiliser Theorem gives $|G| = |[x]|\,|C_G(x)|$, or equivalently, $|[x]| = |G|/|C_g(x)| = |G : C_G(x)|$, for all $x \in G$.

Class equation: If $x_1, x_2, \ldots, x_k$ are representatives of the k distinct conjugacy classes of elements of the group G, then

$$|G| = \sum_{1 \le i \le k} |[x_i]| = \sum_{1 \le i \le k} |G : C_G(x_i)|.$$

Exercise: Use this to show that if G has order p^s for some prime p, then the centre $Z(G)$ of G has at least p elements.

Burnside's Lemma: If the finite group G acts on the set Ω, with exactly m orbits, and $F_\Omega(g) = \{\alpha \in \Omega \mid \alpha^g = \alpha\}$ is the set of fixed points of each $g \in G$, then $m = \frac{1}{|G|} \sum_{g \in G} |F_\Omega(g)|$.

Proof. Simply count pairs $(g, \alpha) \in G \times \Omega$ such that $\alpha^g = \alpha$.

On one hand, counting by g, this number is $\sum_{g \in G} |F_\Omega(g)|$.

On the other hand, for any $\alpha \in \Omega$, the number of g for which $\alpha^g = \alpha$ is $|G_\alpha| = |G|/|\Delta|$, where $\Delta = \alpha^G$ is the orbit of α. When this is counted over all $\alpha \in \Omega$, the term $|G|/|\Delta|$ is counted $|\Delta|$ times (once for each point in Δ), and so each orbit Δ contributes $|\Delta|\,(|G|/|\Delta|) = |G|$ to the total. Hence the total number is $m|G|$, giving $m|G| = \sum_{g \in G} |F_\Omega(g)|$. ■

An application of Burnside's Lemma:

In how many inequivalent ways can the faces of a regular tetrahedron be coloured using up to k given colours with one colour/face (but allowing more than one face/colour)?

Two colourings are considered to be equivalent if one can be obtained from the other by a rotation of the tetrahedron. The rotation group is A_4, acting naturally on the 4 vertices (or 4 faces), so we need the number of orbits on colourings.

The identity fixes all k^4 colourings; a double-transposition $(a,b)(c,d)$ fixes k^2 colourings (with a and b coloured the same and c and d coloured the same); and a 3-cycle (a,b,c) fixes k^2 colourings (with a, b and c all coloured the same).

By Burnside's Lemma, the total number of inequivalent colourings is $\frac{1}{|A_4|}(k^4 + 3k^2 + 8k^2) = \frac{1}{12}k^2(k^2 + 11)$.

A combinatorial proof of some Sylow theory

If G is a finite group whose order $|G|$ is divisible by the prime p, and p^s is the largest power of p that divides $|G|$, then any subgroup of G of order p^s is called a Sylow p-subgroup of G.

Two of the main statements of Sylow theory are that every such G has at least one Sylow p-subgroup and that if n_p is the number of Sylow p-subgroups of G, then $n_p \equiv 1 \mod p$.

Clearly the first property is a consequence of the second one, so we will prove the second one, using a combinatorial proof due to Helmut Wielandt.

Before doing that, we note that if G is cyclic of order n, generated by x, say, then G has exactly one subgroup of order p^s – namely the (cyclic) subgroup generated by x^{n/p^s}.

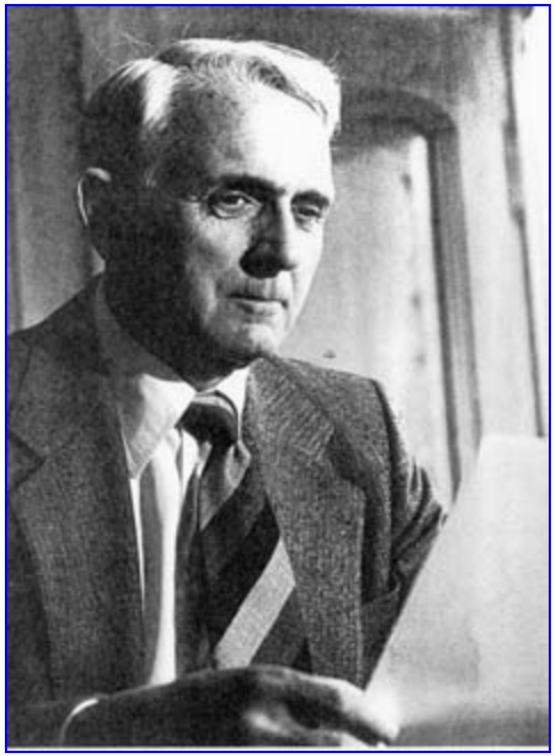

Helmut Wielandt (1910–2001)

Proof (that $n_p \equiv 1 \bmod p$):

Define Ω as the set of $\binom{|G|}{p^s}$ **subsets** of G of size p^s and then let G act on the set Ω by right multiplication, with each $g \in G$ taking S to $Sg = \{xg : x \in S\}$ for every $S \in \Omega$.

Note that if Δ is an orbit of G on Ω, then there exists at least one $S \in \Delta$ such that $1 \in S$ (because if $x \in T$ where $T \in \Delta$, then $1 = xx^{-1} \in Tx^{-1}$ with $Tx^{-1} \in \Delta$).

Now consider the stabiliser G_S of S in G. If $g \in G_S$ then $g = 1g \in Sg = S$ so $g \in S$, hence $G_S \subseteq S$, giving $|G_S| \leq |S|$.

We split the orbits of G on Ω into two types, according to whether or not $G_S = S$. [PTO]

Type (a) Suppose that $G_S = S$. Then S is a subgroup of G, and $\Delta = \{Sg : g \in G\}$ is the right coset space $(G : S)$. In particular, Δ contains only one subgroup (namely S itself), and also we have $|\Delta| = |G : S| = |G|/|S| = (p^s q)/(p^s) = q$. Conversely, if S is a Sylow p-subgroup of G, then Δ has type (a).

Type (b) Suppose that $G_S \subset S$. Then $|G_S| < |S| = p^s$ but also $|G_S| = |G|/|\Delta|$ divides $|G|$, so $|\Delta|$ is divisible by p.

These imply that G has n_p orbits of length q on Ω, with the lengths of all other orbits being divisible by p.

Thus $\binom{|G|}{p^s} = |\Omega| \equiv n_p\, q \mod p$. But as this also holds for the cyclic group of the same order $|G|$, for which $n_p = 1$, we find $n_p\, q \equiv \binom{|G|}{p^s} \equiv q \mod p$ and so $n_p \equiv 1 \mod p$. ■

Further Sylow theory

Let G be a finite group with order $p^s q$ (where p is prime) as before. Then the following hold:

- If P is a Sylow p-subgroup of G, and Q is any subgroup of G with order p^r for some r, then $Q \subseteq x^{-1}Px$ for some $x \in G$.
- If P and Q are Sylow p-subgroups of G, then $Q = x^{-1}Px$ for some $x \in G$. Hence under conjugation, G has a single orbit on Sylow p-subgroups.
- The number of Sylow p-subgroups of G divides q.

My favourite application: If $|G| = pqr$ where p, q and r are distinct primes, then G has a normal subgroup of order p, q or r. Hence in particular, no such group can be simple.

Proof. Assume the contrary. Let n_p, n_q and n_r be the numbers of Sylow p-, q- and r-subgroups (of orders p, q and r), respectively. If $n_p = 1$ then G has a unique Sylow p-subgroup P, and then $x^{-1}Px = P \ \forall x \in G$, so $P \lhd G$, contradiction. Thus $n_p > 1$ and similarly $n_q > 1$ and $n_r > 1$.

Next, without loss of generality, we may suppose $p < q < r$.

By Sylow theory, n_r divides $|G|/r = pq$, so $n_r = p, q$ or pq. But also $n_r \equiv 1$ mod r and hence $n_r > r > q > p$, so $n_r = pq$. Then because any two Sylow r-subgroups intersect trivially (by Lagrange's theorem and since r is prime), it follows that G has $pq(r-1)$ elements of order r.

Similarly, $n_q = r$ or pr, while also $n_p = q, r$ or qr, and hence G has at least $r(q-1)$ elements of order q, and at least $q(p-1)$ elements of order p.

It follows that the number of elements of prime order in G is at least $q(p-1)+r(q-1)+pq(r-1) = pqr-q+qr-r = pqr+(q-1)(r-1)-1$. But this is greater than $pqr = |G|$, contradiction! Hence proof. ■

Exercises

- If $|G| = pq$ where p and q are distinct primes, then G cannot be simple.
- If $|G| = p^2q$ where p and q are distinct primes, then G cannot be simple.

§2. Generating random elements of a group

Q: How do we select elements randomly from a group?

But first, why would we want to? One reason is that many group-theoretic algorithms rely on being able to do this. (Some examples: algorithms for working out the order of a group generated by given elements, or identifying the group itself.)

For some groups, finding random elements is easy:

- Cyclic groups $C_n = \langle\, x \mid x^n = 1 \,\rangle$

Just take x^k where k is a random element of $\{0, 1, 2, \ldots, n-1\}$.

- Symmetric groups S_n

Just take a random permutation of $\{1, 2, 3, \ldots, n\}$.

Exercises (or just possibilities to consider):

What about the following groups?

- Dihedral groups $D_n = \langle\, x, y \mid x^2 = y^n = 1,\ xyx = y^{-1} \,\rangle$

How do we find a random element of this group?

- Alternating groups A_n

How do we find a random **even** permutation of $\{1, 2, 3, \ldots, n\}$?

- A sharply 3-transitive permutation group G – such as PSL$(2, 2^s)$ in its action on the projective line over GF(2^s)? [Base & SGS]

- A group of order p^7 where p is prime? [PC-presentations]

- The Monster simple group? [Worth thinking about]

Clearly finding a good method may depend on the type of description we have for the group.

Digression: Generating sets

Let G be a group, and let X be a subset of G.

We say that X is a generating set for G if every element of G can be expressed as a (finite-length) 'word' $x_1^{e_1}x_2^{e_2}\ldots x_k^{e_k}$ in elements of X and their inverses (with $x_i \in X$ and $e_i = \pm 1$ for $1 \leq i \leq k$).

Also if X is finite, then G is said to be finitely-generated, and the rank of G is defined as the smallest possible value of $|X|$ (over all such X). Otherwise G has infinite rank.

More generally, if S is the set of all such words on $X^{\pm}$, then S is a subgroup of G, called the subgroup generated by X and denoted by $\langle X \rangle$.

Examples:

- Cyclic groups $\equiv$ groups of rank 1
- Dihedral groups have rank 2 (generated by 2 reflections, or by a reflection and a rotation)
- The abelian group $C_m \times C_m \times C_m$ has rank 3 (for all $m > 0$)
- The symmetric group S_n has rank 2 (e.g. generated by the transposition $(1, 2)$ and the n-cycle $(1, 2, \ldots, n)$), and is also generated by $\{(1, 2), (2, 3), \ldots, (n-1, n)\}$
- The alternating group A_n is the subgroup of S_n generated by the 3-cycles (a, b, c) ... and also has rank 2
- Every non-abelian simple group has rank 2 [Needs CFSG]

One further point (for later use):

Let G be a group of order n, with elements $g_1, g_2, \ldots, g_n$.

The multiplication table of a finite group G of order n is an $n \times n$ array with (i,j)th entry equal to the product $g_i\, g_j$. This is a Latin square. But much of its content is redundant!

