Diameter, girth and other properties of highly symmetric graphs

Grahame Erskine

Submitted for the degree of
Doctor of Philosophy in Mathematics
The Open University
Milton Keynes, UK

May 2017

Abstract

We consider a number of problems in graph theory, with the unifying theme being the properties of graphs which have a high degree of symmetry.

In the degree-diameter problem, we consider the question of finding asymptotically large graphs of given degree and diameter. We improve a number of the current best published results in the case of Cayley graphs of cyclic, dihedral and general groups. In the degree-diameter problem for mixed graphs, we give a new corrected formula for the Moore bound and show non-existence of mixed Cayley graphs of diameter 2 attaining the Moore bound for a range of open cases.

In the degree-girth problem, we investigate the graphs of Lazebnik, Ustimenko and Woldar which are the best asymptotic family identified to date. We give new information on the automorphism groups of these graphs, and show that they are more highly symmetrical than has been known to date.

We study a related problem in group theory concerning product-free sets in groups, and in particular those groups whose maximal product-free subsets are complete. We take a large step towards a classification of such groups, and find an application to the degree-diameter problem which allows us to improve an asymptotic bound for diameter 2 Cayley graphs of elementary abelian groups.

Finally, we study the problem of graphs embedded on surfaces where the induced map is regular and has an automorphism group in a particular family. We give a complete enumeration of all such maps and study their properties.

Preface

Acknowledgements

Firstly, I thank my principal supervisor Professor Jozef Širáň for support and guidance throughout the PhD journey, and for allowing me the freedom to pursue my own research interests.

The pure mathematics group at the Open University includes a number of staff and students with interests in combinatorics, group theory and related areas. I thank them for useful discussions, suggestions and encouragement at various points: David Bevan, Robert Brignall, Jay Fraser, Nick Gill, Terry Griggs, Rob Lewis, Kathleen Quinn, Ian Short, Jakub Sliacan, James Tuite and Bridget Webb.

Thanks are due also to other members of the department and students who shared lively conversation, OU anecdotes and history over lunch: Rosie Cretney, Vasso Evdoridou, Matthew Jacques, Tim Lowe, David Marti-Pete, Ben Mestel, Toby O'Neil, Phil Rippon, Hayley Ryder, Gwyneth Stallard and Mairi Walker.

Many of the results in this thesis have evolved from collaborations both within the OU and externally. To my co-authors I owe a particular debt of gratitude: Chimere Anabanti, David Bevan, Dominique Buset, Mourad El Amiri, Sarah Hart, Katarína Hriňáková, Rob Lewis, Mirka Miller, Hebert Pérez-Rosés, Jozef Širáň and James Tuite. In particular, I thank Professor Mirka Miller for encouraging my interest in the mixed graph problem. Mirka's untimely passing has robbed us not only of an inspiring and respected colleague, but also of a friend to so many.

Declaration

Unless otherwise stated, the detailed results presented in this thesis represent my own work. In the case of joint work, I have taken the approach that contributions which are wholly or mainly my own are included here without further comment. Those parts of joint papers which were primarily the work of others are summarised here with acknowledgement of the main author, but detailed proofs of results are omitted unless essential to the exposition (these can be found in the cited publications or preprints).

Contents

Chapter 1

Introduction

The subject of graph theory plays an important role in modern mathematics. At the most basic level, it captures the ideas of objects (the vertices of the graph) and relationships between those objects (the edges of the graph). Graphs may be directed or undirected, respectively reflecting unidirectional or biderectional relationships between objects. As we shall see, graphs may even be mixed, including both directed and undirected components.

Graphs are often a natural model in real world applications such as telecommunications networks, road and transport planning, economics and social media. In purely mathematical terms, graph theory has strong links to other areas of combinatorics and algebra.

The problems we may study in graph theory are many and varied, including for example:

- Distance problems – how far apart are vertices and how large can we make graphs while keeping distances small.
- Cycle problems – what cycles are in graphs, how large are they and can we avoid small cycles.
- Drawing problems – can we embed graphs in the plane or on other surfaces without edge crossings.
- Colouring problems – can we colour vertices or edges in particular ways to include or avoid certain patterns.
- Subgraph problems – what subgraphs does a graph contain and what can we tell about a graph from its subgraphs.
- Extremal problems – how large or small can we make graphs while insisting on certain properties.
- Symmetry problems – which permutations of the graph preserve adjacency.

These and many other problems are the subject of significant ongoing research. We

Vertices	Undirected	Directed
1	1	1
2	2	3
3	4	16
4	11	218
5	34	9608
6	156	1540944
7	1044	882033440
8	12346	$\approx 1.8 \times 10^{12}$
9	274668	$\approx 1.3 \times 10^{16}$
10	12005168	$\approx 3.4 \times 10^{20}$
11	1018997864	$\approx 3.3 \times 10^{25}$
12	165091172592	$\approx 1.1 \times 10^{31}$

Table 1.1: Numbers of unlabelled graphs and digraphs

offer the following quotes from Bollobás in the preface to his book *Modern Graph Theory* [14]:

> Graph theory, more than any other branch of mathematics, feeds on problems. There are a great many significant open problems which arise naturally in the subject: many of these are simple to state and look innocent but are proving to be surprisingly hard to resolve. It is no coincidence that Paul Erdős, the greatest problem-poser the world has ever seen, devoted much of his time to graph theory.
>
> "*As long as a branch of science offers an abundance of problems, so long is it alive*", said David Hilbert in his address to the Congress in Paris in 1900. Judged by this criterion, graph theory could hardly be more alive.

We focus here on only a few of these problems. The greater part of the thesis will be concerned with a couple of extremal type problems. The first is the degree-diameter problem, where we seek to find large graphs subject to constraints on the maximum number of edges incident to any vertex and the maximum distance between vertices. The second is the degree-girth problem, where we try to find small graphs with no short cycles, subject to each vertex requiring to be incident to a fixed number of edges.

Both of these problems are very hard to solve in full generality. Part of the issue is that the population of graphs becomes very large, even for relatively small numbers of vertices. Table 1.1 gives the number of unique unlabelled undirected and directed graphs on up to 12 vertices. (Source: OEIS [68] sequences A000088, A000273.)

Clearly these numbers become unmanageable very quickly. One common approach, and our main tactic here, is to focus on graphs which have a high degree of symmetry. Not only does this reduce the number of graphs we have to consider, but it also allows us to bring algebraic techniques to bear on our problems; most often group theory. Groups will play a key role in the discussions in this thesis. We use groups both to

construct graphs (typically Cayley graphs) and to investigate the properties of graphs via their symmetries (automorphism groups).

1.1 Notation and definitions

Before describing the problems at hand it will be useful to set out our notational conventions and to define some common terms.

All our graphs will be finite. To denote a graph we will most commonly use the letter G, but will employ Γ if there is a risk of confusion with groups. We begin by recalling some basic definitions. We consider an undirected graph G to consist of a set $V(G)$ of vertices and a set $E(G)$ of edges. We think of an edge between vertices u, v as a set $\{u, v\}$ and say u is *adjacent* to v. The *order* of a graph is $|V(G)|$ and the *size* of a graph is $|E(G)|$. Unless otherwise indicated, graphs are *simple*, that is to say they contain no loops (edges from a vertex to itself) or multiple edges between the same pair of vertices.

The *degree* or *valency* of a vertex is the number of edges incident to it; since we have no loops or multiple edges this is the same as the number of adjacent vertices. If all vertices in a graph G have the same degree, we say G is *regular*. A *path* of length ℓ in a graph is a sequence of distinct vertices $v_0, v_1, \ldots, v_\ell$ such that v_i is adjacent to v_{i+1} for each $i = 0, 1, \ldots, \ell - 1$. A *walk* of length ℓ is a similar sequence except we do not require vertices or edges to be distinct. Given two vertices $u, v \in V(G)$, the *distance* $\mathrm{dist}(u, v)$ between them is the smallest length of any path from u to v. (We will generally only consider *connected* graphs, in which there exists a path between any pair of distinct vertices.) The diameter $\mathrm{diam}(G)$ of a graph G is the largest distance between any pair of vertices.

A *cycle* of length ℓ in a graph is a sequence of vertices $v_0, v_1, \ldots, v_\ell$ such that all are distinct except $v_0 = v_\ell$ and v_i is adjacent to v_{i+1} for each $i = 0, 1, \ldots, \ell - 1$. The *girth* $\mathrm{girth}(G)$ of a graph G is the length of its shortest cycle, if any. A connected graph with no cycles is called a *tree*.

A permutation of the vertex set of a graph which preserves adjacency is called an *automorphism*. The automorphisms of a graph G form a natural group structure under composition, which we denote by $\mathrm{Aut}(G)$. We therefore consider $\mathrm{Aut}(G)$ as a group of permutations acting on $V(G)$. We adopt the convention of action on the right. If $\mathrm{Aut}(G)$ acts transitively on $V(G)$ we say that G is *vertex-transitive*. If the natural induced action of $Aut(G)$ on $E(G)$ is transitive, we say that G is *edge-transitive*. If this action is transitive on ordered pairs (u, v) of adjacent vertices,

we say G is *arc-transitive*.

We recall some basic notation of finite group theory. Given a group G, its identity element will simply be denoted by 1, or by 1_G if there is a danger of ambiguity. For abelian groups we will most often use additive notation, except that the direct product of any two groups A, B will be denoted by $A \times B$. The cyclic group of order n will generally be denoted by $\mathbb{Z}_n$ and we think of it as the additive group of residue classes modulo n. The dihedral group of order $2n$ is denoted by D_{2n}.

In Chapters 4 and 5 we construct groups via *semidirect products*. We choose here a notation and definition of the semidirect product convenient for our needs. Given two groups G and K and a group homomorphism $\varphi : K \to \mathrm{Aut}(G)$, the semidirect product $G \rtimes_\varphi K$ is the group with element set the Cartesian product $G \times K$ and multiplication defined by:

$$(g_1, k_1)(g_2, k_2) = (g_1^{\varphi(k_2)} g_2, k_1 k_2)$$

where the superscript on g_1 indicates the image of g_1 under the automorphism $\varphi(k_2)$ of G.

We shall mainly be concerned with graphs which have large or interesting automorphism groups. One such class of graphs is *Cayley* graphs, defined as follows. Given a finite group G and an inverse-closed subset S of $G \setminus \{1\}$, we consider the vertex set of the graph to be the elements of G, with an edge between vertices g, h if and only if $g^{-1}h \in S$. The resulting graph will be denoted $\mathrm{Cay}(G, S)$.

It is immediate from the definition that $\mathrm{Cay}(G, S)$ is a simple graph of order $|G|$, and is regular of degree $|S|$. Since S is inverse-closed, the adjacency relation is symmetric so that $\mathrm{Cay}(G, S)$ is an undirected graph. It is easy to see that $\mathrm{Cay}(G, S)$ is vertex-transitive (for any $x \in G$ the map $g \mapsto xg$ is an automorphism of the graph). It is easy to see that the distance in the graph $\mathrm{Cay}(G, S)$ from the vertex 1_G to an arbitrary vertex g is simply the minimum number of elements of S which we need to multiply together to obtain g. Since the Cayley graph is vertex-transitive, this implies that its diameter is equal to the largest such number for all elements $g \in G$. This translation between the geometric property (diameter) of a Cayley graph and properties of the underlying group will be crucial in our discussions.

An example of a Cayley graph for the group $G = \mathbb{Z}_{21}$ and set $S = \{\pm 1, \pm 4, \pm 11\}$ is shown in Figure 1.1. This graph has diameter 2 since every element in $\mathbb{Z}_{21}$ can be formed by the addition of 0, 1 or 2 elements of S.

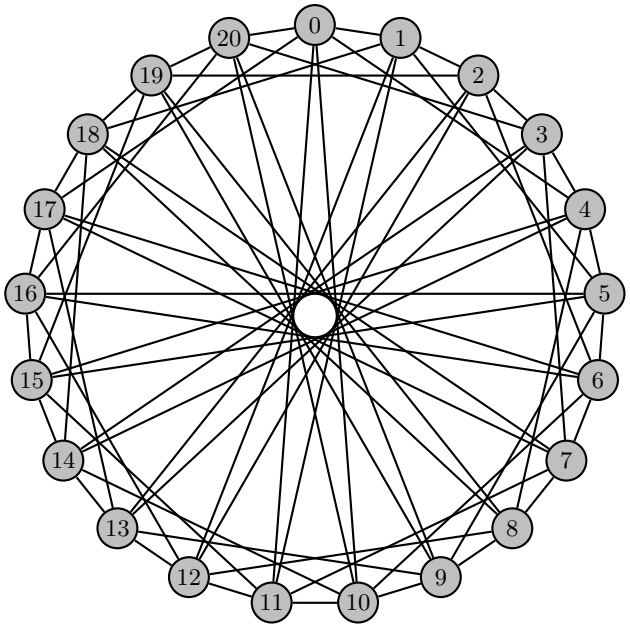

Figure 1.1: The Cayley graph $\mathrm{Cay}(\mathbb{Z}_{21}, \{\pm 1, \pm 4, \pm 11\})$

In the case of directed graphs, we use the same notation as far as possible. In place of an edge set $E(G)$ we have an arc set $A(G)$ with each arc considered as an ordered pair (u, v) of vertices. We must distinguish between the *out-degree* (number of leaving arcs) and *in-degree* (number of entering arcs) of a vertex. A digraph is *out-regular* (resp. *in-regular*) if all its vertices have the same out-degree (resp. in-degree). A digraph which is both out-regular and in-regular is called *diregular*.

Definitions of path, walk, cycle, distance and diameter are similar to the undirected case, but respecting directions of arcs. We may define Cayley graphs of digraphs in exactly the same way as for undirected graphs, except that since we do not require the adjacency relation to be symmetric we drop the condition that the set S must be inverse-closed.

In the case of mixed graphs, we consider the graph to have both undirected edges and directed arcs. All the definitions above carry through in the natural way. A mixed graph whose undirected subgraph is regular and whose directed subgraph is diregular will be called *totally regular*.

We will be very much concerned with parameters such as the order, diameter, degree and girth of graphs. As notation conventions differ between authors, we set out here our usual conventions for how we denote these parameters.

- The letter n will denote the order of a graph.
- The letter k will denote the diameter of a graph.

- The letter g will denote the girth of a graph.
- For regular graphs, the letter d will denote the degree of any vertex.
- For out-regular digraphs, the letter d will denote the out-degree of any vertex.
- For totally regular mixed graphs, the letter r will denote the undirected degree of any vertex and z its directed out-degree.

1.2 The problems

1.2.1 The degree-diameter problem

The degree-diameter problem has its roots in the efficient design of interconnection networks. We try to find the maximum possible number of vertices in a graph where we constrain both the largest degree d of any vertex and the diameter k of the graph. In a communication network, we may think of this as the problem of maximising the number of nodes in the network. Our constraints model the maximum number of interconnections which a single node may have, and the desired maximum number of "hops" required for any two nodes to be able to communicate. Our connections may be bidirectional (in which case we study the problem for undirected graphs), unidirectional (digraphs) or a mixture (mixed graphs).

Our typical approach is to consider the degree-problem restricted to certain families of graphs. Because we concentrate on graphs with a high degree of symmetry, we choose to concentrate for the most part on Cayley graphs.

There are many possible ways to study the degree-diameter problem. One approach is simply to seek the largest possible graph of a given (small) diameter k and degree d. As we have seen though, the combinatorial explosion in the number of graphs of a given order makes this practically impossible except in the smallest cases. Another possibility is to fix the degree d of the graphs under consideration, and investigate how the maximum order of graphs behaves as the diameter k increases. However, we concentrate on tackling the problem in the other direction; that is to say we fix a small diameter k and investigate asymptotic bounds on the largest order of graphs we can construct with a given maximum degree d.

1.2.2 The degree-girth problem

The degree-girth problem is somewhat related to the degree-diameter problem. In this case, we fix a degree d and insist that the minimum degree of vertices in our

graphs should be d, while avoiding cycles shorter than some girth g. We then try to construct graphs with as small an order as possible.

Again, we can tackle the problem in a number of ways. Cayley graphs would be a useful tool here, and indeed there are examples in the literature of this kind of approach. However, in the girth problem it turns out that incidence graphs of finite geometrical structures are a very useful tool, and we will study one such class of graphs in detail. This family of graphs is the best currently known construction in an asymptotic sense, and has a great deal of symmetry reflected in a large and interesting automorphism group which we study.

1.3 Outline

The remainder of the thesis is structured as follows.

In Chapter 2 we give a more detailed account of the degree-diameter problem, including some history and the current best asymptotic results.

In Chapters 3, 4 and 5 we present new asymptotic results in the undirected and directed versions of the degree-diameter problem. Our unifying theme is to use Cayley graphs as a tool to explore constructions of graphs with a high degree of symmetry. However, the results use a variety of techniques and different families of groups in their constructions. Chapter 6 explores the mixed graph version of the problem.

In Chapter 7 we explore the degree-girth problem from the point of view of the best asymptotic family of graphs, which is based on a particular incidence structure. We present new results on the structure and automorphism groups of these graphs.

Chapter 8 explores a related topic in group theory. We may view diameter-constrained Cayley graphs from a group-theoretic perspective as the problem of finding subsets of a group which multiply together to cover the group in an efficient way. An old problem of Street and Whitehead [70] defines a class of groups called *filled* groups and we find a partial classification of all such groups, and make a conjecture on the complete classification. To illustrate the link with the degree-diameter problem, we use the techniques of this chapter to improve the asymptotic bound for diameter two Cayley graphs of elementary abelian 2-groups.

In Chapter 9 we conclude our investigations into the degree-diameter problem by considering arc-transitive graphs.

Chapter 10 departs from the degree-diameter and degree-girth problems to explore

another topic related to highly symmetric graphs. We explore and enumerate the orientably-regular maps having an automorphism group isomorphic to the twisted linear fractional group $M(q^2)$.

Finally, in Chapter 11 we present a revised summary of asymptotic results in the degree-diameter problem, updated to reflect the impact of the constructions described in earlier chapters.

Chapter 2

Background to the degree-diameter problem

2.1 Basic bounds

Recall that our goal is to find the largest possible order of a graph of maximum degree d and diameter k. We deal first with the undirected case. Degree $d = 1$ is a degenerate case so we will assume that $d \geq 2$. A simple counting argument along the following lines yields a natural upper bound called the *Moore bound*. We select an arbitrary vertex u in our graph and construct a spanning tree rooted at u. At distance 1 from u we have a maximum of d vertices. Each of those d vertices has an edge to u, so has $d-1$ other edges available to connect to vertices at distance 2 from u. Thus the maximum possible number of vertices at distance 2 from u is $d(d-1)$. Continuing in this fashion, we see that the Moore bound $M(d,k)$ for graphs of degree d and diameter k is:

$$M(d,k) = \begin{cases} 1 + d\dfrac{(d-1)^k - 1}{d-2} & \text{if } d > 2 \\ 2k+1 & \text{if } d = 2 \end{cases} \tag{2.1}$$

The solid edges of Figure 2.1 illustrate such a tree with parameters $d = 3$ and $k = 2$. So the maximum possible order of such a graph is 10, but this can only be achieved if we can manage to connect the vertices at the lowest level (the dotted edges) in such a

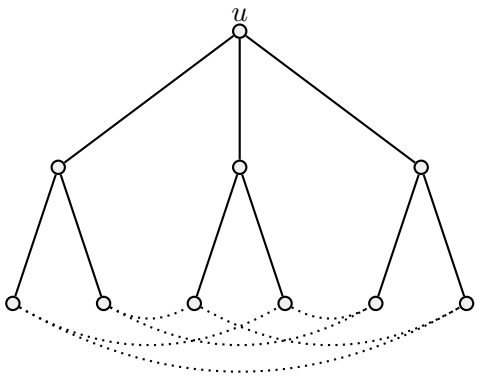

Figure 2.1: The Moore bound for $d = 3, k = 2$

way that the diameter of the whole graph is 2.

It turns out that apart from some trivial examples, this is hardly ever possible and so graphs achieving this bound are exceedingly rare. It is easy to see that at diameter $k = 1$, the bound is achieved for all $d \geq 2$ by the complete graph K_{d+1}. At degree $d = 2$, again the cycle graph C_{2k+1} achieves the bound. The results of Hoffman and Singleton [38] and later Bannai and Ito [8] show that the only non-trivial examples occur at diameter 2. The known graphs are the Petersen graph at degree $d = 3$ and the Hoffman-Singleton graph at degree $d = 7$. The remaining possibility is an unknown graph (or graphs) at degree $d = 57$, whose existence or otherwise is among the most famous open problems in the area.

The classic paper of Hoffman and Singleton [38] uses analysis of the eigenvalues of the adjacency matrix of a Moore graph of diameter two to show that the only possible non-trivial degrees are 3, 7 and 57. This argument was one of the founding papers of the topic of algebraic graph theory.

For directed graphs, we adopt a very similar counting technique based on a spanning tree. In this case it makes sense to consider also the case $d = 1$. This time our vertices at level 1 and beyond have all d out-arcs available to connect to vertices at the next level, and so the Moore bound is:

$$M(d,k) = \begin{cases} 1 + \dfrac{d^{k+1}-1}{d-1} & \text{if } d > 1 \\ k+1 & \text{if } d = 1 \end{cases} \tag{2.2}$$

It turns out that Moore digraphs exist only in the trivial cases when $d = 1$ (directed cycles) or $k = 1$ (complete digraphs). This was first proved by Plesník and Znám [61] in 1974 and independently in 1980 by Bridges and Toueg [16] with an elegant argument, again based on eigenvalues of the adjacency matrix.

A Moore bound can also be defined in a similar way for mixed graphs, where we allow both undirected and directed edges in the graph. The study of Moore graphs in the mixed case was initiated by Bosák [15] in 1979. Bosák gave strong numerical conditions on the parameters for which a mixed graph attaining the bound can exist, but there are very many cases for which the existence of graphs attaining the bound is unknown. Little progress was made on the problem until 2007 when Nguyen, Miller and Gimbert [60, 59] showed that no non-trivial mixed Moore graph of diameter greater than 3 can exist, and gave a formula for the Moore bound in the general case. Much more recently, a number of authors [41, 50] have begun to tackle the problem via computational techniques. We will return to the problem of mixed graphs in

Chapter 6.

Much research in the degree-diameter problem is focused on trying to construct graphs which approach the Moore bound in either an absolute or asymptotic sense. A complete summary of the history and current status of this research is contained in the survey by Miller and Širáň [57], and we mention below only a few of the results most pertinent to our investigations.

2.2 Graphs close to the Moore bound

Since Moore graphs are so rare, it is natural to ask about the existence of graphs which very nearly attain the bound. In this context we speak about the *defect* δ of a graph, and define it to be the shortfall in the order of the graph compared to the relevant Moore bound. So if a graph Γ has maximum degree d and diameter k, then $\delta = M(d,k) - |V(\Gamma)|$.

In the undirected case, it is known that no graphs of defect 1 exist apart from the trivial case of the 4-cycle C_4. For defect 2, five non-trivial graphs are known and various authors have made progress towards placing conditions on the existence of further examples. Not much is known about the situation for larger defects.

For digraphs, defect 1 is attained for diameter 2 in the case of line graphs of complete digraphs, and there are no defect 1 digraphs of diameters 3 or 4. The position for large diameters is unknown.

The survey [57] contains much more detail and references about the above. While further results in this direction would be one way to proceed, we choose instead to focus on the asymptotic version of the problem.

2.3 The asymptotic problem

Our main method of attack in the following chapters will be to try to find families of graphs with good asymptotic properties in the degree-diameter problem. We begin with some definitions and notation.

2.3.1 Definitions

Let $\mathcal{C}$ be a class of graphs. (Typically, we might select $\mathcal{C}$ to be the set of Cayley graphs of a particular class of groups, or the set of graphs with some particular property such as vertex transitivity.)

We define $n_{\mathcal{C}}(d,k)$ to be the largest possible order of a graph in $\mathcal{C}$ with diameter k and maximum degree d. Typically of course, we have no idea what the exact value of $n_{\mathcal{C}}(d,k)$ will be, so our strategy will be to try to find an infinite family $\mathcal{G}$ of graphs within $\mathcal{C}$ which has asymptotically "large" order as we increase d or k. We then use $n_{\mathcal{G}}(d,k)$ as a lower bound on $n_{\mathcal{C}}(d,k)$.

For the most part, we will concentrate on the case where we fix a particular diameter k of interest, and a family $\mathcal{G}$ where our graphs have diameter k but we let the maximum degree d grow as large as we please. From that point of view, the Moore bound (for both the undirected and directed cases) can be expressed as:

$$M(d,k) = d^k + O(d^{k-1}).$$

Thus to measure the usefulness of our family of graphs we define the following two quantities:

$$L^+_{\mathcal{G}}(k) = \limsup_{d\to\infty} \frac{n_{\mathcal{G}}(d,k)}{d^k}\,; \qquad L^-_{\mathcal{G}}(k) = \liminf_{d\to\infty} \frac{n_{\mathcal{G}}(d,k)}{d^k}$$

Our goal will usually be to find lower bounds for these values using some particular construction. Loosely speaking, to bound L^+ we seek a family $\mathcal{G}$ such that for an infinite number of values of d we can show that $n_{\mathcal{G}}(d,k) \geq Kd^k + o(d^k)$ for some constant K. To bound L^-, we do the same thing except we require that the construction must be valid for *all* sufficiently large values of d.