If $X = \{x_1, x_2, \ldots, x_m\}$ is a generating set for G, of size m, then we can reduce the multiplication table to an $m \times n$ array with (i,j)th entry equal to the product $x_i\, g_j$.

We can call this a reduced Cayley table for the pair (G, X).

Note that any element of the full multiplication table, say $g_i\, g_j$, can be obtained by expressing g_i as a word on X, say $x_{i_1}^{e_1} x_{i_2}^{e_2} \ldots x_{i_k}^{e_k}$, and then working out $g_i\, g_j = x_{i_1}^{e_1} x_{i_2}^{e_2} \ldots x_{i_k}^{e_k}\, g_j$ by successive calls to the reduced Cayley table.

Now, back to generating random elements of a group ...

How do we select elements randomly from a finite group G when G does not have a nice canonical form that makes this easy?

One effective method was developed by Celler, Leedham-Green, Murray, Niemeyer and O'Brien (in 1995), and is now called the Product Replacement Algorithm.

The algorithm starts by taking an ordered generating set $X = \{x_1, x_2, \ldots, x_m\}$ for G of size $m > \text{rank}(G)$, and then performs the following basic operation a number of times:

> Choose two random integers i and j from $\{1, 2 \ldots, m\}$,
> and replace x_i by either $x_i x_j$ or $x_j x_i$, to give a new X.

If this is done sufficiently many times, then any element of the resulting set X may be taken as a random element of G.

Example (for a group G of rank less than 4):

$$
\begin{array}{rcllll}
X_1 & = & \{\ x_1, & x_2, & x_3, & x_4\ \}, \\
X_2 & = & \{\ x_1, & x_4x_2, & x_3, & x_4\ \}, \\
X_3 & = & \{\ x_1x_3, & x_4x_2, & x_3, & x_4\ \}, \\
X_4 & = & \{\ x_1x_3x_4, & x_4x_2, & x_3, & x_4\ \}, \\
X_5 & = & \{\ x_1x_3x_4, & x_4x_2, & x_3x_4x_2, & x_4\ \}, \\
X_6 & = & \{\ x_1x_3x_4, & x_4x_4x_2, & x_3x_4x_2, & x_4\ \}, \\
X_7 & = & \{\ x_1x_3x_4, & x_4x_4x_2, & x_3x_4x_2, & x_1x_3x_4x_4\ \}, \\
 & \vdots & & & &
\end{array}
$$

Easy! And very quick!:

Features of the Product Replacement Algorithm:

- Easy to understand
- Very easy to implement
- Works for any finite group
- Very fast! Complexity $O(Nc)$ where $N =$ number of steps and $c =$ cost of a single multiplication in G.
- Elements found are well-distributed according to various criteria and statistical tests.

Let k be the maximal cardinality of a minimal generating set X for G, and suppose $m \geq 2k$. Also let $\mathcal{X}$ be the set of ordered m-tuples of elements of G that generate G, and let X_t be the element of $\mathcal{X}$ obtained by repeating the basic operation t times. Then for every $Y \in \mathcal{X}$, the probability that $X_t = Y$ tends to $1/|\mathcal{X}|$ as $t \to \infty$.

§3. Cayley graphs

A Cayley graph $\mathrm{Cay}(G, X)$ is a graph with vertex-set a group G and edge-set $\{\{g, xg\} : g \in G,\, x \in X\}$ for some $X \subseteq G$.

This gives a diagrammatic representation of multiplication of elements of G by elements of X (on the left), and hence of the rows of the multiplication table for G corresponding to the elements of X.

Usually (but not necessarily) we assume that $\Gamma = \mathrm{Cay}(G, X)$ is finite, undirected, simple and connected – and hence that

- G is finite,
- X does not contain the identity element of G, and
- X generates G.

Also we may assume that X is closed under taking inverses.

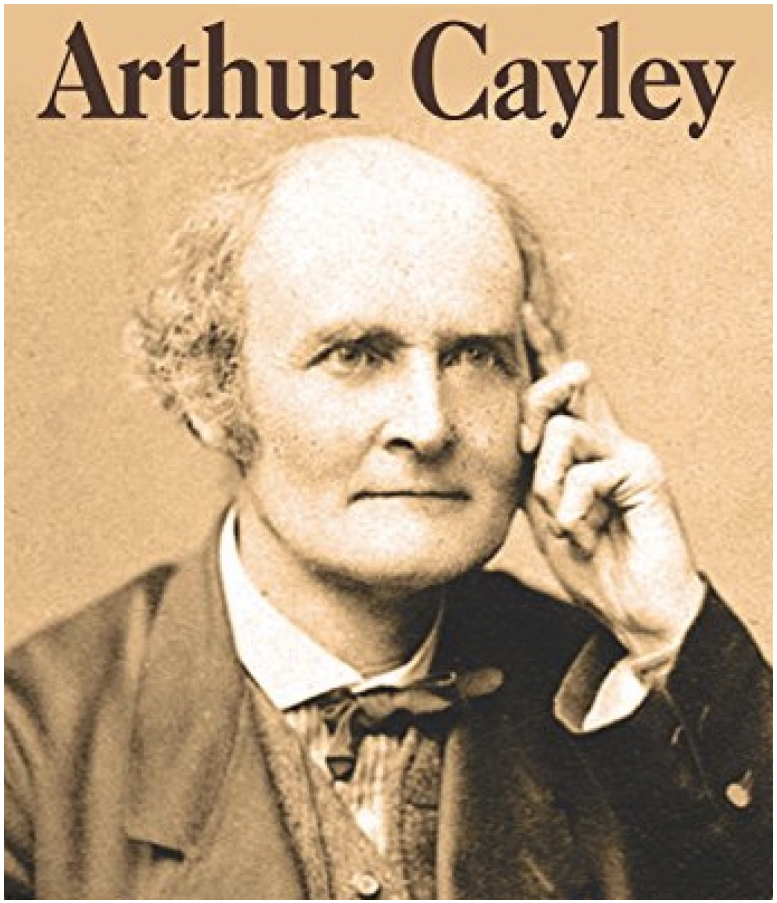

(1821–1895)

Cayley Graph Example:

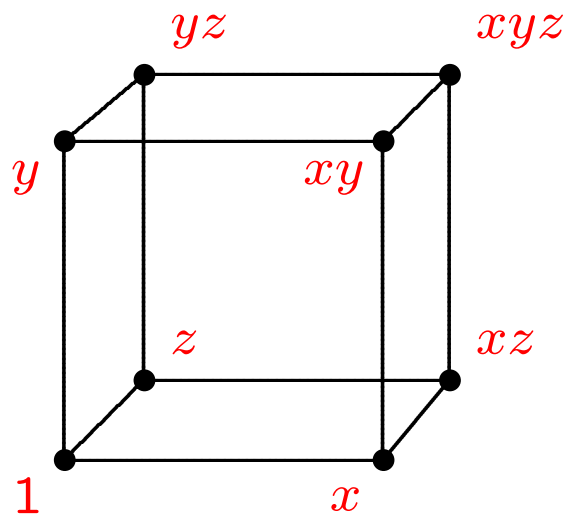

(isomorphic to Q_3)

$G = C_2 \times C_2 \times C_2$ (abelian) with generating set $X = \{x, y, z\}$

Ex: Is Q_3 a Cayley graph for some other group of order 8?

Ex: For which groups of order n is the complete graph K_n a Cayley graph?

Some elementary properties of Cayley graphs

- $\mathrm{Cay}(G, X)$ has order $|G|$
- $\mathrm{Cay}(G, X)$ is connected if and only if $G = \langle X \rangle$
- Every closed walk in $\mathrm{Cay}(G, X)$ corresponds to a relation satisfied by the generators from X
  [Why? $x_k^{e_k} \dots x_2^{e_2} x_1^{e_1} g = g$ if and only if $x_k^{e_k} \dots x_2^{e_2} x_1^{e_1} = 1$]
- $\mathrm{Cay}(G, X)$ is regular of valency $|X^{\pm}| = |X \cup X^{-}|$
- The group G acts transitively on the vertices of $\Gamma = \mathrm{Cay}(G, X)$ as a group of automorphisms, and it follows that every Cayley graph is vertex-transitive.

  Proof. Right-multiplication by any element $h \in G$ takes an edge $\{g, xg\}$ to the edge $\{gh, xgh\}$, and hence gives an automorphism of Γ taking vertex 1 to vertex h. ■

Conversely:

Theorem Let Γ be a connected finite graph, and suppose $\mathrm{Aut}(\Gamma)$ contains a subgroup G that acts regularly (sharply-transitively) on the vertices of Γ. Then Γ is a connected Cayley graph.

Sketch proof. Take any vertex v of Γ, and label it 1. Now label each vertex w of Γ with the unique element $g \in G$ that takes v to w, and let X be the set all such g for which w $(= v^g)$ is a neighbour of v. It follows fairly easily that Γ is isomorphic to $\mathrm{Cay}(G, X)$. ■

Another example: a Cayley graph for A_5

Let $G = A_5$ (the alternating group on 5 points)

and let $X = \{(1,2)(3,4), (1,2,3,4,5), (1,5,4,3,2)\}$.

Note that $(1,2)(3,4) \cdot (1,2,3,4,5) = (1,3,5)$, and it follows easily that the given set X generates G. Also X does not contain the identity element, and X is closed under inverses.

Hence the Cayley graph $\Gamma = \mathrm{Cay}(G, X)$ is a vertex-transitive, 3-valent connected graph of order 60.

This Cayley graph is the same as the **chemical molecule** C_{60} (also known as the Buckminsterfullerene):

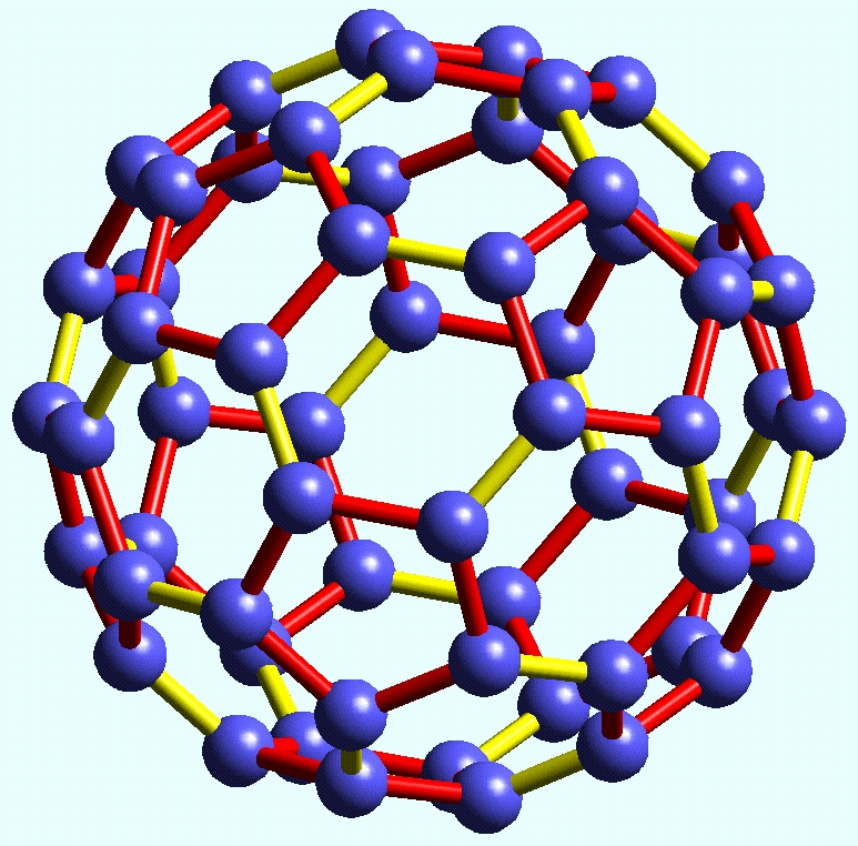

Here it is again, in a possibly more recognisable form:

Some (other) common properties of Cayley graphs

- Symmetry (as we have seen: vertex-transitive)
- Rigidity (e.g. the C_{60} molecule)
- Some were used by Max Dehn (in the early 1900s) to solve the 'word problem' for the fundamental group of an orientable surface of genus ≥ 2.
- Good broadcast properties – many Cayley graphs have large order-to-diameter ratio or small order-to-girth ratio

Note: Some improvements to the 'degree-diameter table' for simple undirected graphs have been made in just the last few months using Cayley graphs, as shown in the next slide ...