In the following, where the class or family of graphs under consideration is clear from the context we will often omit the subscripts $\mathcal{C}$, $\mathcal{G}$ from the notation.

2.4 Existing asymptotic results

In the remainder of this section we set the context for the following few chapters by reviewing the best asymptotic results in the literature to date. Since the majority of these results relate to the undirected version of the problem, we will focus on that case.

2.4.1 General graphs

In this case we take $\mathcal{G}$ to be the most general family possible: the set of all undirected graphs.

At diameter 2 we have the construction of Brown [17]. Let q be any prime power and

let P be the set of points of $PG(2,q)$, which we view as the set of 1-dimensional subspaces of $GF(q)^3$. We define the graph $B(q)$ to have vertex set P, with adjacency defined by orthogonality of (representative non-zero vectors of) the corresponding subspaces. It is easy to see that $B(q)$ has diameter 2, order q^2+q+1 and maximum degree $d = q+1$. This means that $L^+(2) = 1$ in this case.

Brown's construction is only directly applicable for degrees d of the form $q+1$ where q is a prime power. However, it was shown by Šiagiová, Širáň and Ždimalová in [66] that this can be extended to all degrees by using arguments from analytic number theory on the distribution of prime numbers. So in fact $L^-(2) = 1$. We note that this extension from degrees based on prime numbers to all possible degrees is a useful technique, and we will return to it several times.

Delorme [26] gives a construction showing that $L^+(3) = L^+(5) = 1$ and quotes other references to show $L^-(3) \geq 8/27$, $L^-(4) \geq 3/16$ and $L^-(5) \geq 4^4/5^5$. In a different paper [27], Delorme shows $L^+(4) \geq 1/4$.

For general diameter k, the survey [57] indicates that the best known bound for $L^-(k)$ is 2^{-k} using De Bruijn or Kautz graphs. (We describe the directed version of Kautz graphs in the diameter 2 case in Section 6.2.) We have $L^+(k) \geq 1.6^{-k}$ from a paper by Canale and Gómez [21].

2.4.2 Cayley graphs

We note here the best results available for Cayley graphs, without restriction on the families of groups considered. At diameter 2, a construction of Šiagiová and Širáň [65] yields, for infinitely many degrees d, a Cayley graph of asymptotic order d^2 so that $L^+(2) = 1$. More recently, the same authors with Bachratý [5] obtained a similar result at diameter 3 so that $L^+(3) = 1$.

No equivalent results are available for larger diameters. Using the available results for diameters 2 and 3, it is straightforward to find a direct product construction yielding $L^+(4) \geq \frac{1}{16}$ and $L^+(5) \geq \frac{1}{32}$, but these are likely to be poor bounds.

The optimal constructions above for diameters 2 and 3 are only valid for a very sparse set of degrees, so cannot as they stand be extended to all degrees to provide a bound on L^-. In the L^- version of the problem, a recent result of Abas [2] yields $L^-(2) \geq 0.684$. The best available results for diameters 3 to 5 come from Vetrík [73] who shows that $L^-(3) \geq \frac{3}{16}$, $L^-(4) \geq \frac{32}{625}$ and $L^-(5) \geq \frac{25}{1024}$.

We are also interested in results valid for arbitrary diameter k. The best we currently

have is $L^+(k) \geq L^-(k) \geq \frac{k}{3^k}$ from Macbeth, Šiagiová, Širáň and Vetrík [54].

We consider the problem of general Cayley graphs in Chapter 5.

2.4.2.1 Circulant graphs

The most obvious starting point for an analysis of groups of a particular family is cyclic groups. Cayley graphs of cyclic groups are sometimes called *circulant* graphs, and we will use both terms. At diameter 2, we have a trivial upper bound on $L^+(2)$ of $1/2$ since in an abelian group, commutativity of the generators leads to duplication of paths in the Cayley graph. In fact, for diameter k it is not hard to see (see for example Dougherty and Faber [28]) that $L^+(k) \leq 1/k!$. Vetrík [74] has a diameter 2 construction which gives $L^+(2) \geq \frac{13}{36} \approx 0.36111$. A trivial lower bound on $L^-(2)$ is $1/4$ (see Lemma 4.2 for an explanation of this).

Specific results for diameters 3 and above do not currently appear in the literature. A lower bound on $L^-(k)$ for arbitrary diameter $k \geq 2$ of $1/k^k$ can be deduced in a similar way to the diameter 2 limit of $1/4$ above.

We consider the circulant graph version of the problem in Chapter 3.

2.4.2.2 General abelian groups

The next natural family to consider after cyclic groups is the general abelian case. At diameter 2, the best result for a long time was $L^+(2) \geq \frac{3}{8}$ from Macbeth, Šiagiová and Širáň [53], extended to L^- using the prime gaps method in [66]. However we now have the very recent results in Pott and Zhou [62] giving $L^-(2) \geq \frac{25}{64}$ and $L^+(2) \geq \frac{4}{9}$. At diameter 3 we have $L^+(3) \geq \frac{9}{128}$ from Vetrík [74], again extended to L^- by [66].

There is nothing in the literature beyond diameter 3, so the best bounds are determined by those for cyclic or elementary abelian groups. We discuss abelian groups further in Chapter 8.

2.4.3 Vertex-transitive graphs

For L^+ at diameters 2 and 3 we simply note the Cayley graph results above, giving a limit of 1. For larger diameters, good results are provided (by suppressing directions on Faber-Moore-Chen graphs) by Macbeth, Šiagiová, Širáň and Vetrík in [54]. These graphs have order $\frac{((d+3)/2)!}{((d+3)/2-k)!}$ yielding $L^+(k) \geq 1/2^k$. Unfortunately this construction is only valid for odd degrees and there seems to be no obvious way to

extend it to a result on L^-. However, this result is interesting because the graphs are in general not Cayley graphs, and provide a better bound than the Cayley graph bound noted above.

For $L^-(2)$, we can again use the results for Cayley graphs. In the non-Cayley case, good candidates would be the graphs of McKay, Miller and Širáň [56]. Unfortunately, these are only valid for degrees related to prime powers in certain congruence classes and there is no obvious way to add edges to cover all possible d while maintaining transitivity. So at present, there are no better asymptotic results than those for Cayley graphs.

Our discussions of this version of the problem in the following chapters will focus on the case of Cayley graphs.

2.4.4 Arc-transitive graphs

This class of graphs has received very little attention in the diameter problem, with the first result in Zhou's paper [76] giving $n(d, 2) \geq d^{5/3} + O(d)$ for infinitely many d. Unfortunately, the exponent of 5/3 is smaller than 2, so this result does not currently provide a useful bound on L^+ in this category. We return to this problem in Chapter 9.

2.4.5 Summary table

Table 2.1 collects the results of the above discussion into a single reference.

Type		**Diam 2**	**Diam 3**	**Diam 4**	**Diam 5**	**Diam k**
General graphs						
All graphs	L^-	1.00000	0.29629	0.18750	0.08192	$1/2^k$
	L^+	1.00000	1.00000	0.25000	1.00000	$1/1.6^k$
Vertex-transitive	L^-	0.68400	0.18750	0.05120	0.02441	$k/3^k$
	L^+	1.00000	1.00000	0.06250	0.03125	$1/2^k$
Arc-transitive	L^-	—	—	—	—	—
	L^+	—	—	—	—	—
Cayley graphs						
All groups	L^-	0.68400	0.18750	0.05120	0.02441	$k/3^k$
	L^+	1.00000	1.00000	0.06250	0.03125	$k/3^k$
Cyclic	L^-	0.25000	0.03703	0.00390	0.00032	$1/k^k$
	L^+	0.36111	0.03703	0.00390	0.00032	$1/k^k$
General abelian	L^-	0.39062	0.07031	0.00390	0.00032	$1/k^k$
	L^+	0.44444	0.07031	0.00390	0.00032	$1/k^k$

Table 2.1: Asymptotic lower bounds on orders of undirected graphs

Chapter 3

The degree-diameter problem for circulant graphs

3.1 Introduction

If Cayley graphs are a natural way to study the degree-diameter problem, then in some sense Cayley graphs of cyclic groups are an obvious starting point for that study. Cayley graphs of cyclic groups are often called *circulant* graphs, and we will use these terms interchangeably. We begin our exploration of circulant graphs with some specific notation and definitions.

All our groups in this chapter will be abelian (indeed cyclic) and so we use additive notation for the group operation. We are interested as usual in the largest graph we can construct of given degree and diameter, and we will use the following notation:

- $CC(d,k)$ is the largest order of an undirected circulant graph with degree d and diameter k.
- $DCC(d,k)$ is the largest order of a directed circulant graph with degree d and diameter k.

For a given diameter k, we are interested in determining the asymptotics of $CC(d,k)$ and $DCC(d,k)$ as the degree d tends to infinity. We make use of the following limits as introduced in Chapter 2:

- $L_C^-(k) = \liminf_{d\to\infty} CC(d,k)/d^k$; $L_C^+(k) = \limsup_{d\to\infty} CC(d,k)/d^k$.
- $L_D^-(k) = \liminf_{d\to\infty} DCC(d,k)/d^k$; $L_D^+(k) = \limsup_{d\to\infty} DCC(d,k)/d^k$.

We begin with some trivial bounds on L^- and L^+. The following asymptotic upper bound is easily obtained; see for example the survey paper [57]:

Observation 3.1 (Trivial upper bound)**.** $L_C^+(k) \le L_D^+(k) \le \frac{1}{k!}$.

For a lower bound, consider $\mathbb{Z}_{r^k}$ with generators $\{hr^\ell : 0 < |h| \le \lfloor \frac{r}{2} \rfloor, 0 \le \ell < k\}$:

Observation 3.2 (Trivial lower bound)**.** $L_D^-(k) \ge L_C^-(k) \ge \frac{1}{k^k}$.

For larger diameters, the trivial bounds become numerically small, and the ratio between the upper and lower bound becomes arbitrarily large. Therefore, in order more easily to assess the success of our constructions, we make use of the following measure which records improvement over the trivial lower bound.

Let $R_C^-(k) = kL_C^-(k)^{1/k}$, and define $R_C^+(k)$, $R_D^-(k)$ and $R_D^+(k)$ analogously. Thus, $R_C^-(k) \geq 1$, with equality if the trivial lower bound is approached asymptotically for large degrees. For each k, these R values thus provide a useful indication of the success of our constructions in exceeding the trivial lower bound. In Section 3.5, we show how to construct a cyclic Cayley graph from two smaller ones in such a way that the R values are preserved.

The R values are bounded above by $R_{\max}(k) = k(k!)^{-1/k}$. Using the asymptotic version of Stirling's approximation, $\log k! \sim k \log k - k$, we see that as the diameter tends to infinity,

$$1 \;\leq\; \liminf_{k\to\infty} R_C^-(k) \;\leq\; \liminf_{k\to\infty} R_C^+(k) \;\leq\; e,$$

and similarly for $R_D^-(k)$ and $R_D^+(k)$.

The structure of most of this chapter follows closely our joint paper with Bevan and Lewis [11]. However we begin with an introductory section containing some crucial lemmas which we need both in this chapter and in subsequent chapters, and which allow us in some circumstances to extend a result for $L^+(k)$ to $L^-(k)$.

In the following section, we use these lemmas to extend a result of Vetrík [74] to deduce new lower bounds for $L_C^-(2)$ and $R_C^-(2)$. In Section 3.4, we describe a direct product construction and use it to build large circulant graphs of small diameter and arbitrarily large degree. We also prove that this construction is unable to yield values that exceed the trivial lower bound for large diameter. However, in Section 3.5, we demonstrate a method of building a circulant graph from two smaller ones, and show how the application of this method to the constructions from Section 3.4 enables us to exceed the trivial lower bound for every diameter.

Section 3.6 contains an application of our constructions to obtain upper bounds on the minimum size of a set $A \subseteq \mathbb{Z}_n$ such that the k-fold sumset kA is equal to $\mathbb{Z}_n$. We conclude, in Section 3.7, by presenting a revised table of the largest known circulant graphs of small degree and diameter, including a number of new largest orders.