The Degree-Diameter Table (Sept 2024, revised July 2026)

New entries found this year using Cayley graphs, including for $(d, k) = (5, 5)$

$d\backslash k$	2	3	4	5	6	7	8	9	10
3	10	20	38	70	132	196	360	600	1 250
4	15	41	98	364	740	1 320	3 243	7 575	17 703
5	24	72	212	~~624~~ 648	2 772	5 516	17 030	57 840	187 056
6	32	111	390	1 404	7 917	19 383	76 461	331 387	1 253 615
7	50	168	672	2 756	11 988	52 768	249 660	1 223 050	6 007 230
8	57	253	1 100	5 060	39 672	131 137	734 820	4 243 100	24 897 161
9	74	585	1 550	8 268	75 893	279 616	1 697 688	12 123 288	65 866 350
10	91	650	2 286	13 140	134 690	583 083	4 293 452	27 997 191	201 038 922
11	104	715	3 200	19 500	156 864	1 001 268	7 442 328	72 933 102	600 380 000
12	133	786	4 680	29 470	359 772	1 999 500	15 924 326	158 158 875	1 506 252 500
13	162	856 979	6 560	40 260	531 440	3 322 080	29 927 790	249 155 760	3 077 200 700
14	183	~~916~~	8 200	57 837	816 294	6 200 460	55 913 932	600 123 780	7 041 746 081
15	187	1 215	11 712	76 518	1 417 248	8 599 986	90 001 236	1 171 998 164	10 012 349 898
16	200	1 600	14 640	132 496	1 771 560	14 882 658	140 559 416	2 025 125 476	12 951 451 931
17	274	1 610	19 040	133 144	3 217 872	18 495 162	220 990 700	3 372 648 954	15 317 070 720
18	307	1746 ~~1 620~~	23 800	171 828	4 022 340	26 515 120	323 037 476	5 768 971 167	16 659 077 632
19	338	1950 ~~1 638~~	23 970	221 676	4 024 707	39 123 116	501 001 000	8 855 580 344	18 155 097 232
20	381	2202 ~~1 958~~	34 952	281 820	8 947 848	55 625 185	762 374 779	12 951 451 931	78 186 295 824

How do we find Cayley graphs?

– e.g. all connected Cayley graphs of order n and valency d?

For small n, one way is to use the database of groups of order up to 2000 (excepting 1024), as created by Besche, Eick and O'Brien (2000), and available in GAP and MAGMA.

For valency 3, for example, we can search over two types of generating set X for the group G:

- $X = \{a, b, c\}$ where a, b, c have order 2, and
- $X = \{x, y\}$ where x has order 2 and y has order > 2.

Also we can use conjugacy within G (or within Aut(G)) to reduce the number of possibilities.

Another way:

For the first type of Cayley graphs of valency 3, the generating set X consists of three involutions, so the group G is a quotient of the finitely-presented group

$$\mathcal{G} = \langle\, a, b, c \mid a^2 = b^2 = c^2 = 1 \,\rangle.$$

Instead of checking possibilities for G from among groups of order up to n, we can (directly) find all quotients of $\mathcal{G}$ of order up to n using an algorithm for finding all normal subgroups of up to given index in a finitely-presented group. [This algorithm will be described later.]

Finding all such 3-valent Cayley graphs of order up to 100 takes only a few minutes (using MAGMA).

Note: This method doesn't find the Petersen graph!

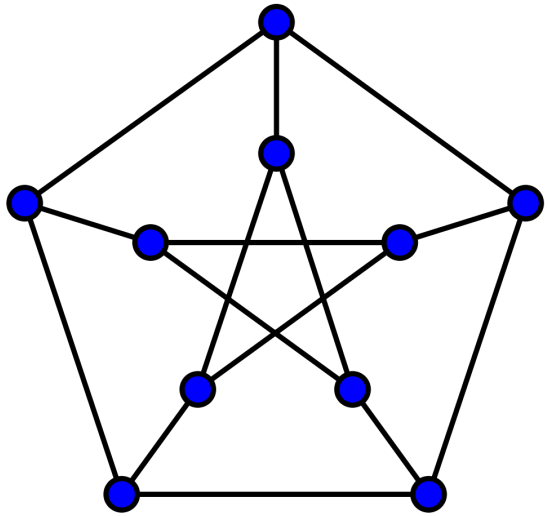

Reason: The Petersen graph is not a Cayley graph.

Why not? What are the possibilities for the generating set X in each of the groups of order 10? How does this relate to properties of the Petersen graph?

Answer(s):

By Sylow theory, every group of order 10 has at least one element of order 2, and has a normal subgroup of order 5. It follows easily that there are just two groups of order 10 – the cyclic group C_{10} and the dihedral group D_5.

If $G = C_{10}$ and u and v are two different elements of X, then $uv = vu$ so $u^{-1}v^{-1}uv = 1$, which gives a 4-cycle in the Cayley graph, but the Petersen graph has girth 5, so this is impossible. This leaves $G = D_5$ as the Cayley group.

Next, since the valency is odd, at least one element of X has order 2, say a. If the other elements of X are b and b^{-1} of order 5, then ab is a product of a reflection and a rotation and hence is a reflection (of order 2), and so $(ab)^2 = 1$, but again this gives a

4-cycle, contradiction. Thus X consists of three involutions, say a, b, c. But this too is impossible, as:

- the Petersen graph has no proper 3-edge-colouring, or
- $w = ba$ has order 5, and the Cayley graph has a Hamilton cycle $(1, a, w, aw, w^2, aw^2, w^{-2}, aw^{-2}, w^{-1}, aw^{-1})$, or
- the involution c must be aw or aw^2 or aw^{-2}, and each of these three possibilities gives a 4-cycle (viz. $(1, a, w, aw)$ or $(1, a, w^{-2}, aw^2)$ or $(1, aw^{-2}, w^{-1}, aw^{-1})$, or easier:
- the resulting Cayley graph is bipartite!

Hence the Petersen graph is not a Cayley graph. ■

But: The Petersen graph is symmetric, and has some very nice properties, and is constructible as a double-coset graph.

[See later]

§4. Schreier coset graphs & their applications

Let G be a group generated by a finite set $X = \{x_1, x_2, \ldots, x_d\}$.

Given any transitive permutation representation of G on a set Ω of size n, we may form a graph with vertex-set Ω, and with edges of the form $\alpha \text{ — } \alpha^{x_i}$ for $1 \leq i \leq d$.

e.g.

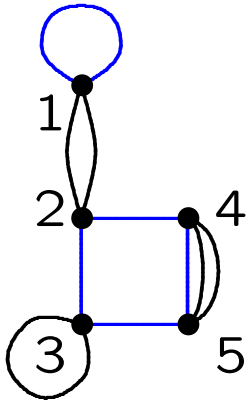

when $x_1 \mapsto (1,2)(4,5)$ and $x_2 \mapsto (2,3,5,4)$.

[Given any transitive permutation representation of G on a set Ω of size n, we may form a graph with vertex-set Ω, and with edges of the form $\alpha \text{ — } \alpha^{x_i}$ for $1 \leq i \leq d$.]

Similarly, if H is a subgroup of index n in G, we may form a graph whose vertices are the right cosets of H and whose edges are of the form $Hg \text{ — } Hgx_i$ for $1 \leq i \leq d$.

These two graphs are exactly the same when Ω is the coset space $(G : H)$ or when H is the stabiliser of a point of Ω.

The latter one is called the Schreier coset graph $\mathbf{\Sigma}(G, X, H)$.

Each is a generalisation of a Cayley graph (which occurs when the subgroup H is trivial).

Otto Schreier (1901–1929)

Some elementary properties of Schreier coset graphs:

- The graph is connected (as the group action is transitive)
- There can be loops and/or multiple edges
- Edges may be directed or labelled/coloured ... or not!
- The action of G can be recovered from the graph
  – or indeed defined by it

e.g.

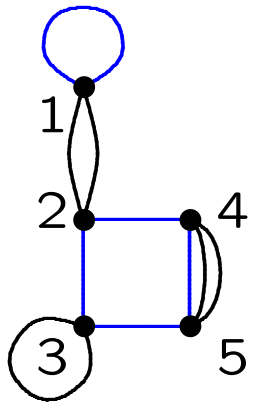

gives $x_1 \mapsto (1,2)(4,5)$ and $x_2 \mapsto (2,3,5,4)$.

• Every circuit in Σ based at the vertex labelled H gives an element of H expressed as a word on X

Why? Any path in Σ corresponds to a word $w = w(X)$ in the generators of G, and such a path from H is closed whenever $Hw = H$, which occurs if and only if $w \in H$.

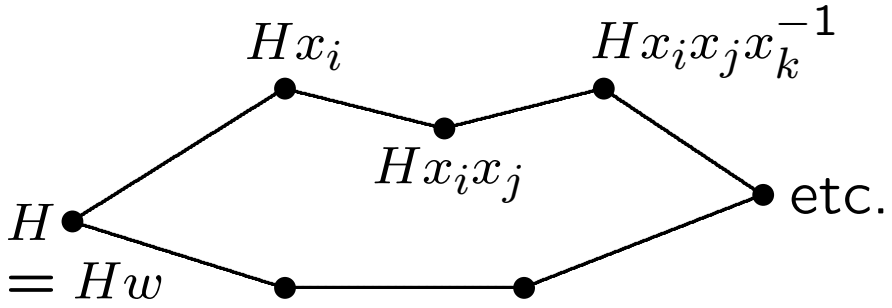

Importance/use of Schreier coset graphs [to follow]

- Visual representation (easier to see/understand than the permutations or 'coset table')
- Coset graphs can be used to construct representations (and build new ones from known representations)
- They give an easy proof of the Ree-Singerman theorem on necessary conditions for transitivity of a permutation representation of a finitely-generated group
- Depiction of Schreier transversals and Schreier generators
- Use in the Reidemeister-Schreier process (for finding a presentation for a given subgroup of finite index in a finitely-presented group)

Schreier coset graphs (cont.)

The Schreier coset graph $\Sigma(G, X, H)$ gives a diagrammatic representation of the natural action of G on cosets of H.

This action can also be given by a coset table, e.g. as on right in the following:

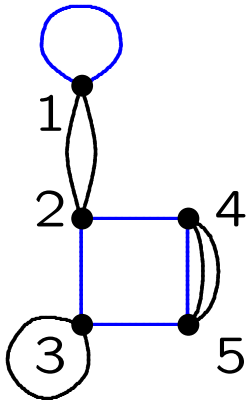

	x_1	x_2	x_1^{-1}	x_2^{-1}
1	2	1	2	1
2	1	3	1	4
3	3	5	3	2
4	5	2	5	5
5	4	4	4	3

when $x_1 \mapsto (1,2)(4,5)$ and $x_2 \mapsto (2,3,5,4)$.

Coset diagrams — simplified coset graphs

To make a Schreier coset graph easier to work with, we can sometimes simplify it by

- deleting loops (that occur for fixed points of generators)
- using single edges for 2-cycles of involutory generators
- ignoring the effect of redundant generators.

Special case: Triangle groups

Coset graphs for actions of the $(2, k, m)$ triangle group

$$\langle x, y, z \mid x^2 = y^k = z^m = xyz = 1 \rangle$$

can be simplified by deleting z-edges, and using heavy dots for fixed points of y, and polygons for non-trivial cycles of y.

The resulting figures are called (Schreier) coset diagrams rather than coset graphs.

Example:

Below is a coset diagram for an action of the $(2,3,7)$ triangle group $\langle\, x, y, z \mid x^2 = y^3 = z^7 = xyz = 1 \,\rangle$ on 7 points:

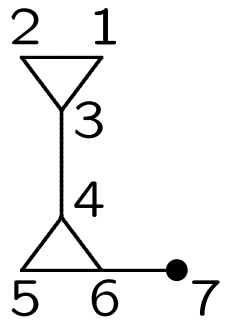

$x \mapsto (3,4)(6,7)$

$y \mapsto (1,2,3)(4,5,6)$

$z \mapsto (1,4,7,6,5,3,2)$

Composition of coset diagrams:

Often two coset diagrams for the same group G on (say) m and n points can be composed to produce a transitive permutation representation of larger degree $m + n$,

e.g.

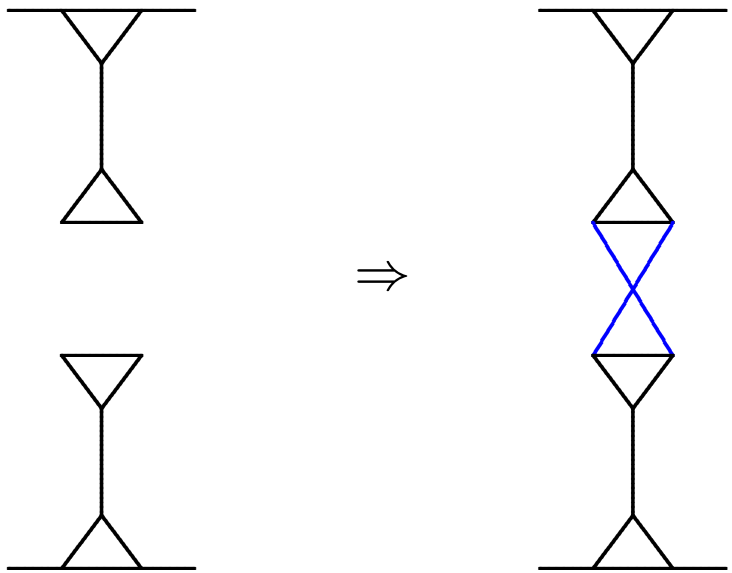

Ex: Check that the orders of x, y and xy are preserved.

What effect does this have?

Strange things can happen! For example, consider coset diagrams for transitive actions of the $(2, 3, 7)$ triangle group.

One can join together three such diagrams D1, D2, D3

D1 on 14 points, where the permutations generate a group isomorphic to PSL$(2, 13)$

D2 on 64 points, where the permutations generate a group isomorphic to A_{64}

D3 on 22 points, where the permutations generate a group isomorphic to A_{22}

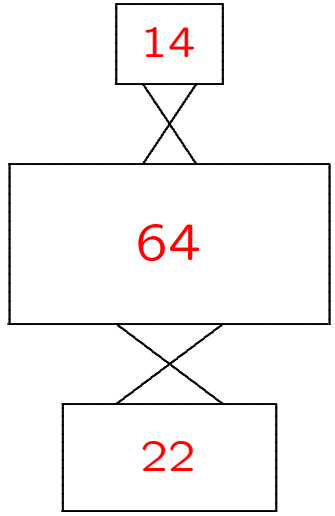

to get a diagram on 14+64+22 = 100 points, where the permutations generate the Hall-Janko group of order 604800.

Abelian extensions:

In some cases, where a coset diagram D for a group G may be joined together to another copy of itself in two different places, it is possible to string together n copies of the diagram into a circular chain (like a necklace) and get a new diagram in which the permutations generate a larger group B with an abelian normal subgroup K of exponent n such that $B/K \cong G$.

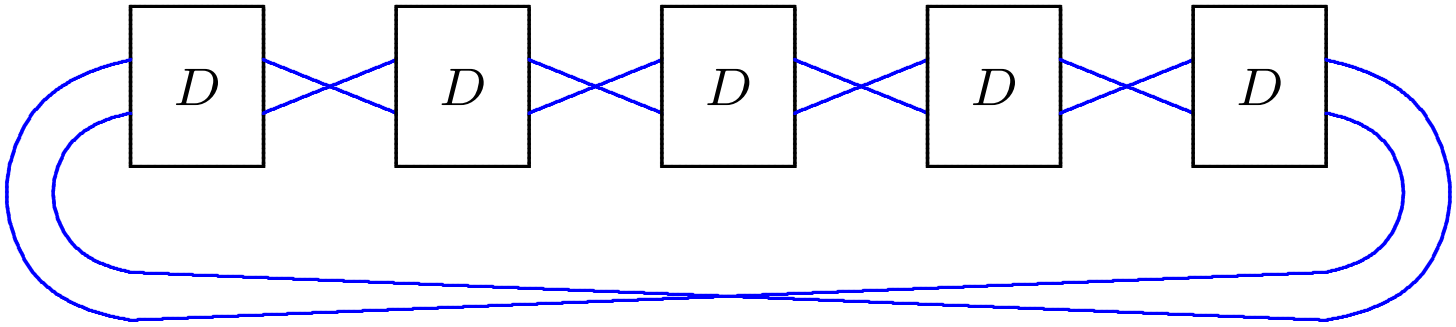

Proving groups are infinite:

When the previous construction is possible, one can string together an infinite number of copies of the given diagram D

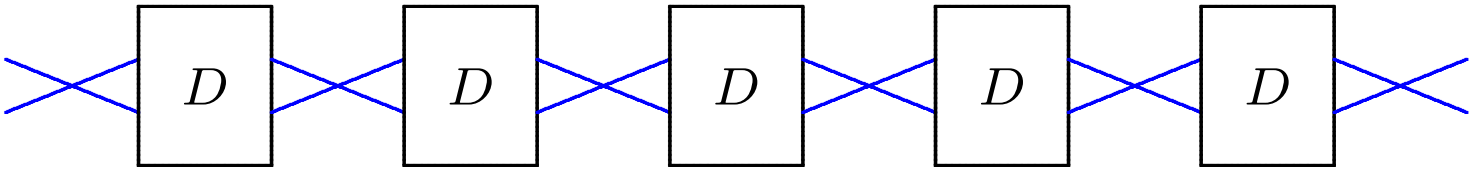

and get an infinite group!

This method can be used to prove that certain finitely-presented groups are infinite. It is equivalent to showing that some subgroup of finite index has infinite abelianisation (... also achievable by the Reidemeister-Schreier process, to be explained soon).

Exercise: Use a coset diagram with just 6 vertices to prove that the $(2, 3, 6)$ triangle group is infinite.

Alternating and symmetric quotients:

- If Diagrams P and Q with p points and q points have two 'handles' each (for attaching to other diagrams), then we can string together 'a' copies of P and 'b' copies of Q and get a much larger diagram on $m = ap + bq$ points.

- In particular, if $\gcd(p, q) = 1$ then the degree $m = ap + bq$ can be any sufficiently large positive integer.

- We can sometimes adjoin a single copy of an extra diagram R (with r points) to 'break symmetry', and make the permutations induced on the larger diagram generate the alternating group A_{m+r} or the symmetric group S_{m+r}.

- In this way, we can sometimes obtain all but finitely many A_n or S_n as quotients of a given finitely-presented group.

Some applications [by MC and students]

- For each $m \geq 7$, the $(2, 3, m)$ triangle group has all but finitely many alternating groups A_n among its quotients
- Every Fuchsian group has all but finitely many alternating groups A_n among its homomorphic images [Brent Everitt]
- There are infinitely many 5-arc-transitive connected finite 3-valent graphs
- There are infinitely many 7-arc-transitive connected finite 4-valent graphs [MC & Cameron Walker]
- There are infinitely many 'locally 9-arc-transitive' connected finite graphs [2026]
- There are infinitely many 5-arc-transitive Cayley graphs of valency 3, and infinitely many 7-arc-transitive Cayley graphs of valency 3^t+1 for each $t \geq 1$.

Further properties of the Schreier coset graph $\Sigma(G, X, H)$

First, recall that every circuit in Σ based at the vertex H gives an element of the subgroup H, written as a word on X.

Why? Any walk in Σ corresponds to a word $w = w(X)$ in the generators of G, and such a path from H is closed whenever $Hw = H$, which occurs if and only if $w \in H$.

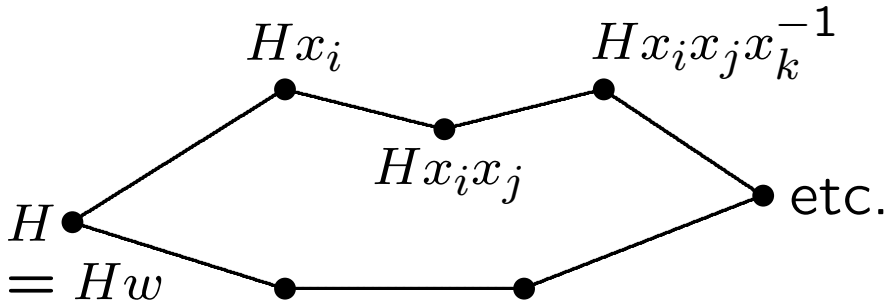

It follows that a set of generators for H can be created from the circuits in Σ based at the vertex labelled H.

Why? On one hand, all of the words coming from those circuits are elements of H. On the other hand, specifying all of them is enough to define the graph Σ, and hence the action on G on cosets, and hence H (as point-stabiliser).

e.g. 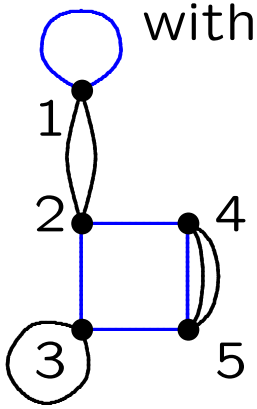

with $x \mapsto (1,2)(4,5)$ and $y \mapsto (2,3,5,4)$

gives $H = \langle\, y, x^2, xyxy^{-1}x^{-1}, xy^4x^{-1}, xy^{-1}xy^{-2}x, xy^2xyx \,\rangle$.

Q: How/when do we know that we have enough generators?

Reidemeister-Schreier Theory (when $|G:H|$ is finite)

First, a Schreier transversal for H in G is defined as a complete set $T = \{w_1, w_2, \ldots, w_n\}$ of representatives of the $n = |G:H|$ right cosets Hg of H, each expressed as a word in the generating set X, with the Schreier property, namely that for each $w \in T$, every 'initial sub-word' of w also lies in T.

e.g. 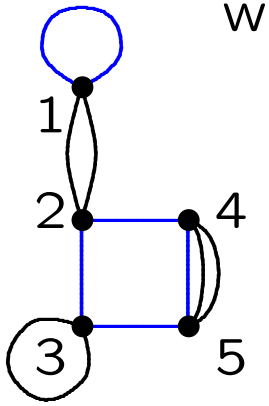

with $x \mapsto (1,2)(4,5)$ and $y \mapsto (2,3,5,4)$

gives $T = \{1, x, xy, xy^{-1}, xy^2\}$ as a Schreier transversal.