3.2 Preliminary results

In this and subsequent chapters, we will encounter Cayley graph constructions in the degree-diameter problem based on finite fields, which means we can only directly use the construction for degrees which are related to some prime power. Examples of this type of construction will be found in Theorems 4.5 and 5.4. Other examples in the literature are Šiagiová, Širáň and Ždimalová [66] and Vetrík [73], amongst others.

While these constructions yield graphs which are valid for an infinite number of degrees and hence can be used to obtain a lower bound on $L^+(k)$, we would ideally like to extend the validity to all degrees and hence obtain a bound on $L^-(k)$. Our strategy is to use results from analytic number theory on the distribution of prime numbers to prove that for all sufficiently large degrees d, we can find a prime number such that we can build a graph $\mathrm{Cay}(G, S)$ of degree $d' \leq d$ using our chosen construction. We then add $d - d'$ generators to our set S yielding a graph of the same order, no larger diameter and degree d. The method hinges on being able to find a prime p such that $d - d'$ is small enough not to affect the asymptotic value of the result.

The method was first used by Šiagiová, Širáň and Ždimalová [66]. Because this is such a useful technique we give here a general version of their idea in the form of a lemma.

Lemma 3.3. *Let $\mathcal{G}$ be a family of groups. Let $k \geq 2$ and suppose that there exists some N such that for all primes $p \geq N$, we can find a group $G(p) \in \mathcal{G}$ and an inverse-closed subset $S(p) \subseteq G(p)$ such that $\mathrm{Cay}(G(p), S(p))$ has diameter k. Suppose further that there exists a positive constant C, D such that as $p \to \infty$, $|G(p)| = Cp^k + o(p^k)$, $|S(p)| = Dp + o(p)$ and that for all p, $G(p) \setminus S(p)$ contains at least one involution.*

Then in the class of Cayley graphs of the group family $\mathcal{G}$, $L^-_{\mathcal{G}}(k) \geq \dfrac{C}{D^k}$.

Proof. It suffices to show that for any sufficiently large degree d, we can find a Cayley graph of a group in $\mathcal{G}$ with degree d, diameter k and order $\frac{C}{D^k}d^k + o(d^k)$. Let d be a degree large enough so that there exists a prime p such that we can find a group $G(p)$ and a set $S(p)$ satisfying the conditions. We choose p to be the largest such prime so that $|S(p)| \leq d$. We now add any inverse-closed set of size $d - |S(p)|$ chosen from $G(p) \setminus S(p)$ to our generating set to obtain a new generating set $S'(p)$. Note that we can always do this because if we need to add an odd number of generators, we have an involution in $G(p) \setminus S(p)$.

Let $d' = |S(p)|$. Then $d' = Dp + o(p)$. Now we use the result of Baker, Harman and

Pintz [7] which states that for sufficiently large x, we are guaranteed a prime in the interval $(x, x + x^\theta]$ where $\theta = 0.525$. This means that $p = \frac{1}{D}d' + o(d') = \frac{1}{D}d + o(d)$. Then $\mathrm{Cay}(G(p), S'(p))$ has the required properties. □

Lemma 3.3 is applicable for constructions where we are free to choose any sufficiently large prime p. However, in Section 3.3 we will encounter a construction in which we are able to select p only from a restricted set of congruence classes modulo 13. To handle this situation, we now derive a more general result which in some circumstances allows us still to move from L^+ to L^- without reduction in the asymptotic value. In what follows we use the usual notation $\pi(x)$ to mean the number of primes not exceeding x. We use ϕ for Euler's totient function, so that for $n \geq 2$, $\phi(n)$ is the number of positive integers less than n and coprime to it.

Our basic tool is a strong version of the Brun-Titchmarsh theorem, which was proved by Montgomery and Vaughan and can be stated in the following form.

Lemma 3.4. *[58, Theorem 2] Given a non-trivial congruence class* $C = \{a + nq : n \in \mathbb{Z}\}$ *with* $\gcd(a, q) = 1$ *we denote by* $\pi(x; C)$ *the number of primes not exceeding* x *which are in the class* C*. Then for any function* f *with* $f(x) \to \infty$ *as* $x \to \infty$*, and for sufficiently large* x*:*

$$\pi(x + f(x); C) - \pi(x; C) \leq \frac{2f(x)}{\phi(q) \log(f(x)/q)}$$

We use this result to derive a further lemma in a form more suitable for our needs.

Lemma 3.5. *Let* $q > 1$ *and let* a *be a positive integer with* $\gcd(a, q) = 1$*. Let* θ *be a positive real number. Then as* $x \to \infty$*, an asymptotic upper bound on the number of primes in the interval* $(x, x + x^\theta]$ *which are congruent to* $a \mod q$ *is:*

$$\frac{2}{\theta\phi(q)} \frac{x^\theta}{\log x}$$

Proof. Take $f(x) = x^\theta$ in Lemma 3.4. □

What we would like to do now is to find a lower bound on *all* the primes in such an interval. To do this we revisit the paper of Baker, Harman and Pintz [7] which we used in the proof of Lemma 3.3. In that proof we used only a rather weak version of their result. What they actually proved is that for sufficiently large x,

$$\pi(x + x^\theta) - \pi(x) \geq \frac{9}{100} \frac{x^\theta}{\log x} \text{ where } \theta = 0.525$$

So if $\phi(q) \geq \frac{200}{9\theta} \approx 42.3$, for sufficiently large x we are guaranteed a prime p in the interval $(x, x + x^\theta]$ which lies outside any single congruence class $\bmod q$, for example $p \not\equiv 1 \bmod q$. This is exactly the kind of result we want, but unfortunately we need it to work for smaller values of q for our diameter 2 construction.

The idea now is that $\theta = 0.525$ gives a tighter interval than we really need. In fact for our argument about adding edges to graph constructions in the proof of Lemma 3.3, we only really need $\theta < 1$. The hope would be that we might find a bound of the form

$$\pi(x + x^\theta) - \pi(x) \geq K \frac{x^\theta}{\log x}$$

for some $\theta > 0.525$ and $K > \frac{9}{100}$. To do this we turn to the predecessor paper by Baker and Harman [6, Eq(2)], which says that for sufficiently large x,

$$\pi(x) - \pi(x - y) \geq \frac{2y}{5 \log x} \text{ for } y \geq x^{0.54}$$

This lets us use $K = \frac{2}{5}$, $\theta = 0.54$. Using these values we can make the argument work for any q with $\phi(q) > 9.3$. Thus using exactly the same argument as in the proof of Lemma 3.3, we have the following lemma which extends that result to the case where we exclude a single congruence class.

Lemma 3.6. *Let q be a positive integer with $\phi(q) \geq 10$. Let a be any integer coprime to q. Let $\mathcal{G}$ be a family of groups. Let $k \geq 2$ and suppose that there exists some N such that for all primes $p \geq N$ such that $p \not\equiv a \pmod q$, we can find a group $G(p) \in \mathcal{G}$ and an inverse-closed subset $S(p) \subseteq G(p)$ such that $\mathrm{Cay}(G(p), S(p))$ has diameter k. Suppose further that there exists a positive constant C, D such that as $p \to \infty$, $|G(p)| = Cp^k + o(p^k)$, $|S(p)| = Dp + o(p))$ and that for all p, $G(p) \setminus S(p)$ contains at least one involution.*

Then in the class of Cayley graphs, $L^-_{\mathcal{G}}(k) \geq \frac{C}{D^k}$.

Ideally we would try to extend this technique further to be valid for even smaller q. However there is an inherent limitation in this method which means that we must always have $\phi(q) > 2$. To see why, note that the combination of the inequalities means that we need $\phi(q) > \frac{2}{K\theta}$. We seek the best possible values for K and θ, but we are constrained by $\theta < 1$ to make our construction work. In addition, we must have $K \leq 1$ by the Prime Number Theorem, since the density of primes in the interval from x is close to $1/\log x$. So we can never make this technique work for very small q, for example 3 or 4.

3.3 Bounds for diameter 2 circulant graphs

In common with the degree-diameter problem in general, much of the study to date for the restricted circulant graph problem has concentrated on the diameter 2 undirected case. In this instance, the trivial lower bound on $L^-(2)$ is 1/4 and the trivial upper bound on $L^+(2)$ is 1/2.

The best published asymptotic result to date for circulant graphs of diameter 2 is by Vetrík [74] (building on Macbeth, Šiagiová and Širáň [53]) who presents a construction which proves that $L_C^+(2) \geq \frac{13}{36} \approx 0.36111$.

Vetrík's result [74] is valid for all degrees of the form $6p-2$ where p is a prime such that $p > 14$, $p \not\equiv 1 \pmod{13}$. The construction satisfies all the conditions of Lemma 3.6, and so we can directly apply the lemma to obtain our first result on circulant graphs as follows.

Theorem 3.7. *In the class of circulant graphs,*

$$L_C^-(2) \geq \frac{13}{36} \approx 0.36111$$

3.4 A new direct product construction

In this section, we construct large undirected circulant graphs of diameters $k = 3, 4, 5$ and arbitrary large degree. We also construct large directed circulant graphs of diameters $k = 2, \ldots, 9$ and arbitrary large degree. We then prove that the approach used is insufficient to yield values that exceed the trivial lower bound for large diameter.

The diameter 2 constructions of Vetrík [74] and earlier similar ideas from other authors construct cyclic groups of the form $F_p^+ \times F_p^* \times \mathbb{Z}_w$ for some fixed w and variable p, where F_p^+ and F_p^* are the additive and multiplicative groups of the Galois field $GF(p)$. Thus the first two components of their constructions are very tightly coupled, and this coupling is a key to their success. However, a significant limitation of this method is that it is only applicable in the diameter 2 case.

In contrast, the constructions considered here have components that are as loosely coupled as possible. For diameter k, they have the form $\mathbb{Z}_{r_1} \times \mathbb{Z}_{r_2} \times \ldots \times \mathbb{Z}_{r_k} \times \mathbb{Z}_w$ for some fixed w and variable pairwise coprime r_i. This gives us greater flexibility, especially in terms of the diameters we can achieve. The price for this is that we lose the inherent structure of the finite field, which consequently places limits on the bounds we can achieve.

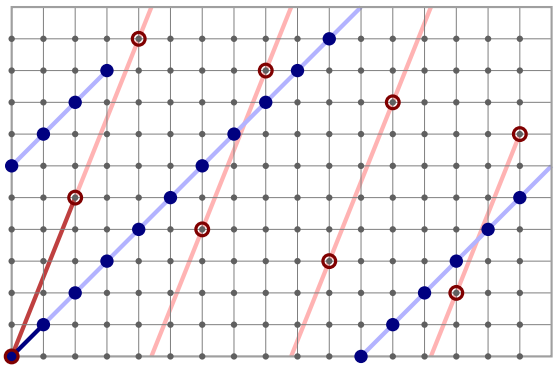

Figure 3.1: Every element of $\mathbb{Z}_{17} \times \mathbb{Z}_{11}$ is the sum of one of the 21 solid elements and one of the 9 circled elements.

The constructions in this section make use of the following result concerning the representation of each element of the cyclic group $\mathbb{T} = \mathbb{Z}_r \times \mathbb{Z}_s$ (r and s coprime) as the sum of a small multiple of the element $(1,1)$ and a small multiple of another element (u,v). It can be helpful to envisage $\mathbb{T}$ as a group of vectors on the $r \times s$ discrete torus.

Lemma 3.8. *Let u, d, s and m be positive integers with $s > 1$ and coprime to md. Let $v = u + d$. Suppose $s \geq mv(u-1)$. Then, for every element (x,y) of $\mathbb{T} = \mathbb{Z}_{s+md} \times \mathbb{Z}_s$, there exist nonnegative integers $h < s + mv$ and $\ell < s - m(u-1)$ such that $(x,y) = h(1,1) + \ell(u,v)$.*

Observe that the construction ensures that $(s+mv)(1,1) = m(u,v)$. Figure 3.1 illustrates the case with parameters $u = 2$, $v = 5$, $s = 11$, $m = 2$.

Proof (Bevan). Let $t = s - m(u-1)$. Since s is coprime to md, $(1,1)$ generates $\mathbb{T}$. Hence, it suffices to show that, in the list $(0,0)$, $(1,1)$, $(2,2)$, ..., the gaps between members of $\{\ell(u,v) : 0 \leq \ell < t\}$ are not "too large".

Specifically, we need to show that, for each nonnegative $\ell < t$, there is some positive $h' \leq s + mv$ and nonnegative $\ell' < t$ such that $\ell(u,v) + h'(1,1) = \ell'(u,v)$.