Generally, a Schreier transversal for H in G corresponds to a spanning tree for the coset graph Σ.

Why?

In any spanning tree T for Σ, we can label each vertex v with the product of the elements of X that label the edges of a path from 1 to v in T. The resulting n words form a transversal for H in G with the Schreier property, because every 'initial sub-word' of any one of them traces an initial sub-path from 1 to some vertex u of the corresponding path.

Exercise: Convince yourself that the converse is also true.

Example (as before)

Perm. representation $x \mapsto (1,2)(4,5)$ and $y \mapsto (2,3,5,4)$

Schreier transversal $T = \{1, x, xy, xy^{-1}, xy^2\}$

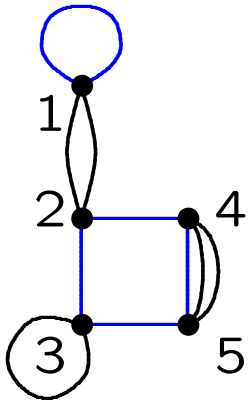

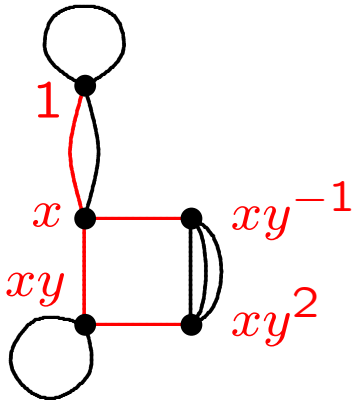

Check: Every 'initial sub-word' of each $w \in T$ also lies in T

Reidemeister-Schreier Theory (cont.)

Next, suppose T is a Schreier transversal for H in G, and for every element $g \in G$, denote by $\overline{g}$ the representative in T of the (right) coset Hg. In other words, $\overline{g}$ is the unique element of T for which $Hg = H\overline{g}$.

Schreier's Subgroup Lemma: The subgroup H is generated by the set $\{\, ux(\overline{ux})^{-1} : u \in T,\, x \in X \,\}$.

Sketch proof. First, each of the elements $ux(\overline{ux})^{-1}$ lies in H, because $Hux = H\overline{ux}$ for all $u \in T$, $x \in X$.

Indeed the x-edge from vertex Hu to vertex Hux creates a circuit in Σ based at vertex H, running from H to Hu along edges of T, then across to Hux $(= H\overline{ux})$, and back to H along edges of T again.

Conversely, every edge of Σ that is not in T has this form.

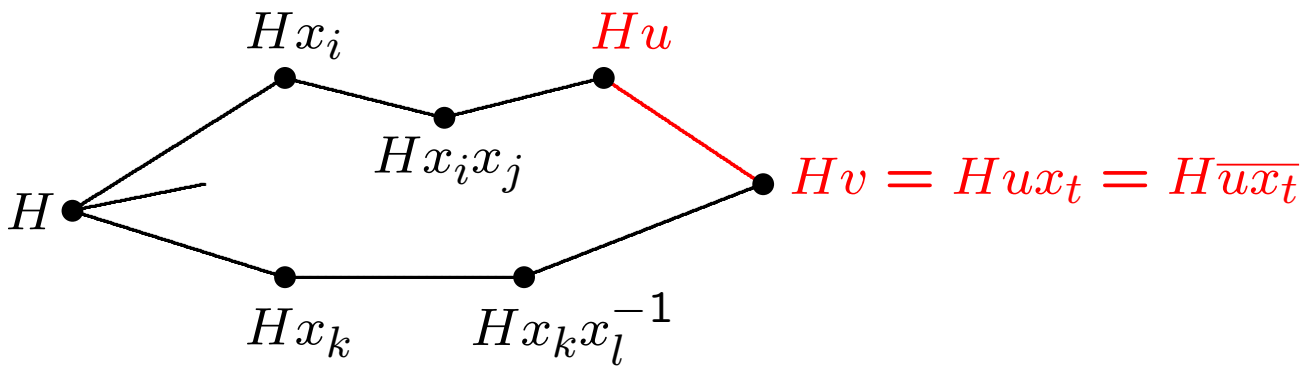

Hence these 'Schreier edges' are enough to add back all the edges of Σ, completing Σ, and determining H as the stabiliser of the vertex 'H' in the corresponding permutation representation of G (of degree $n = |G:H|$).

Summary: A Schreier generating-set for H in G corresponds to edges of the coset graph that are **not used in the spanning tree**.

Example

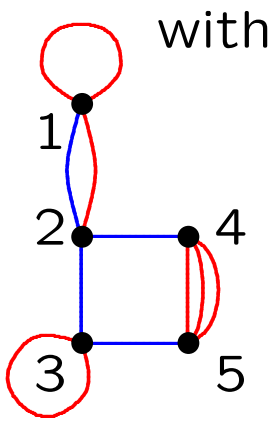

with $x \mapsto (1,2)(4,5)$ and $y \mapsto (2,3,5,4)$

A spanning tree has $5-1=4$ edges; for example, the edges $\{1,2\}$, $\{2,3\}$, $\{2,4\}$ and $\{4,5\}$ give us $\{1, x, xy, xy^{-1}, xy^2\}$ as a Schreier transversal. The total number of edges is 10, so $10-4=6$ of them do not lie in the spanning tree.

These give the six non-trivial Schreier generators y, x^2, $xyxy^{-1}x^{-1}$, xy^4x^{-1}, $xy^{-1}xy^{-2}x^{-1}$ and xy^2xyx^{-1} for H.

Corollary: Every subgroup of index n in a k-generator group can be generated by at most $nk - n + 1$ elements.

Proof. The coset graph has up to nk edges, and $n - 1$ of these are used in any spanning tree, so there remain at most $nk - (n - 1)$ edges for the non-trivial Schreier generators. ■

In fact, for a free group (which has no non-trivial relations), this bound is attained: every subgroup of index n in a free group of rank k is free of rank $nk - n + 1$. [See later]

Further application: the Reidemeister-Schreier process

Given a finitely-presented group $G = \langle X \,|\, R \rangle$, where X is the set of generators and R is the set of defining relations, and a subgroup H of finite index in G, how do we find a presentation for H (in terms of generators and relations)?

This can be done using further Reidemeister-Schreier theory, but the theory is very algebraic and can be difficult to follow. It is easier to explain (and implement) using coset graphs:

1) Construct the coset graph Σ, using the coset table
2) Use a spanning tree to give a Schreier transversal
3) Label the unused edges with Schreier generators
4) Trace each of the relators from R around the coset graph, starting from each of the vertices in turn, to obtain the Reidemeister-Schreier relations for H.

Example 1:

Let $G = \langle\, x, y \mid x^2 = y^3 = 1 \,\rangle$, and let H be the stabiliser of 1 in the permutation representation $x \mapsto (2,3)$, $y \mapsto (1,2,3)$:

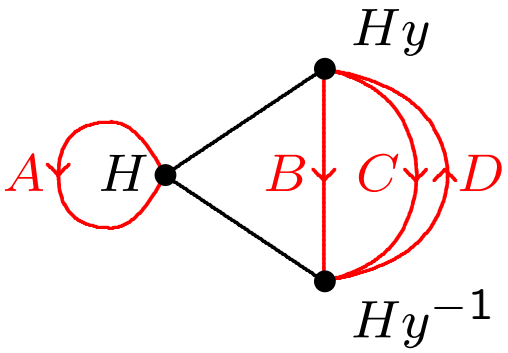

Schreier generators

$A = x$

$B = y^3 \ (= 1)$

$C = yxy$

$D = y^{-1}xy^{-1} \ (= C^{-1})$

Relation $x^2 = 1$ gives new relations $A^2 = 1$ and $CD = 1$

Relation $y^3 = 1$ gives new relation $B = 1$

So H has presentation $\langle\, A, C \mid A^2 = 1 \,\rangle$ via $(A, C) = (x, yxy)$.

Example 2:

Let $F_k = \langle x_1, x_2, \ldots, x_k \,|\, - \rangle$, the free group of rank k, and let H be any subgroup of index n in F_k.

Any spanning tree (with $n-1$ edges) in the Schreier coset graph $\Sigma(F_k, X, H)$ gives a Schreier transversal for H in F_k, and $nk-n+1$ Schreier generators. And then since F_n is free (which means that it has no non-trivial relations), the Reidemeister-Schreier process gives no relations for H.

Hence H is a free group of rank $nk-n+1$.

In particular, every subgroup of finite index in a finitely-generated free group is free.

Another application: the Ree-Singerman theorem

Theorem: Let G be a group generated by permutations $x_1, x_2, ..., x_d$ on a set Ω of size n, such that $x_1 x_2 ... x_d = 1$ (identity), and let c_i be the number of orbits of $\langle x_i \rangle$ on Ω. If G is transitive on Ω, then $c_1 + c_2 + \cdots + c_d \leq (d-2)n + 2$.

Note that if $d = 2$ then $x_2 = x_1^{-1}$, and the theorem gives $c_1 + c_2 \leq 2$, so $c_1 = c_2 = 1$, making x_1 a single n-cycle.

The case $d = 3$ was proved by Singerman (1970), and the general case ($d \geq 2$) was proved by Ree in 1971, using the theory of Riemann surfaces.

But: it can be proved easily using coset diagrams.

Proof. Embed the coset graph for the given permutation representation in an orientable surface with the edges to neighbours at each vertex Hg ordered naturally as follows:

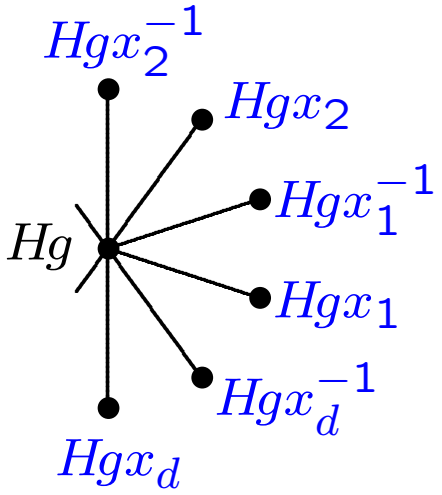

Here $|V| = n$ and $|E| = nd$, and we have

(1) a face bounded by x_i-edges for every cycle of each x_i

(2) a face bounded by edges $x_1, x_2, \ldots, x_d$ at each vertex,

and so $|F| = \sum_{1 \le i \le d} c_i + n.$

If the permutation representation is transitive, then the coset graph is connected, and then the Euler characteristic χ of the surface into which the graph has been embedded is given by the Euler-Poincaré formula:

$$\chi = |V| - |E| + |F| = n - nd + \sum_{1 \leq i \leq d} c_i + n.$$

But $\chi \leq 2$ for every orientable surface, and so from the above equation(s) it follows that

$$\sum_{1 \leq i \leq d} c_i \leq 2 - 2n + nd = (d-2)n + 2$$

as required. ■

Application: No transitive group of degree 167 can be generated by elements x, y, z such that $x^2 = y^3 = z^7 = xyz = 1$.