There are two cases. If $\ell < t - m$, then we can take $h' = s + mv$ and $\ell' = \ell + m$:

$$\begin{aligned} \ell(u,v) \;+\; (s+mv)(1,1) &= (\ell u + s + mu + md,\ \ell v + s + mv) \\ &= (\ell u \quad + mu, \qquad \ell v \quad + mv) \\ &= (\ell + m)(u,v). \end{aligned}$$

If $\ell \geq t - m$, then we can take $h' = muv$ and $\ell' = \ell + m - t = \ell + mu - s$:

$$\begin{aligned} \ell(u,v) \,+\, muv(1,1) &= (\ell u + mu^2 + mud, \qquad\qquad \ell v + muv) \\ &= (\ell u + mu^2 + mud - u(s+md),\ \ell v + muv - vs) \\ &= (\ell + mu - s)(u,v). \end{aligned}$$

The requirement that $muv \leq s + mv$ is clearly equivalent to the condition on s in the statement of the lemma. □

In our direct product constructions, we make use of Lemma 3.8 as follows:

Lemma 3.9. *Let $\mathbb{T} = \mathbb{Z}_{r_1} \times \mathbb{Z}_{r_2} \times \ldots \times \mathbb{Z}_{r_k}$ such that $r_1 > r_2 > \ldots > r_k$, and for each i, j with $1 \leq i < j \leq k$,*

- *r_i and r_j are coprime,*
- *r_i is coprime with i,*
- *there is a positive integer $m_{i,j}$ such that $r_i - r_j = m_{i,j}(j-i)$ and $r_j \geq m_{i,j}(i-1)j$.*

Let $\mathbf{o} = (1, 1, \ldots, 1)$, $\mathbf{u} = (1, 2, \ldots, k)$ and, for each i, $\mathbf{e}_i = (0, \ldots, 0, 1, 0, \ldots, 0)$ be elements of $\mathbb{T}$, where only the ith coordinate of $\mathbf{e}_i$ is 1, and let the set A consist of these $k+2$ elements.

Let $c_{\mathbf{o}} = \max\limits_{i<j}(r_j + jm_{i,j})$, $c_{\mathbf{u}} = r_1$, and for each i, $c_{\mathbf{e}_i} = r_i$.

Then, for every element $\mathbf{x}$ of $\mathbb{T}$ and every k-element subset S of A, there exist nonnegative integers $h_{\mathbf{s}} < c_{\mathbf{s}}$ for each $\mathbf{s} \in S$, such that $\mathbf{x} = \sum_{\mathbf{s}} h_{\mathbf{s}}\mathbf{s}$.

Proof (Bevan). There are four cases. If S contains neither $\mathbf{o}$ nor $\mathbf{u}$, the result follows trivially.

If S contains $\mathbf{o}$ but not $\mathbf{u}$, omitting $\mathbf{e}_i$, then we can choose $h_{\mathbf{o}}$ to be the ith coordinate of $\mathbf{x}$. Note that, as required, $c_{\mathbf{o}} \geq r_2 + 2(r_1 - r_2) = r_1 + (r_1 - r_2) > r_i$ for all i.

If S contains $\mathbf{u}$ but not $\mathbf{o}$, omitting $\mathbf{e}_i$, then, since i and r_i are coprime, we can choose $h_{\mathbf{u}}$ such that $ih_{\mathbf{u}} \pmod{r_i}$ is the ith coordinate of $\mathbf{x}$.

Finally, if S contains both $\mathbf{o}$ and $\mathbf{u}$, omitting $\mathbf{e}_i$ and $\mathbf{e}_j$, then we can choose $h_{\mathbf{o}}$ and $h_{\mathbf{u}}$ by applying Lemma 3.8 to $\mathbb{Z}_{r_i} \times \mathbb{Z}_{r_j}$ with $(u, v) = (i, j)$. □

3.4.1 Undirected constructions

We can use Lemma 3.9 to construct undirected circulant graphs of any diameter by means of the following theorem.

Theorem 3.10. *Let w and k be positive integers and suppose that there exist sets B and T of positive integers with the following properties:*

- *$B = \{b_1, \ldots, b_{k+2}\}$ has cardinality $k+2$ and the property that every element of $\mathbb{Z}_w$ can be expressed as the sum of exactly k distinct elements of $B \cup -B$, no two of which are inverses.*
- *$T = \{r_1, r_2, \ldots, r_k\}$ has cardinality k and the properties that all its elements are coprime to w, and it satisfies the requirements of Lemma 3.9, i.e. for each $i < j$:*
  - *$r_i > r_j$*
  - *$(r_i, r_j) = 1$*
  - *$(r_i, i) = 1$*
  - *There is a positive integer $m_{i,j}$ such that $r_i - r_j = m_{i,j}(j-i)$ and $r_j \geq m_{i,j}(i-1)j$.*

Let $c_{\mathbf{o}} = \max_{i<j}(r_j + jm_{i,j})$ and $c_{\mathbf{u}} = r_1$ as in Lemma 3.9.

Then there exists an undirected circulant graph of order $w\prod_{i=1}^{k} r_i$, degree at most $2\left(\sum_{i=1}^{k} r_i + c_o + c_u\right)$ and diameter k.

Proof (Bevan). Let $\mathbb{T} = \mathbb{Z}_{r_1} \times \mathbb{Z}_{r_2} \times \ldots \times \mathbb{Z}_{r_k} \times \mathbb{Z}_w$. Then $\mathbb{T}$ is a cyclic group since all its factors have coprime orders.

Let X be the generating set consisting of the following elements.

- $(x, 0, 0, \ldots, 0, \pm b_1)$, $x \in \mathbb{Z}_{r_1}$
- $(0, x, 0, \ldots, 0, \pm b_2)$, $x \in \mathbb{Z}_{r_2}$
  $\vdots$
- $(0, 0, \ldots, 0, x, \pm b_k)$, $x \in \mathbb{Z}_{r_k}$
- $\pm(x, x, \ldots, x, x, b_{k+1})$, $0 \leq x < c_{\mathbf{o}}$
- $\pm(x, 2x, \ldots, (k-1)x, kx, b_{k+2})$, $0 \leq x < c_{\mathbf{u}}$

Then by construction and by Lemma 3.9, every element of $\mathbb{T}$ is the sum of at most k elements of X. Since $|\mathbb{T}| = w\prod_{i=1}^{k} r_i$ and $|X| = 2\left(\sum_{i=1}^{k} r_i + c_o + c_u\right)$, the result follows. □

For small diameters this technique results in the following asymptotic bounds.

Theorem 3.11. *For diameters $k = 3, 4, 5$, we have the following lower bounds on $L_C^-(k)$ and $R_C^-(k)$:*

(a) $L_C^+(3) \geq \frac{57}{1000}$ *and* $L_C^-(3) \geq \frac{7}{125}$, *so* $R_C^+(3) > 1.15455$ *and* $R_C^-(3) > 1.14775$.

(b) $L_C^+(4) \geq L_C^-(4) \geq \frac{25}{3456}$, *so* $R_C^+(4) \geq R_C^-(4) > 1.16654$.

(c) $L_C^+(5) \geq L_C^-(5) \geq \frac{109}{134456}$, *so* $R_C^+(5) \geq R_C^-(5) > 1.20431$.

Proof. Given a diameter k, the strategy is to find an optimal value of w which admits a set B satisfying the conditions of Theorem 3.10. We then seek an infinite family of positive integers q and a set $\Delta = \{\delta_1, \delta_2, \ldots, \delta_{k-1}\}$ such that for each of our values of q, the set $T = \{q, q - \delta_1, \ldots, q - \delta_{k-1}\}$ satisfies the conditions of the theorem. We illustrate for $k = 3$.

To prove *(a)* we take $w = 57$ and $B = \{1, 2, 7, 8, 27\}$. It is easily checked that every element of $\mathbb{Z}_{57}$ is the sum of three distinct elements of $B \cup -B$, no two of which are inverses. Now we let $\Delta = \{4, 6\}$. For any $q \geq 17, q \equiv 5 \pmod 6, q \not\equiv 0, 4, 6 \pmod{19}$ it is straightforward to verify that the set $T = \{q, q - 4, q - 6\}$ satisfies the conditions of Theorem 3.10. In the notation of Lemma 3.9, we have $c_{\mathbf{o}} = q + 4$.

Taking a generating set X as defined in Theorem 3.10 we may construct a circulant graph of diameter 3, degree $d = |X| = 10q - 12$ and order
$57q(q - 4)(q - 6) = \frac{57}{1000}(d + 12)(d - 28)(d - 48)$.

We can do this for an infinite number of values of q, and hence for an infinite number of values of $d = 10q - 12$ we have

$$CC(d, 3) \geq \frac{57}{1000}(d + 12)(d - 28)(d - 48).$$

This yields $L_C^+(3) \geq \frac{57}{1000}$. Now we need to consider $L_C^-(3)$. The strategy will be to try to add "few" edges to our graphs to cover all possible degrees. Observe that we can use this construction for any $q \equiv 17 \pmod{114}$ and hence for any $d \equiv 158 \pmod{1140}$. Given any arbitrary *even* degree d, we can therefore find some d' no smaller than $d - 1140$ for which the construction works. We can then add $d - d'$ generators to our graph to obtain a graph of the same order, degree d and diameter 3.

However our graphs always have odd order, and so we are unable to obtain an odd degree graph by this method. To get round this problem we may use $w = 56$, $B = \{1, 2, 7, 14, 15\}$, $\Delta = \{2, 4\}$ and $c_{\mathbf{o}} = q + 2$. Again it is easy to check that the

relevant conditions are satisfied for any $q \geq 15$ such that $q \equiv 3, 5 \pmod 6$ and $q \equiv 1, 3, 5, 6 \pmod 7$. Then for $d = 10q - 8$ we can construct a graph of order $\frac{7}{25}(d+8)(d-12)(d-32)$, degree d and diameter 3. We can do this for any $q \equiv 15 \pmod{42}$ and hence for any $d \equiv 142 \pmod{420}$. So given any arbitrary degree d, we can therefore find some d' no smaller than $d - 420$ for which the construction works, and then add $d - d'$ generators to our graph to obtain a graph of the same order and diameter 3. (Since our graphs now have even order it is possible to add an odd number of generators.) Since the number of added generators is bounded above (by 419), the order of the graph is $\frac{7}{125}d^3 + O(d^2)$. Result *(a)* for $L_C^-(3)$ follows.

For *(b)* and *(c)* we adopt a similar method. For brevity we show only the relevant sets in the construction, summarised as follows.

(b) $(k = 4)$ – Take $w = 150, B = \{1, 7, 16, 26, 41, 61\}$ and $\Delta = \{6, 8, 12\}$ so $c_{\mathbf{o}} = q + 6$. Then for $q \geq 49, q \equiv 19 \pmod{30}$ and $d = 12q - 40$, we have

$$CC(d, 4) \geq \frac{25}{3456}(d+40)(d-32)(d-56)(d-104).$$

(c) $(k = 5)$ – Take $w = 436, B = \{1, 15, 43, 48, 77, 109, 152\}$ and $\Delta = \{0, 4, 10, 12, 16\}$ so $c_{\mathbf{o}} = q + 8$. Then for $q \geq 77, q \equiv 5 \pmod 6, q \not\equiv 0, 1 \pmod 5, q \not\equiv 0, 4, 10, 12, 16 \pmod{109}$ and $d = 14q - 68$, we have

$$CC(d, 5) \geq \frac{109}{134456}(d+68)(d+12)(d-72)(d-100)(d-156). \qquad \square$$

3.4.2 Directed constructions

An analogous method yields directed circulant graphs via the following theorem.

Theorem 3.12. *Let w and k be positive integers and suppose that there exist sets B and T of positive integers with the following properties:*

- *$B = \{0, b_2, \ldots, b_{k+2}\}$ has cardinality $k + 2$ and the property that every element of $\mathbb{Z}_w$ can be expressed as the sum of exactly k distinct elements of B.*
- *$T = \{r_1, r_2, \ldots, r_k\}$ has cardinality k and the properties that all its elements are coprime to w, and it satisfies the requirements of Lemma 3.9, i.e. for each $i < j$:*
  - *(a) $r_i > r_j$*
  - *(b) $(r_i, r_j) = 1$*
  - *(c) $(r_i, i) = 1$*
  - *(d) There is a positive integer $m_{i,j}$ such that $r_i - r_j = m_{i,j}(j - i)$ and $r_j \geq m_{i,j}(i - 1)j$.*

Let $c_{\mathbf{o}} = \max_{i<j}(r_j + jm_{i,j})$ *and* $c_{\mathbf{u}} = r_1$ *as in Lemma 3.9.*

Then we may construct a directed circulant graph of order $w\prod_{i=1}^{k} r_i$*, degree* $\sum_{i=1}^{k} r_i + c_o + c_u - 1$ *and diameter* k*.*

Proof. Let $\mathbb{T} = \mathbb{Z}_{r_1} \times \mathbb{Z}_{r_2} \times \ldots \times \mathbb{Z}_{r_k} \times \mathbb{Z}_w$. Then $\mathbb{T}$ is a cyclic group since all its factors have coprime orders.