Why not? Assume that there **is** such a group. Let c_x, c_y and c_z be the number of cycles of x, y and z respectively.

Then the smallest possible number of cycles of z occurs when z has $\lfloor \frac{167}{7} \rfloor = 23$ cycles of length 7 and $167 - 23 \cdot 7 = 6$ fixed points, so $c_z \geq 29$. Similarly $c_y \geq 55 + 2 = 57$ and $c_x \geq 83 + 1 = 84$ (but in fact we can improve this to $c_x \geq 82 + 3 = 85$ because $x = (yz)^{-1}$ is an even permutation).

This gives $c_x + c_y + c_z \geq 84 + 57 + 29 = 170$ while on the other hand $(d-2)n + 2 = n + 2 = 169$, contradicting the Ree-Singerman theorem. ■

Summary:

Schreier coset graphs

- depict transitive actions of groups
- may help build large quotients of a group from small ones
- can be used to prove certain groups are infinite
- illustrate/assist the Reidemeister-Schreier process
- help prove the Ree-Singerman theorem.

Next:

How do we find good examples in the first place?

§5. Back-track search for subgroups

Let G be a group generated by $X = \{x_1, x_2, \ldots, x_m\}$.

Recall that right multiplication of right cosets of a subgroup H in G by elements of $X \cup X^{-1}$ can be depicted by a coset table like this:

	x_1	x_2	$\ldots$	x_1^{-1}	x_2^{-1}	$\ldots$
1	2	3		4		
2				1		
3					1	
4	1					
$\vdots$						

Subgroups of index $\leq n$ in G can be found (up to conjugacy) by a systematic enumeration of coset tables with $\leq n$ rows.

The low index subgroups algorithm

- Given $G = \langle X \mid R \rangle$ finitely-presented group
- Algorithm (due to Dietze & Schaps (but also Sims), in the 1970s) finds a representative of every conjugacy class of subgroups of index $\leq n$ (for given n) in G
- Backtrack search through a tree, with nodes at level k corresponding to k-generator (pseudo-)subgroups H
- Enumeration (by Todd-Coxeter) of cosets of H
- Create branches to new nodes at the next level (if necessary) by identifying pairs of cosets: forcing $Hg_i = Hg_j$ is equivalent to adding $g_i g_j^{-1}$ to a set of generators for H

Low index subgroups algorithm (cont.)

- Output includes generators for the subgroup H, and/or permutations induced by generators (in X) on cosets of H

- Schreier's theorem ensures that every subgroup of given index m in $G = \langle X \,|\, R \rangle$ is finitely-generated (so will be found)

- Conjugates of subgroups found earlier may be eliminated easily (by a test on the coset table)

- This can facilitated by normal ordering of cosets in the coset table – where each 'new' coset is labelled with the smallest unused positive integer

- Example: (PTO)

	x_1	x_2	x_3	x_1^{-1}	x_2^{-1}	x_3^{-1}
1	2	3	4	5		
2				1		
3					1	
4						1
5	1					
:						

Cosets 1: H
2: Hx_1
3: Hx_2
4: Hx_3
5: Hx_1^{-1}

Transversal $\{1, x_1, x_2, x_3, x_1^{-1}, \dots\}$

Processing: e.g. $1x_1^{-1} = 5 \Rightarrow 5x_1 = 1$

Forcing coincidences

The key to the standard 'Low index subgroups algorithm' is to define more than n cosets, and then force coincidences between them, using the fact that $Ha = Hb \Leftrightarrow ab^{-1} \in H$.

The algorithm starts with the identity subgroup and attempts to enumerate its right cosets, constructing a partial transversal, say $\{u_1, u_2, u_3, \dots\}$. Then (or at any stage) if more than n cosets are defined, all possible coincidences between two cosets Hu_i and Hu_j are considered, for $1 \leq i < j \leq n + 1$.

Often such a coincidence is found to produce a subgroup H that is conjugate to one found previously, in which case that coincidence is rejected and the next one is looked at.

If not rejected, then $u_i u_j^{-1}$ is added to a (partial) set of generators for H and the search continues:

	x_1	x_2	x_3	x_1^{-1}	x_2^{-1}	x_3^{-1}
1	2	3	4	5		
2				1		
3					1	
4						1
5	1					
:						

Level	Coincidence	Additional generator
1	$1 = 2$	x_1^{-1}
	$1 = 3$	x_2^{-1}
	$2 = 3$	$x_1 x_2^{-1}$
	:	:
	$1 = 5$	x_1
	$2 = 5$	x_1^2
	$3 = 5$	$x_2 x_1$
	:	:

Branching/backtrack process

Systematic enumeration of coincidences between cosets (and adding new generators for H) sets up a branching process:

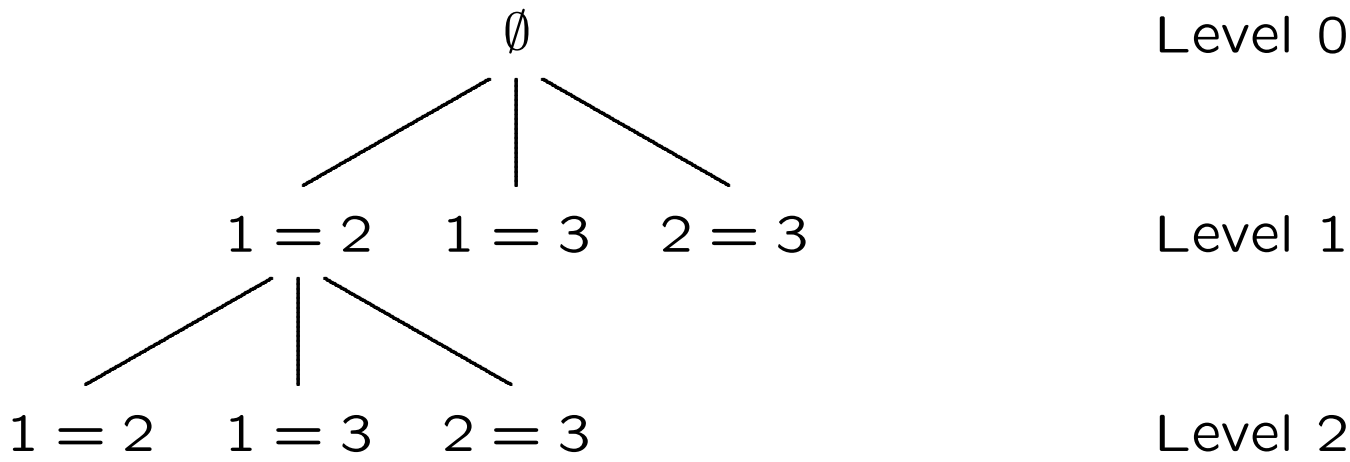

A backtrack search will terminate (given sufficient time and memory), by Schreier's theorem: every subgroup of finite index in a finitely-generated group is itself finitely-generated.

Example:

Let $G = \langle\, x, y \,|\, x^2 = y^3 = 1 \,\rangle$, which is the modular group, and look for subgroups of index up to 4. We get these:

#	Coincidences	Index	Generators
1	$1 = 2,\ 1 = 2$	1	$x,\ y$
2	$1 = 2,\ 2 = 4,\ 3 = 4$	3	$x,\ yxy^{-1},\ y^{-1}xy$
3	$1 = 2,\ 2 = 4,\ 4 = 5$	4	$x,\ yxy^{-1},\ y^{-1}xy^{-1}xy$
4	$1 = 2,\ 3 = 4$	3	$x,\ y^{-1}xy^{-1}$
5	$1 = 3,\ 2 = 3$	2	$y^{-1},\ xy^{-1}x$
6	$1 = 3,\ 4 = 5$	4	$y^{-1},\ xy^{-1}xy^{-1}x$

Low index normal subgroups

Small homomorphic images of a finitely-presented group G can be found as the groups of permutations induced by G on cosets of subgroups of small index. This gives G/K where K is the core of H, but produces only images that have small degree faithful permutation representations.

Alternatively, the (standard) low index subgroups method can be adapted to produce only normal subgroups.

A new method was developed in 2005 by Derek Holt and his student David Firth, which systematically enumerates the possibilities for the composition series of the factor group G/K, for any normal subgroup K of small index in G.

This method has produced the automorphism groups of lots of symmetric structures (including graphs, maps & polytopes).

Summary:

- Standard 'Low Index Subgroups' algorithm
- New variant for finding normal subgroups only
- These two methods can help find all small degree transitive permutation representations and all small quotients of a given finitely-presented group.

§6. Double-coset graphs

Let G be a group, H be a subgroup of G, and a be an element of G such that $a^2 \in H$. Now define a graph $\Gamma = \Gamma(G, H, a)$ by

$$V(\Gamma) = \{Hg : g \in G\}$$
$$E(\Gamma) = \{\{Hx, Hy\} : x, y \in G \mid xy^{-1} \in HaH\}.$$

For example, vertex H is adjacent to Hah for all $h \in H$.

Now G induces a group of automorphisms of Γ by right multiplication, since $(xg)(yg)^{-1} = xgg^{-1}y^{-1} = xy^{-1} \in HaH$ whenever $\{Hx, Hy\} \in E(\Gamma)$).

This action is vertex-transitive (since $(Hx)x^{-1}y = Hy$), with vertex-stabiliser $G_H = \{g \in G : Hg = H\} = H$ itself, which acts transitively on the neighbours Hah of H.

Thus $\Gamma = \Gamma(G, H, a)$ is arc-transitive, or 'symmetric'.

Example: a double-coset graph for A_5

Let $G = A_5$ (the alternating group on 5 points), and take $H = \langle\,(1,2,3),(1,2)(4,5)\,\rangle \cong S_3$ and $a = (1,4)(2,5)$.

Then the coset graph $\Gamma(G,H,a)$ is a connected arc-transitive 3-valent graph of order 10 which turns out to be isomorphic to the Petersen graph ... as promised earlier.

This graph can also be constructed as a double-coset graph for $G = S_5$, using the subgroup $H = \langle\,(1,2,3),(1,2),(4,5)\,\rangle \cong S_3 \times C_2$ and the element $a = (1,4)(2,5)$.

In fact the full automorphism group of the Petersen graph is S_5.

Special case: 3-**valent arc-transitive graphs**

- By work of Tutte and Djoković & Miller, it is known that if Γ is a finite connected 3-valent arc-transitive graph, then $G = \mathrm{Aut}\,\Gamma$ is a quotient of one of seven finitely-presented groups now known as G_1, G_2^1, G_2^2, G_3, G_4^1, G_4^2, G_5.

- Conversely, if G is any non-degenerate quotient of one of those groups, then we can use the double-coset graph construction to obtain a finite connected 3-valent arc-transitive graph Γ on which G acts as a group of automorphisms.

- For example, $G = A_5$ is a quotient of the group

$$G_2^1 = \langle\, h, a, p \,|\, h^3 = a^2 = p^2 = 1,\ apa = p,\ php = h^{-1} \,\rangle$$

via $h \mapsto (1,2,3)$, $a \mapsto (1,4)(2,5)$, $p \mapsto (1,2)(4,5)$, and from this we get the Petersen graph (as before).

Small 3-valent symmetric graphs

- Take each one of G_1, G_2^1, G_2^2, G_3, G_4^1, G_4^2, G_5 in turn

- Use the 'Low Index Normal Subgroups' algorithm (e.g. in MAGMA) to find all normal subgroups of index up to n

- For each normal subgroup K, let G be the quotient of the given group by K, and use the double-coset graph construction to obtain a finite connected 3-valent symmetric graph Γ on which G acts as a group of automorphisms

- Check whether G is the full automorphism group of Γ (using GAP or MAGMA); discard the graph if it's not

- This approach has provided a census of all such graphs on up to 10000 vertices.

Arc-transitive maps (sometimes called **regular maps**)

These are maximally symmetric 2-cell embeddings of connected graphs or multigraphs on surfaces:

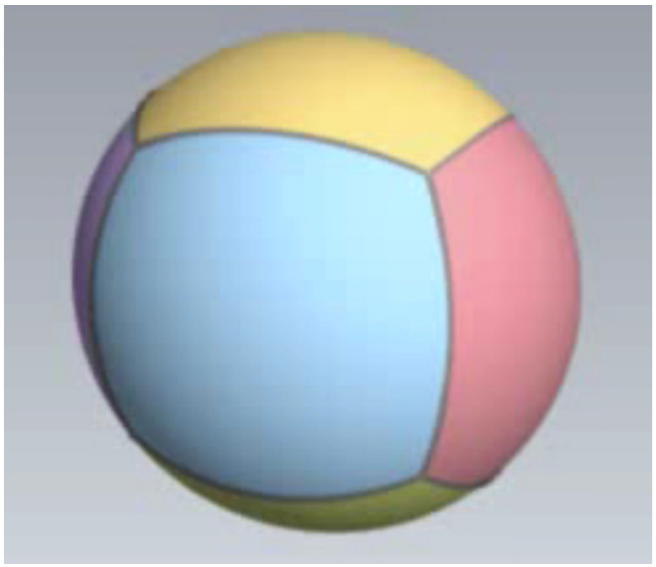

Q_3 on sphere

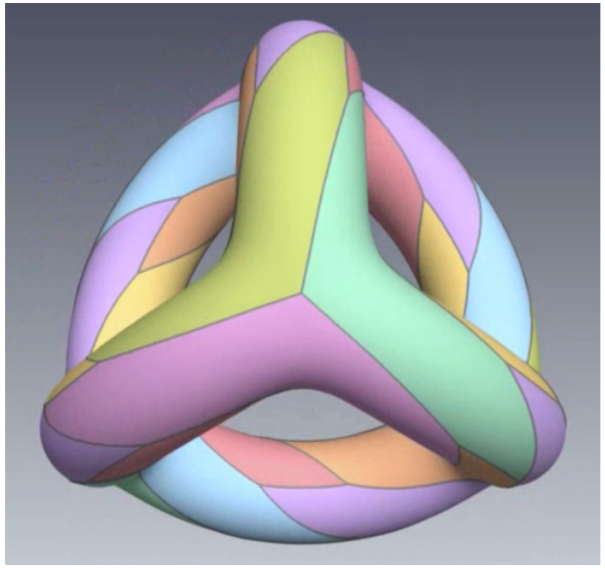

Klein map (genus 3)

Hurwitz map of genus 7

[Image created by Jarke van Wijk (Eindhoven)]

Some key points:

A flag of a map is an incident vertex-edge-face triple (v, e, f).

An automorphism of a map M is a bijection taking vertices to vertices, edges to edges, and faces to faces, preserving incidence between them.

Apart from some weird exceptions, every automorphism of a map M is uniquely determined by its effect on a given flag, and it follows that $|\mathrm{Aut}(M)| \leq 4|E|$ where E is the edge set.

A map M is (fully) regular if $\mathrm{Aut}(M)$ is transitive on flags.

A map M on an orientable surface is orientably-regular if the orientation-preserving subgroup $\mathrm{Aut}^+(M)$ of $\mathrm{Aut}(M)$ is transitive on arcs, and in that case $|\mathrm{Aut}^+(M)| \leq 2|E|$.

Moreover, such a map M is either reflexible (when $\mathrm{Aut}(M)$ is transitive on all flags), or otherwise chiral (irreflexible).

Construction for arc-transitive maps

Let G be a group that is generated by two elements R and S such that $R^m = S^k = (RS)^2 = 1$, where $k, m \geq 2$.

Now construct the graph $\Gamma = \Gamma(G, \langle S \rangle, RS)$, on which the group G acts as an arc-transitive group of automorphisms, and define an embedding of this graph into a surface by taking as faces the cycles induced by R on the vertices of Γ.

This turns Γ into an arc-transitive map $M = M(G, R, S)$ on some orientable surface, with the given group G inducing its group $\mathrm{Aut}^o M$ of orientation-preserving automorphisms.

The cyclic subgroups of G generated by S, RS and R are stabilisers of a vertex, edge and face, respectively.

Similar construction for fully regular maps:

Let G be a group that is generated by three involutions a, b, c such that $(ab)^2 = (bc)^k = (ca)^m = 1$, where $k, m \geq 3$.

Now construct the graph $\Gamma = \Gamma(G, \langle b, c\rangle, a)$, and turn this into a regular map $M = M(G, a, b, c)$, by taking as faces the cycles induced by ca on the vertices of Γ.

The dihedral subgroups $\langle b, c\rangle$, $\langle a, b\rangle$ and $\langle a, c\rangle$ are stabilisers of a vertex, edge and face, respectively.

This map is orientable if and only if the subgroup generated by $R = ca$ and $S = bc$ has index 2 in G; otherwise, the latter subgroup has index 1 and the surface is non-orientable.

Regular maps of small genus

- Take the extended $(2, \infty, \infty)$ triangle group

$$\Delta = \langle\, a, b, c \mid a^2 = b^2 = c^2 = (ab)^2 = 1 \,\rangle$$

— that is, without specifying the orders of bc and ca

- Use the 'Low Index Normal Subgroups' algorithm to find all normal subgroups of small index in Δ

- For each torsion-free normal subgroup K, let $G = \Delta/K$, and use the previous construction to obtain a regular map (on which G acts as a group of automorphisms)

- Check the genus and orientability of the map

- This approach has provided a census of all regular maps of Euler characteristic -1 to -3000 (available online since 2025).

Chiral arc-transitive maps

Use the arc-transitive maps construction, and check whether the kernel K in Δ^o is normal in the full triangle group Δ:

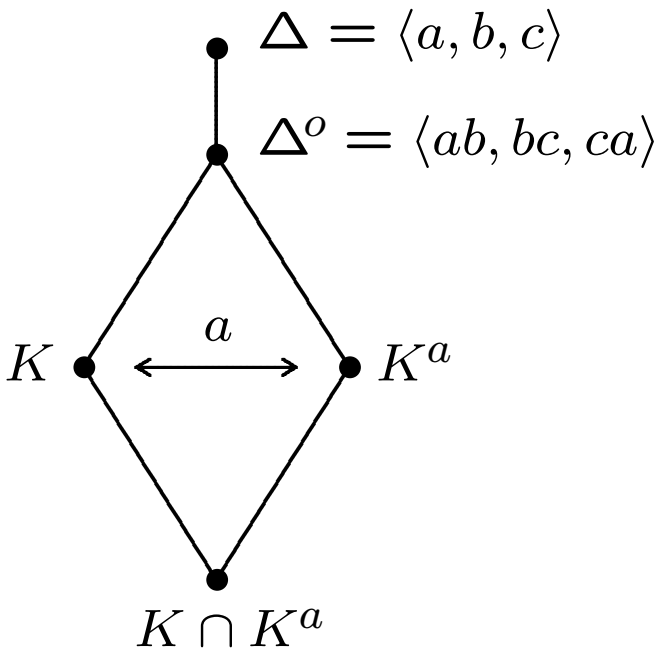

$K = K^a$ iff map M is reflexible

(when $\mathrm{Aut}^o M \cong \Delta^o / K$)

Regular and chiral polytopes

A similar approach can be taken to construct polytopes with large automorphism groups. Polytopes are generalisations of maps and polyhedra to higher dimensions.

The largest possible automorphism groups are quotients of Coxeter groups (or of subgroups of index 2 in Coxeter groups), satisfying an additional test on the intersections of certain subgroups — in order to meet criteria of 'strong flag connectedness' and the 'diamond condition'.

This method was used in 2008 to construct the first known examples of finite chiral polytopes (with maximum possible rotational symmetry) of ranks 5, 6, 7 and 8. The latter also enabled a direct construction of chiral n-polytopes for all $n \geq 3$ in 2024.

§7. Möbius Inversion

In number theory, the Möbius function $\mu\colon \mathbb{N} \to \mathbb{Z}$ is defined by

$$\mu(n) = \begin{cases} 1 & \text{if } n \text{ is a product of an even no. of distinct primes} \\ -1 & \text{if } n \text{ is a product of an odd no. of distinct primes} \\ 0 & \text{if } n \text{ is divisible by the square of some prime.} \end{cases}$$

It has this nice property, called the Möbius inversion formula:

Theorem: If $f\colon \mathbb{N} \to \mathbb{C}$ and $g\colon \mathbb{N} \to \mathbb{C}$ are functions with f defined in terms of g by $f(n) = \sum_{d\,|\,n} g(d)$ for all $n \in \mathbb{N}$, then

$$g(n) = \sum_{d\,|\,n} \mu(d)\, f\left(\frac{n}{d}\right) \quad \text{for all } n \in \mathbb{N}.$$

In other words, μ helps determine g from the values of f.

Möbius Inversion (cont.)

The same thing can be proved in a more general setting, thanks to work by Louis Weisner and Philip Hall in the 1930s.

Setting: Let $\mathcal{P}$ be any partially ordered set with a maximum element M, and define the function $\mu\colon \mathcal{P} \to \mathbb{Z}$ by setting

$$\mu(x) = \sum_{s=0}^{\infty} (-1)^s \, n_s(x)$$

where $n_s(x)$ is the total number of descending chains of the form $M = x_0 > x_1 > \cdots > x_s = x$ from M to x in $\mathcal{P}$.

Note: Such a chain may be refineable, but increasing its length s.

Observe also that $\mu(M) = 1$, while $\mu(x) = -1$ for every maximal element x of $\mathcal{P}$, and so on.

Key property of μ : $\sum_{x \geq y} \mu(x) = 0$ for every $y \in \mathcal{P} \setminus \{M\}$.

Proof. Observe that $\sum_{x \geq y} \mu(x) = \sum_{x \geq y} \sum_{s=0}^{\infty} (-1)^s n_s(x)$ can be split as the sum of $\sum_{s=0}^{\infty} (-1)^s n_s(y)$ and $\sum_{s=0}^{\infty} (-1)^s \sum_{x > y} n_s(x)$.

Every chain $M = y_0 > y_1 > \cdots > y_{s-1} > y_s = y$ of length s from M to y corresponds uniquely to a chain of length $s-1$ from M to the element $x = y_{s-1}$, and this is counted just once in the last of the above sums, with coefficient $(-1)^{s-1}$.

Hence the terms of $\sum_{s=0}^{\infty} (-1)^s n_s(y)$ and $\sum_{s=0}^{\infty} (-1)^s \sum_{x > y} n_s(x)$ cancel each other and give 0. ■

The Hall-Weisner Theorem for Möbius Inversion:

If $f : \mathcal{P} \to \mathbb{C}$ and $g : \mathcal{P} \to \mathbb{C}$ are functions s.t. f is defined in terms of g by $f(x) = \sum_{y \leq x} g(y)$ for all $x \in \mathcal{P}$, then $g(M) = \sum_{x \in \mathcal{P}} \mu(x)\, f(x)$.

In other words, μ helps determine $g(M)$ from the values of f.