Let X be the generating set consisting of the following elements.

- $(x, 0, 0, \ldots, 0, 0)$, $x \in \mathbb{Z}_{r_1} \setminus \{0\}$
- $(0, x, 0, \ldots, 0, b_2)$, $x \in \mathbb{Z}_{r_2}$
  $\vdots$
- $(0, 0, \ldots, 0, x, b_k)$, $x \in \mathbb{Z}_{r_k}$
- $(x, x, \ldots, x, x, b_{k+1})$, $0 \leq x < c_{\mathbf{o}}$
- $(x, 2x, \ldots, (k-1)x, kx, b_{k+2})$, $0 \leq x < c_{\mathbf{u}}$

Then by construction and by Lemma 3.9, every element of $\mathbb{T}$ is the sum of at most k elements of X. Since $|\mathbb{T}| = w\prod_{i=1}^{k} r_i$ and $|X| = \sum_{i=1}^{k} r_i + c_o + c_u - 1$, the result follows. □

For small diameters this technique results in the following asymptotic bounds.

Theorem 3.13. *For diameters* $k = 2, \ldots, 9$*, we have the following lower bounds on* $L_D^-(k)$ *and* $R_D^-(k)$

(a) $L_D^-(2) \geq \frac{3}{8}$*, so* $R_D^-(2) > 1.22474$*.*

(b) $L_D^-(3) \geq \frac{9}{125}$*, so* $R_D^-(3) > 1.24805$*.*

(c) $L_D^-(4) \geq \frac{13}{1296}$*, so* $R_D^-(4) > 1.26588$*.*

(d) $L_D^-(5) \geq \frac{17}{16807}$*, so* $R_D^-(5) > 1.25881$*.*

(e) $L_D^-(6) \geq \frac{3}{32768}$*, so* $R_D^-(6) > 1.27378$*.*

(f) $L_D^-(7) \geq \frac{10}{1594323}$*, so* $R_D^-(7) > 1.26436$*.*

(g) $L_D^-(8) \geq \frac{9}{25000000}$*, so* $R_D^-(8) > 1.25206$*.*

(h) $L_D^-(9) \geq \frac{42}{2357947691}$*, so* $R_D^-(9) > 1.23939$*.*

Proof. The method is exactly the same as the proof of Theorem 3.11 and we summarise as follows.

(a) $(k = 2)$ – Take $w = 6$, $B = \{0, 1, 2, 4\}$, $\Delta = \{2\}$ so $c_{\mathbf{o}} = q + 2$. Then for $q \geq 7, q \equiv 1 \pmod 6$ and $d = 4q - 1$, we have

$$DCC(d, 2) \geq \frac{3}{8}(d + 1)(d - 7).$$

(b) $(k = 3)$ – Take $w = 9$, $B = \{0, 1, 2, 3, 6\}$, $\Delta = \{4, 6\}$ so $c_{\mathbf{o}} = q + 4$. Then for $q \geq 17, q \equiv 5 \pmod 6$ and $d = 5q - 7$, we have

$$DCC(d, 3) \geq \frac{9}{125}(d + 7)(d - 13)(d - 23).$$

(c) $(k = 4)$ – Take $w = 13$, $B = \{0, 1, 3, 5, 7, 8\}$, $\Delta = \{2, 4, 6\}$ so $c_{\mathbf{o}} = q + 2$. Then for $q \geq 23, q \equiv 5 \pmod 6, q \not\equiv 0, 2, 4, 6 \pmod{13}$ and $d = 6q - 11$, we have

$$DCC(d, 4) \geq \frac{13}{1296}(d + 11)(d - 1)(d - 13)(d - 25).$$

(d) $(k = 5)$ – Take $w = 17$, $B = \{0, 1, 2, 3, 4, 8, 13\}$, $\Delta = \{4, 10, 12, 16\}$ so $c_{\mathbf{o}} = q + 8$. Then for $q \geq 77, q \equiv 5 \pmod 6, q \not\equiv 0, 1 \pmod 5, q \not\equiv 0, 4, 10, 12, 16 \pmod{17}$ and $d = 7q - 35$, we have

$$DCC(d, 5) \geq \frac{17}{16807}(d + 35)(d + 7)(d - 35)(d - 49)(d - 77).$$

(e) $(k = 6)$ – Take $w = 24$, $B = \{0, 1, 2, 4, 8, 13, 18, 22\}$, $\Delta = \{6, 12, 18, 24, 30\}$ so $c_{\mathbf{o}} = q + 6$. Then for $q \geq 181, q \equiv 1, 5 \pmod 6, q \not\equiv 0, 4 \pmod 5$ and $d = 8q - 85$, we have

$$DCC(d, 6) \geq \frac{3}{32768}(d + 85)(d + 37)(d - 11)(d - 59)(d - 107)(d - 155).$$

(f) $(k = 7)$ – Take $w = 30$, $B = \{0, 1, 2, 6, 9, 12, 16, 17, 18\}$, $\Delta = \{0, 2, 6, 18, 20, 30, 42\}$ so $c_{\mathbf{o}} = q + 42$. Then for $q \geq 529, q \equiv 1 \pmod 6, q \equiv 4 \pmod 5, q \not\equiv 0, 2, 6 \pmod 7, q \not\equiv 9 \pmod{11}$ and $d = 9q - 77$, we have

$$DCC(d, 7) \geq \frac{10}{1594323}(d + 77)(d + 59)(d + 23)(d - 85)(d - 103)(d - 193)(d - 301).$$

(g) $(k = 8)$ – Take $w = 36$, $B = \{0, 1, 2, 3, 6, 12, 19, 20, 27, 33\}$, $\Delta = \{0, 6, 12, 18, 24, 30, 36, 42\}$ so $c_{\mathbf{o}} = q + 6$. Then for $q \geq 353, q \equiv 1, 5$

$(\text{mod } 6), q \equiv 3 \ (\text{mod } 5), q \not\equiv 0, 1 \ (\text{mod } 7)$ and $d = 10q - 163$, we have

$$DCC(d, 8) \geq \frac{9}{25000000}(d+163)(d+103)(d+43)(d-17)(d-77)$$
$$(d-137)(d-197)(d-257).$$

(h) $(k = 9)$ – Take $w = 42$, $B = \{0, 1, 2, 3, 4, 9, 16, 20, 26, 30, 37\}$,
$\Delta = \{0, 2, 6, 12, 20, 30, 42, 56, 72\}$ so $c_{\mathbf{o}} = q + 72$. Then for $q \geq 1093, q \equiv 1$
$(\text{mod } 6), q \equiv 3, 4 \ (\text{mod } 5), q \equiv 1, 3, 4 \ (\text{mod } 7), q \not\equiv 1, 6, 9 \ (\text{mod } 11), q \not\equiv 4, 7 \ (\text{mod } 13)$
and $d = 11q - 169$, we have

$$DCC(d, 9) \geq \frac{42}{2357947691}(d+169)(d+147)(d+103)(d+37)(d-51)$$
$$(d-161)(d-293)(d-447)(d-623). \qquad \square$$

3.4.3 Limitations

In [49], Lewis showed that an analogous class of constructions using finite fields to create graphs of diameter 2 is limited by the bound $L_C^-(2) \leq \frac{3}{8}$. By extending this logic, we can show that the constructions in this section have a similar limitation:

Theorem 3.14. *Let k be a positive integer. The direct product constructions of Theorems 3.10 and 3.12 can never yield a lower bound on $L_C^-(k)$ or $L_D^-(k)$ that exceeds $\frac{k+1}{2(k+2)^{k-1}}$.*

Proof (Lewis). First we consider the undirected case. Suppose the requirements of Theorem 3.10 hold and for each $i = 1, \ldots, k$, we have $r_i = q - a_i$, where $a_1 < a_2 < \ldots < a_k$. Let $\mathbb{T} = \mathbb{Z}_{q-a_1} \times \ldots \times \mathbb{Z}_{q-a_k} \times \mathbb{Z}_w$ and X be its generating set as in the proof of Theorem 3.10.

Since every element of $\mathbb{Z}_w$ is a sum of k distinct elements of B, no pair of which are inverses, we must have $w \leq \binom{k+2}{k} 2^k = (k+1)(k+2)2^{k-1}$.

By the requirements of Lemma 3.9, for any $i < j$, we have $m_{i,j} \leq r_i - r_j$ and $c_{\mathbf{o}} = \max_{i<j}(r_j + j m_{i,j})$. Hence, since $r_i = q - a_i$, we have $m_{i,j} \leq a_k - a_1$, and so $c_{\mathbf{o}} \leq q + k a_k$.

Thus X is the generating set for a Cayley graph on $\mathbb{T}$ with diameter k, degree d no greater than $2(k+2)q - 2\sum_{i=1}^{k} a_i + 2k a_k - 2a_1$, and order $n = w(q-a_1)(q-a_2)\ldots(q-a_k)$.

Hence:

$$n = \frac{w}{(2(k+2))^k}d^k + O(d^{k-1}) \leq \frac{(k+1)(k+2)2^{k-1}}{((2(k+2))^k}d^k + O(d^{k-1})$$
$$= \frac{k+1}{2(k+2)^{k-1}}d^k + O(d^{k-1})$$

as required.

The directed case is analogous. We follow Theorem 3.12 and its proof. In this case, every element of $\mathbb{Z}_w$ is the sum of k distinct elements of B, so $w \leq \binom{k+2}{k} = (k+1)(k+2)/2$, and X is the generating set for a Cayley graph on $\mathbb{T}$ with diameter k, degree $d \leq (k+2)q - \sum_{i=1}^{k} a_i + ka_k - a_1 - 1$, and order $n = w(q-a_1)(q-a_2)\dots(q-a_k)$.

Hence, $n = \frac{w}{(k+2)^k}d^k + O(d^{k-1}) \leq \frac{(k+1)(k+2)}{2(k+2)^k}d^k + O(d^{k-1}) = \frac{k+1}{2(k+2)^{k-1}}d^k + O(d^{k-1})$. □

Observe that, in the limit,

$$\lim_{k\to\infty} k\left(\frac{k+1}{2(k+2)^{k-1}}\right)^{1/k} = 1.$$

As a consequence, these direct product constructions themselves can never yield an improvement on the trivial lower bound for the limiting value of $R_C^-(k)$ or $R_D^-(k)$. However, it is possible to combine graphs of small diameter to produce larger graphs in such a way that we can improve on the trivial lower bound in the limit as the diameter increases. The next section introduces this idea.

3.5 A general graph product construction

The following theorem gives a simple way to combine two cyclic Cayley graphs to obtain a third cyclic Cayley graph. It is valid in both the directed and undirected cases.

Theorem 3.15. *Let Γ_1 and Γ_2 be two cyclic Cayley graphs of diameters k_1 and k_2, orders n_1 and n_2, and degrees d_1 and d_2 respectively. In the case of undirected graphs where d_1 and d_2 are both odd let $\delta = 1$, otherwise $\delta = 0$. In the directed case let $\delta = 0$ always. Then there exists a cyclic Cayley graph with diameter $k_1 + k_2$, degree at most $d_1 + d_2 + \delta$, and order $n_1 n_2$.*

Proof. Let S_1 be the connection set of Γ_1 so that $|S_1| = d_1$ and similarly for Γ_2. For convenience we consider each S_i to consist of elements within the interval

$(-n_i/2, n_i/2]$. Let G be the cyclic group $\mathbb{Z}_{n_1 n_2}$ and consider the connection set $S' = n_2 S_1 \cup S_2$. Then $|S'| \leq n_1 + n_2$.

We now construct a connection set S for the group G such that the Cayley graph $\mathrm{Cay}(G, S)$ has diameter $k_1 + k_2$. In the directed case we may simply take $S = S'$. In the undirected case we need to ensure that $S = -S$. If at least one of d_1, d_2 is even we may assume without loss of generality that d_2 is even and then we may again let $S = S'$ and $S = -S$ by construction.

It remains to consider the undirected case when d_1 and d_2 are both odd (the case $\delta = 1$). In that case we know $n_2/2 \in S_2 \subset S'$ and we let $S = S' \cup \{-n/2\}$ so that $S = -S$.

It is then clear that the Cayley graph $\mathrm{Cay}(G, S)$ has degree at most $d_1 + d_2 + \delta$, diameter $k_1 + k_2$ and order $n_1 n_2$. □

We can use this construction to obtain lower bounds on our L and R values for large diameters, given values for smaller diameters.

Corollary 3.16. *If $L(k)$ is one of $L_C^-(k)$, $L_C^+(k)$, $L_D^-(k)$ or $L_D^+(k)$ and $R(k)$ is one of $R_C^-(k)$, $R_C^+(k)$, $R_D^-(k)$ or $R_D^+(k)$, then*

(a) $$L(k_1 + k_2) \geq \frac{L(k_1) L(k_2) k_1^{k_1} k_2^{k_2}}{(k_1 + k_2)^{k_1 + k_2}}$$

(b) $$R(k_1 + k_2) \geq \left(R(k_1)^{k_1} R(k_2)^{k_2} \right)^{\frac{1}{k_1 + k_2}}$$

Proof. (a) Let $d > 1$. For $i = 1, 2$ we may construct graphs Γ_i of diameter k_i, degree $k_i d$ and order $L(k_i)(k_i d)^{k_i} + o(d^{k_i})$. Theorem 3.15 yields a product graph of diameter $k_1 + k_2$, degree at most $(k_1 + k_2)d + 1$ and order $L(k_1)L(k_2)k_1^{k_1} k_2^{k_2} d^{k_1 + k_2} + o(d^{k_1 + k_2})$.

Part *(b)* follows by straightforward algebraic manipulation. □

In particular, we note that the general product construction of Theorem 3.15 preserves lower bounds on the R values: $R(mk) \geq R(k)$ for every positive integer m.

We may use this idea to obtain better bounds for some particular diameters; for example we may improve on the undirected diameter 4 construction in Theorem 3.11:

Corollary 3.17.

$$L_C^+(4) \geq L_C^-(4) \geq \frac{169}{20736} \approx 0.0081501$$

	Diameter (k)							
	2	3	4	5	6	7	8	9
$R_{\max}(k) \approx$	1.41421	1.65096	1.80720	1.91926	2.00415	2.07100	2.12520	2.17016
$R_C^+(k) >$	1.20185^a	1.15455^d	1.20185^c	1.20431^d	1.20185^f	1.20360^f	1.20185^f	1.20321^f
$R_C^-(k) >$	1.20185^b	1.14775^d	1.20185^c	1.20431^d	1.20185^f	1.20360^f	1.20185^f	1.20321^f
$R_D^-(k) >$	1.22474^e	1.24805^e	1.26588^e	1.25881^e	1.27378^e	1.26436^e	1.26588^f	1.26514^f

Table 3.1: The best R values for diameter $k \leq 9$
a. Vetrík [74]; *b.* Theorem 3.7; *c.* Corollary 3.17;
d. Theorem 3.11; *e.* Theorem 3.13; *f.* Corollary 3.16

and hence

$$R_C^+(4) \geq R_C^-(4) > 1.20185$$

Proof. We note $L_C^-(2) \geq \frac{13}{36}$ from Theorem 3.7 and apply Corollary 3.16 with $k_1 = k_2 = 2$. □

Theorem 3.15 can be iterated to produce a construction for any desired diameter, and Corollary 3.16 then gives us a lower bound for the R values for that diameter. We illustrate the results for small diameter k in Table 3.1. As an indicator of progress we show also the largest possible value of R for a particular k, given by $R_{\max}(k) = k(k!)^{-1/k}$.

It is worth noting that the method of Corollary 3.16 may be used to produce values of R which are larger than those achievable from the direct product constructions of Section 3.4. For example, the limitations noted in Theorem 3.14 show that the maximum possible value of $R_D^-(10)$ we could achieve using Theorem 3.12 is approximately 1.26699. However, combining the results for diameters 4 and 6 in Table 3.1 yields $R_D^-(10) > 1.27061$.

Next we use our previous results to show that R is well-behaved in the limit.

Theorem 3.18. *Let $L(k)$ be one of $L_C^-(k)$, $L_C^+(k)$, $L_D^-(k)$ or $L_D^+(k)$, and let $R(k) = kL(k)^{1/k}$. The limit $R = \lim_{k\to\infty} R(k)$ exists and is equal to $\sup R(k)$.*

Proof (Bevan). $R(k)$ is bounded above (by e), so $s = \sup R(k)$ is finite. Hence, given $\varepsilon > 0$, we can choose k so that $s - R(k) < \varepsilon/2$. By Corollary 3.16(b), $R(mk) \geq R(k)$ for every positive integer m. Moreover, for any fixed $j < k$, since $R(j) \geq 1$, we have $R(mk + j) \geq R(k)^{mk/(mk+j)} \geq R(k)^{m/(m+1)}$, which, by choosing m large enough, can be made to differ from $R(k)$ by no more than $\varepsilon/2$. □

Corollary 3.19.

(a) $\lim_{k\to\infty} R_C^-(k) \geq \dfrac{5 \times 109^{1/5}}{7 \times 2^{3/5}} > 1.20431$

(b) $\lim_{k\to\infty} R_D^-(k) \geq \dfrac{3^{7/6}}{2^{3/2}} > 1.27378$

Proof. We choose the largest entry in the relevant row in Table 3.1. For *(a)* we know from Theorem 3.11 that $L_C^-(5) \geq \dfrac{109}{2^3 \times 7^5}$. For *(b)* we know from Theorem 3.13 that $L_D^-(6) \geq \dfrac{3}{2^{15}}$. □

We conclude this section by using the foregoing to derive new lower bounds for the maximum possible orders of circulant graphs of given diameter and sufficiently large degree.

Theorem 3.20.

(a) For any diameter $k \geq 2$ and any degree d large enough, $CC(d,k) > \left(1.14775\frac{d}{k}\right)^k$.

(b) For any diameter k that is a multiple of 5 or sufficiently large, and any degree d large enough, $CC(d,k) > \left(1.20431\frac{d}{k}\right)^k$.

(c) For any diameter $k \geq 2$ and any degree d large enough, $DCC(d,k) > \left(1.22474\frac{d}{k}\right)^k$.

(d) For any diameter k that is a multiple of 6 or sufficiently large, and any degree d large enough, $DCC(d,k) > \left(1.27378\frac{d}{k}\right)^k$.

Proof.

(a) Corollary 3.19 shows that for any k large enough, $R_C^-(k) > 1.20431$. We cannot choose a constant larger that 1.14775 because this value of R_C^- appears in Table 3.1 at diameter 3.

(b) For k a multiple of 5, we know from Theorem 3.11 and Corollary 3.16 that $R_C^-(k) > 1.20431$. The result for sufficiently large k follows from Corollary 3.19.

(c) and (d) follow by using similar logic in the directed case. □

These represent significant improvements over the trivial bound of $\left(\dfrac{d}{k}\right)^k$.

3.6 Application to sumsets covering $\mathbb{Z}_n$

Our constructions of directed circulant graphs can be used to obtain an upper bound on the minimum size, $SS(n,k)$, of a set $A \subset \mathbb{Z}_n$ for which the sumset

$$kA = \underbrace{A + A + \ldots + A}_{k} = \mathbb{Z}_n.$$

The trivial bound is $SS(n,k) \leq kn^{1/k}$ which follows in the same way as the trivial lower bound for the directed circulant graph (see Observation 3.2). Improvements to this trivial bound do not appear to have been investigated in the literature.

The idea is that, given $S \subseteq \mathbb{Z}_n$ such that $\mathrm{Cay}(\mathbb{Z}_n, S)$ has diameter k, if we let $A = S \cup \{0\}$ then $kA = \mathbb{Z}_n$. Our constructions thus enable us to bound $SS(n,k)$ for fixed k and infinitely many values of n. For example, if we let $L_S^-(k) = \liminf_{n\to\infty} SS(n,k)/n^{1/k}$, then the following new result for $k = 2$ follows from Theorem 3.13(a):

Corollary 3.21. $L_S^-(2) \leq \sqrt{\frac{8}{3}} \approx 1.63299$.

More generally, Corollary 3.19 shows that for large enough k and infinitely many values of n, $SS(n,k)$ is at least 21 percent smaller than the trivial bound:

Corollary 3.22. $\lim_{k\to\infty} k^{-1} L_S^-(k) \leq \frac{2^{3/2}}{3^{7/6}} \approx 0.78506$.

These covering sumsets are an interesting area of study in their own right, and we will return to the topic in Chapter 8.

3.7 Searching for optimal graphs

We can use the construction of Theorem 3.15 to obtain large undirected circulant graphs for small degrees and diameters. Recently in [32], Feria-Puron, Pérez-Rosés and Ryan published a table of largest known circulant graphs with degree up to 16 and diameter up to 10. Their method uses a construction based on graph Cartesian products which is somewhat similar to ours. In contrast, however, Theorem 3.15 does not in general result in a graph isomorphic to the Cartesian product of the constituents. Furthermore, our construction does not require the constituent graph orders to be coprime, which allows more graphs to be constructed.

Using Theorem 3.15 allowed us to improve many of the entries in the published table. However, at the same time we developed a computer search algorithm which allows us

to find circulant graphs of given degree, diameter and order. It turns out that this search is able to find larger graphs (at least in the range $d \leq 16, k \leq 10$) than the Theorem 3.15 method. We therefore describe this algorithm and present a much improved table of largest known circulant graphs.

3.7.1 The algorithm

We begin with a given order n, degree d and diameter k. We attempt to find a subset $S \subseteq \mathbb{Z}_n$ with $|S| = d$ such that $\mathrm{Cay}(\mathbb{Z}_n, S)$ has diameter k. As usual, this involves trying to find a generating set $S = \{\pm s_1, \pm s_2, \ldots, \pm s_f\}$ such that any element of $\mathbb{Z}_n$ can be expressed as a sum of at most k elements of S. (If the degree d is odd, we must have the unique involution $n/2$ as an additional member of any generating set S.) The number $f = \lfloor \frac{d}{2} \rfloor$ is the *dimension* of the problem; that is, the number of choices of elements s_i we need to make. It will be convenient for us to consider the *reduced* form $r(S)$ of a generating set S, where $r(S) = \{s_1, s_2, \ldots, s_f\}$ consists only of elements s_i in the range $1 \ldots \lfloor \frac{n-1}{2} \rfloor$. For clarity, we always write $r(S)$ in numerical order so that $s_1 < s_2 < \cdots < s_f$. It is clear that the set S can be recovered from $r(S)$ uniquely.

A naive approach to the search would be simply to carry out an exhaustive enumeration of possible reduced generating sets $r(S)$ and then test whether each leads to a diameter k graph. However, for moderately large d, k or n, this procedure quickly becomes infeasible. To reduce the search space we employ two specific techniques. The first is a variation on a traditional "branch and bound" algorithm.

We view the search space as a tree, with each branch of the tree represented by a partial reduced set $\{s_1, s_2, \ldots, s_m\}$ for some $m \leq f$. The subtree of the search space represented by this branch consists of all the full reduced sets of size f which begin with these m elements. We search the tree depth first in *lexicographic* order; that is to say, we order the reduced sets first by their smallest element, then by their second smallest if there is a tie and so on. In every branch, the algorithm keeps track of how many elements of $\mathbb{Z}_n$ have been covered so far by sums of up to k of the reduced set of size m. It is straightforward to calculate the maximum possible number of uncovered elements which may become covered by adding the remaining $f - m$ choices to the reduced set. If this is insufficient to cover the whole of $\mathbb{Z}_n$, we can discard the branch and avoid searching the subtree.

The second search reduction technique uses the idea of isomorphism avoidance. Given a Cayley graph $\mathrm{Cay}(G, S)$ on some group G and some automorphism ϕ of G, it can be shown (for example Biggs [12, Proposition 16.2]) that the graph $\mathrm{Cay}(G, \phi(S))$ is isomorphic to $\mathrm{Cay}(G, S)$. In the case of cyclic groups, the automorphisms of $\mathbb{Z}_n$ are

precisely the maps $\phi_\alpha : x \mapsto \alpha x$ where $\gcd(\alpha, n) = 1$. An automorphism ϕ_α of $\mathbb{Z}_n$ acts on a reduced set $r(S)$ in a natural way by defining:

$$\phi_\alpha(r(S)) = r(\{\alpha s : s \in S\}).$$

In this way, the orbit of a reduced set under the action of the automorphism group of $\mathbb{Z}_n$ is the set:

$$\mathrm{Orb}(r(S)) = \{\phi_\alpha(r(S)) : \gcd(\alpha, n) = 1\}.$$

The idea now is that a branch represented by a reduced set need not be searched if we have already searched another branch in the same orbit. To that end we define the canonical orbit representative CR of a reduced set $r(S)$ to be:

$$CR(r(S)) = \min(\mathrm{Orb}(r(S)))$$

where of course the minimum is taken with respect to the lexicographic ordering.