Proof. First we have $\sum_{x \in \mathcal{P}} \mu(x) f(x) = \sum_{x \in \mathcal{P}} \mu(x) \sum_{y \leq x} g(y)$.

Now the right hand side rearranges to $\sum_{y \in \mathcal{P}} g(y) \sum_{x \geq y} \mu(x)$,

and then this splits according to whether $y = M$ or $y < M$.

Indeed since $\mu(M) = 1$ and $\sum_{x \geq y} \mu(x) = 0$ for all $y < M$, it

simplifies to just $g(M)\mu(M)$, which is $g(M)$. ■

The above theorem was proved by Weisner for lattices (with maximum element and minimum element, and 'meets' and 'joins'), but Hall proved it more generally for posets with a maximum element, using the 'key property' given earlier.

We will soon see how Philip Hall and others applied Möbius inversion to the poset of subgroups of a group. This can even be applied in some case where the group is infinite.

Exercise: Show how the number-theoretic Möbius inversion formula follows from the Hall-Weisner theorem. [Note: you may use symmetry of the lattice of divisors of n.]

Optional exercise: Look up how the Inclusion-Exclusion Principle relates to Möbius inversion on the lattice of subsets of a finite set.

The Möbius function on a subgroup lattice:

For any finite group G, let $\mathcal{S}$ be the lattice of all subgroups of G, with $A \wedge B = A \cap B$ (intersection), and $A \vee B = \langle A, B \rangle$ (the subgroup generated by A and B, or equivalently, the intersection of all subgroups of G that contain A and B). Also let μ be the Möbius function on $\mathcal{S}$.

Exercise: Show how the values of μ on the subgroup lattice of an infinite cyclic group are related to those of the values of the number-theoretic Möbius function.

Exercise: The group A_4 has three subgroups of order 2, four of order 3, and one each of orders 1, 4 and 12, and no others. Find the values of μ on each of these subgroups.

Exercise: Find μ on the subgroup lattice of A_5.

Lemma: For every proper subgroup $H \leq G$, either $\mu(H) = 0$ or H is the intersection of some maximal subgroups of G.

Proof. We use induction on the index $|G:H|$.

First, if H itself is maximal in G, then $\mu(H) = -1$.

Now let J be the intersection of all maximal subgroups of G containing H, and suppose $H < J$. Then by the earlier 'key property' of μ, we know that $\mu(H) = -\sum_{K > H} \mu(K)$.

Next, by induction, we may suppose that $\mu(K) = 0$ for every subgroup $K > H$ that does not contain J (for then K cannot be the intersection of maximal subgroups of G), and from this we find that $\sum_{K > H} \mu(K) = \sum_{K \geq J} \mu(K)$, which is 0, by the 'key property'.

Hence either $H = J$, or $\mu(H) = -\sum_{K \geq J} \mu(K) = 0$. ■

Application of the Hall-Weisner theorem to groups:

Suppose we are interested in knowing whether the group G can be generated by a subset X that has a given property (P)
– e.g. X is a set $\{a, b\}$ of two elements of orders 2 and 3.

Then we can define two functions f and g on the subgroup lattice $\mathcal{S}$ of G by setting for each subgroup H of G

$f(H) =$ number of subsets of type (P) contained in H, and
$g(H) =$ number of subsets of type (P) that generate H.

Since every subset generates a unique subgroup, we have
$f(H) = \sum_{K \leq H} g(K)$ for all $H \leq G$, so by Möbius inversion, the total number of subsets of type (P) that generate G is
$g(M) = \sum_{H \leq G} \mu(H)\, f(H)$ – which is often easily computable.

Banal example: Is S_4 generated by two elements of order 4?

in S_4 there are 6 elements of order 4 in S_4 (the 4-cycles), and the only subgroups containing an element of order 4 are the three copies of C_4, the three copies of D_4, and S_4 itself. Hence the number $f(H)$ of subsets of a subgroup H consisting of two elements of order 4 is

$$f(H) = \begin{cases} 1 & \text{when } H \cong C_4 \\ 1 & \text{when } H \cong D_4 \\ {}^6C_2 = 15 & \text{when } H = S_4 \\ 0 & \text{otherwise.} \end{cases}$$

Also $\mu(S_4) = 1$ and $\mu(H) = -1$ for all $H \cong D_4$ (maximal), while $\mu(H) = 0$ for all $H \cong C_4$. Thus

$$g(S_4) = \sum_{H \leq S_4} \mu(H)\, f(H) = 3(0 \cdot 1) + 3(-1 \cdot 1) + 1(1 \cdot 15) = 12,$$

so there are 12 generating pairs $\{x, y\}$ for S_4 with $o(x) = o(y) = 4$.

Some further applications/examples:

- Philip Hall found the no. of generating pairs (x, y) for PSL$(2, p)$ s.t. $x^2 = y^3 = 1$ for every prime p [Hall (1936)]

- Martin Downs extended this to PSL$(2, q)$ for every prime-power q [Downs (1991)]

- Martin Downs extended this further to generating pairs (x, y) for PSL$(2, q)$ s.t. $x^k = y^\ell = (xy)^m = 1$ [Downs (1997)]

- For every $s > 1$, the Suzuki group $Sz(2^s)$ can be generated by two elements x and y such that $x^2 = y^4 = (xy)^5 = 1$ (and moreover, is the automorphism group of a chiral map of type $\{4, 5\}$). [Jones & Silver (1993)]

Thank you for listening!

Some references

Cayley graphs

L. Babai & C.D. Godsil, On the automorphism groups of almost all Cayley graphs, *European J. Combin.* 3 (1982), 9–15.

M.D.E. Conder, On symmetries of Cayley graphs and the graphs underlying regular maps, *J. Algebra* 321 (2009), 3112–3127.

P.R. Hafner, Large Cayley graphs and digraphs with small degree and diameter, in *Computational algebra and number theory (Sydney, 1992)*, Kluwer Acad. Publ. (Dordrecht), 1995, pp.291–302.

M-C. Heydemann, Cayley graphs and interconnection networks, in *Graph symmetry (Montreal, PQ, 1996)*, Kluwer Acad. Publ. (Dordrecht), 1997, pp.167–224.

W. Imrich & M.E. Watkins, On automorphism groups of Cayley graphs. *Period. Math. Hungar.* 7 (1976), 243–258.

C.H. Li, On isomorphisms of finite Cayley graphs—a survey, *Discrete Math.* 256 (2002), 301–334.

Schreier coset graphs

M. Conder, Schreier coset graphs and their applications, *RIMS Kokyuroku* 794 (1992), 169–175.

M. Conder & J. McKay, A necessary condition for transitivity of a finite permutation group, *Bull. L.M.S.* 20 (1988), 235–238.

H.S.M. Coxeter & W.O.J. Moser, *Generators and Relations for Discrete Groups*, 4th ed. Springer-Verlag (Berlin), 1980.

Low index subgroup methods

M. Conder & P. Dobcsányi, Applications and adaptations of the low index subgroups procedure, *Mathematics of Computation* 74 (2005), 485–497.

A. Dietze & M. Schaps, Determining subgroups of a given finite index in a finitely presented group, *Canadian J. Math.* 26 (1974), 769–782.

D.F. Holt, B. Eick & E.A. O'Brien, *Handbook of Computational Group Theory*, CRC Press, 2005.

D. Firth, *An algorithm to find normal subgroups of a finitely presented group up to a given index*, PhD Thesis, University of Warwick, 2005.

C.C. Sims, *Computation with Finitely-Presented Groups* (Cambridge Univ. Press, 1994).

Edge-transitive graphs and double-coset graphs

M. Conder, An infinite family of 5-arc-transitive cubic graphs, *Ars Combinatoria* 25A (1988), 95–108.

M. Conder, Schreier coset graphs and their applications, *RIMS Kokyuroku* 794 (1992), 169–175.

M.D.E. Conder, The infinitude of locally 9-arc-transitive graphs, *J. Algebra* 692 (2026), 582–591.

M. Conder & P. Dobcsányi, Trivalent symmetric graphs on up to 768 vertices, *J. Combin. Math. Combin. Computing* 40 (2002), 41–63.

M. Conder & P. Lorimer, Automorphism groups of symmetric graphs of valency 3, *J. Combin. Theory Ser. B* 47 (1989), 60–72.

M. Conder & P. Potočnik, Edge-transitive cubic graphs: analysis, cataloguing and enumeration, *J. Algebra* 685 (2026), 703–737.

M.D.E. Conder & C.G. Walker, The infinitude of 7-arc-transitive graphs, *J. Algebra* 208 (1998), 619–629.

D.Z. Djoković & G.L. Miller, Regular groups of automorphisms of cubic graphs, *J. Combin. Theory Ser. B* 29 (1980), 195–230.

V. Trofimov, On the action of a group on a graph, *Acta Appl. Math.* 29 (1992), 161–170.

W.T. Tutte, A family of cubical graphs, *Proc. Camb. Phil. Soc.* 43 (1947), 459–474.

W.T. Tutte, On the symmetry of cubic graphs, *Canad. J. Math.* 11 (1959), 621–624.

Richard Weiss, The non-existence of 8-transitive graphs, *Combinatorica* 1 (1981), 309–311.

Richard Weiss, Presentations for (G, s)-transitive graphs of small valency, *Math. Proc. C bridge Philos. Soc.* 101 (1987), 7–20.

Construction of regular/chiral maps and polytopes

M.D.E. Conder, Regular maps and hypermaps of Euler characteristic -1 to -200, J. Combin. Theory Ser. B 99 (2009), 455–459.

M. Conder & P. Dobcsányi, Determination of all regular maps of small genus, *J. Combin. Theory Ser. B* 81 (2001), 224–242.

M. Conder, I. Hubard & T. Pisanski, Constructions for chiral polytopes, *J. London Math. Soc. (2)* 77 (2008), 115–129.

M. Conder, I. Hubard & E. O'Reilly-Regueiro) Construction of chiral polytopes of large rank with alternating or symmetric automorphism group, *Advances in Mathematics* 452 (2024), 109819 (22 pages).

M. Conder & P. Potočnik, Census of rotary maps by genus and number of edges, `https://rotarymaps.graphsym.net`.

Möbius inversion

M. Downs, The Möbius function of $\mathrm{PSL}_2(q)$, with application to the maximal normal subgroups of the modular group, *J. London Math. Soc.* 43 (1991), 61–75.

M. Downs, Some enumerations of regular hypermaps with automorphism group isomorphic to $\mathrm{PSL}_2(q)$, *Quart. J. Math. Oxford* **48** (1997), 39–58.

P. Hall, The Eulerian functions of a group. *Quart. J. Math. Oxford* **7** (1936), 134–151.

G.A. Jones and S.A. Silver, Suzuki groups and surfaces *J. London Math. Soc.* **48** (1993), 117–125.

L. Weisner, Abstract theory of inversion of finite series, *Trans. Amer. Math. Soc.* 38 (1935), 474–484.

Other references

W. Bosma, J. Cannon & C. Playoust, The Magma Algebra System I: the user language, *J. Symbolic Comput.* 24 (1997), 235–265.

M. Conder, Combinatorial and computational group-theoretic methods in the study of graphs, maps and polytopes with maximal symmetry, in: *Applications of Group Theory to Combinatorics* (ed. J. Koolen, J.H. Kwak & M.Y. Xu), Taylor & Francis, London, 2008, pp. 1–11.

J. Neubüser, *An elementary introduction to coset-table methods in computational group theory*, Groups – St. Andrews 1981, London Math. Soc. Lecture Note Series, vol. 71, 1982, pp. 1–45.