We implement this technique by computing, at the start of processing of a branch, its canonical orbit representative. If this is smaller, we can skip the branch. For practical purposes, we only implement this test down to a fixed depth in the search tree, since otherwise the overhead of computation would outweigh the gains.

The combination of these techniques has allowed us to make substantial improvements to the table of largest known circulant graphs, which we now present.

3.7.2 Revised table

In Table 3.2, we show the largest known circulant graphs of degree $d \leq 16$ and diameter $k \leq 10$. In Table 3.3 we give a reduced generating set $r(S)$ for each new record largest graph found. The computer search described has been completed as an exhaustive search in the diameter 2 case up to degree 23, and these results are included in Table 3.3 for completeness.

$d \setminus k$	1	2	3	4	5	6	7	8	9	10
2	3	5	7	9	11	13	15	17	19	21
3	4	8	12	16	20	24	28	32	36	40
4	5	13	25	41	61	85	113	145	181	221
5	6	16	36	64	100	144	196	256	324	400
6	7	21	55	117	203	333	515	737	1027	1393
7	8	26	76	160	308	536	828	1232	1764	2392
8	9	35	104	248	528	984	1712	2768	4280	6320
9	10	42	130	320	700	1416	2548	4304	6804	10320
10	11	51	177	457	1099	2380†	4551†	8288†	14099†	22805†
11	12	56	210	576	1428†	3200†	6652†	12416†	21572†	35880†
12	13	67	275	819†	2040†	4283†	8828†	16439†	29308†	51154†
13	14	80	312	970†	2548†	5598†	12176†	22198†	40720†	72608†
14	15	90	381	1229†	3244†	7815†	17389†	35929†	71748†	126109†
15	16	96	448	1420†	3980†	9860†	22584†	48408†	93804†	177302†
16	17	112	518†	1717†	5024†	13380†	32731†	71731†	148385†	298105†

Table 3.2: Largest known circulant graphs of degree $d \leq 16$ and diameter $k \leq 10$
† new record largest value

d	k	Order	Generators
6	2	21*	1, 2, 8
6	3	55*	1, 5, 21
6	4	117*	1, 16, 22
6	5	203*	1, 7, 57
6	6	333*	1, 9, 73
6	7	515*	1, 46, 56
6	8	737*	1, 11, 133
6	9	1027*	1, 13, 157
6	10	1393*	1, 92, 106
7	2	26*	1, 2, 8
7	3	76*	1, 27, 31
7	4	160*	1, 5, 31
7	5	308*	1, 7, 43
7	6	536*	1, 231, 239
7	7	828*	1, 9, 91
7	8	1232*	1, 11, 111
7	9	1764*	1, 803, 815
7	10	2392*	1, 13, 183
8	2	35*	1, 6, 7, 10
8	3	104*	1, 16, 20, 27
8	4	248*	1, 61, 72, 76
8	5	528*	1, 89, 156, 162
8	6	984*	1, 163, 348, 354
8	7	1712*	1, 215, 608, 616
8	8	2768	1, 345, 1072, 1080
8	9	4280	1, 429, 1660, 1670
8	10	6320	1, 631, 2580, 2590
9	2	42*	1, 5, 14, 17
9	3	130*	1, 8, 14, 47
9	4	320*	1, 15, 25, 83
9	5	700*	1, 5, 197, 223
9	6	1416	1, 7, 575, 611
9	7	2548	1, 7, 521, 571
9	8	4304	1, 9, 1855, 1919
9	9	6804	1, 9, 1849, 1931
9	10	10320	1, 11, 4599, 4699
10	2	51*	1, 2, 10, 16, 23
10	3	177*	1, 12, 19, 27, 87
10	4	457*	1, 20, 130, 147, 191
10	5	1099*	1, 53, 207, 272, 536
10	6	2380	1, 555, 860, 951, 970
10	7	4551	1, 739, 1178, 1295, 1301
10	8	8288	1, 987, 2367, 2534, 3528
10	9	14099	1, 1440, 3660, 3668, 6247
10	10	22805	1, 218, 1970, 6819, 6827
11	2	56*	1, 2, 10, 15, 22
11	3	210*	1, 49, 59, 84, 89
11	4	576*	1, 9, 75, 155, 179
11	5	1428	1, 169, 285, 289, 387
11	6	3200	1, 259, 325, 329, 1229
11	7	6652	1, 107, 647, 2235, 2769
11	8	12416	1, 145, 863, 4163, 5177
11	9	21572	1, 663, 6257, 10003, 10011

d	k	Order	Generators
11	10	35880	$1, 2209, 5127, 5135, 12537$
12	2	67*	$1, 2, 3, 13, 21, 30$
12	3	275*	$1, 16, 19, 29, 86, 110$
12	4	819	$7, 26, 119, 143, 377, 385$
12	5	2040	$1, 20, 24, 152, 511, 628$
12	6	4283	$1, 19, 100, 431, 874, 1028$
12	7	8828	$1, 29, 420, 741, 2727, 3185$
12	8	16439	$1, 151, 840, 1278, 2182, 2913$
12	9	29308	$1, 219, 1011, 1509, 6948, 8506$
12	10	51154	$1, 39, 1378, 3775, 5447, 24629$
13	2	80*	$1, 3, 9, 20, 25, 33$
13	3	312*	$1, 14, 74, 77, 130, 138$
13	4	970	$1, 23, 40, 76, 172, 395$
13	5	2548	$1, 117, 121, 391, 481, 1101$
13	6	5598	$1, 12, 216, 450, 1204, 2708$
13	7	12176	$1, 45, 454, 1120, 1632, 1899$
13	8	22198	$1, 156, 1166, 2362, 5999, 9756$
13	9	40720	$1, 242, 3091, 4615, 5162, 13571$
13	10	72608	$1, 259, 4815, 8501, 8623, 23023$
14	2	90*	$1, 4, 10, 17, 26, 29, 41$
14	3	381*	$1, 11, 103, 120, 155, 161, 187$
14	4	1229	$1, 8, 105, 148, 160, 379, 502$
14	5	3244	$1, 108, 244, 506, 709, 920, 1252$
14	6	7815	$1, 197, 460, 696, 975, 2164, 3032$
14	7	17389	$1, 123, 955, 1683, 1772, 2399, 8362$
14	8	35929	$1, 796, 1088, 3082, 3814, 13947, 14721$
14	9	71748	$1, 1223, 3156, 4147, 5439, 11841, 25120$
14	10	126109	$1, 503, 4548, 7762, 9210, 9234, 49414$
15	2	96*	$1, 2, 3, 14, 21, 31, 39$
15	3	448*	$1, 10, 127, 150, 176, 189, 217$
15	4	1420	$1, 20, 111, 196, 264, 340, 343$
15	5	3980	$1, 264, 300, 382, 668, 774, 1437$
15	6	9860	$1, 438, 805, 1131, 1255, 3041, 3254$
15	7	22584	$1, 1396, 2226, 2309, 2329, 4582, 9436$
15	8	48408	$1, 472, 2421, 3827, 4885, 5114, 12628$
15	9	93804	$1, 3304, 4679, 9140, 10144, 10160, 13845$
15	10	177302	$1, 2193, 8578, 18202, 23704, 23716, 54925$
16	2	112*	$1, 4, 10, 17, 29, 36, 45, 52$
16	3	518	$1, 8, 36, 46, 75, 133, 183, 247$
16	4	1717	$1, 46, 144, 272, 297, 480, 582, 601$
16	5	5024	$1, 380, 451, 811, 1093, 1202, 1492, 1677$
16	6	13380	$1, 395, 567, 1238, 1420, 1544, 2526, 4580$
16	7	32731	$1, 316, 1150, 1797, 2909, 4460, 4836, 16047$
16	8	71731	$1, 749, 4314, 7798, 10918, 11338, 11471, 25094$
16	9	148385	$1, 6094, 6964, 10683, 11704, 14274, 14332, 54076$
16	10	298105	$1, 5860, 11313, 15833, 21207, 26491, 26722, 99924$
17	2	130*	$1, 7, 26, 37, 47, 49, 52, 61$
18	2	138*	$1, 9, 12, 15, 22, 42, 27, 51, 68$
19	2	156*	$1, 15, 21, 23, 26, 33, 52, 61, 65$
20	2	171*	$1, 11, 31, 36, 37, 50, 54, 47, 65, 81$
21	2	192*	$1, 3, 15, 23, 32, 51, 57, 64, 85, 91$
22	2	210*	$2, 7, 12, 18, 32, 35, 63, 70, 78, 91, 92$
23	2	216*	$1, 3, 5, 17, 27, 36, 43, 57, 72, 83, 95$

Table 3.3: Largest circulant graphs of small degree d and diameter k found by computer search
* proven extremal

Chapter 4

The degree-diameter problem for dihedral Cayley graphs

If cyclic groups are the most obvious starting point for an investigation of Cayley graphs, then natural next steps would be general abelian groups or dihedral groups. We will return briefly in Chapter 8 to a particular class of non-cyclic abelian groups. Here we study the problem of dihedral groups. We saw in Chapter 3 that even in the simplest case of diameter 2, we do not have a complete understanding of the asymptotic behaviour of Cayley graphs of cyclic groups. In contrast, we will now see that we do have such an understanding for dihedral groups.

4.1 Diameter 2

Recall that we have an upper limit of $d^2 + 1$ (the Moore bound) for a graph of maximum degree d and diameter 2. Thus for any family of diameter 2 graphs the largest possible asymptotic order is d^2. A recent result of Abas [1] shows that a Cayley graph of diameter 2 and asymptotic order $d^2/2$ can be constructed for any degree d using direct products of dihedral and cyclic groups.

In this section we show that the asymptotic limit for dihedral groups is precisely $d^2/2$, first by obtaining a lower bound by way of a construction involving Galois fields, and then by finding an upper bound for generalised dihedral groups by a counting argument. We follow the structure of our published paper [29].

We denote the dihedral group of order $2n$ by D_{2n}. We will view the usual dihedral group as an example of a *generalised dihedral group* $G \rtimes C_2$ which is a semidirect product of an abelian group G with the multiplicative group $\{\pm 1\}$ where the action on G is via its inversion automorphism. For a group G and a subset $S \subseteq G$ which is inverse-closed and identity-free, recall that the graph $\mathrm{Cay}(G, S)$ is vertex-transitive and hence regular, with degree $d = |S|$, and has diameter at most k if and only if each element of G can be expressed as a product of no more than k elements of the generating set S.

By $DC(d, k)$ we mean the largest number of vertices in a Cayley graph of a dihedral

group having degree d and diameter k.

4.1.1 Results

Our first result uses a construction based on finite fields to obtain a lower bound for $DC(d, 2)$ for certain values of d. The method is similar to constructions in [53]. We also make use of a folklore result for cyclic groups which we include for completeness.

Lemma 4.1. *Let $n > 1$. Then the cyclic group $\mathbb{Z}_n$ has a diameter 2 Cayley graph with a generating set of size at most $2\lceil\sqrt{n}\rceil$.*

Proof. Let $K = \lceil\sqrt{n}\rceil, M = \lfloor\frac{K}{2}\rfloor$ and take a generating set consisting of $\{\pm 1, \pm 2, \ldots, \pm M, \pm K, \pm 2K, \ldots, \pm MK\}$. □

Lemma 4.2. *If p is any prime and $d = 2(p + \lceil\sqrt{p}\rceil - 1)$, then $DC(d, 2) \geq 2p(p - 1)$.*

Proof. Let F be the Galois field $GF(p)$ where p is a prime. The additive and multiplicative groups F^+ and F^* are cyclic of coprime orders so that $F^+ \times F^*$ is a cyclic group of order $n = p(p - 1)$. Consider the dihedral group D_{2n} as a semidirect product $G = (F^+ \times F^*) \rtimes C_2$ where the cyclic group C_2 is thought of as the multiplicative group $\{\pm 1\}$ and acts on $F^+ \times F^*$ by inversion. Specifically, the multiplication rule is:

$$(a, b, c)(\alpha, \beta, \gamma) = (a + \alpha c, b\beta^c, c\gamma)$$

The subgroup $C = \langle(1, 1, 1)\rangle$ is cyclic of order p and so by Lemma 4.1 it has a diameter 2 Cayley graph with respect to some generating set $\{c_1, c_2, \ldots, c_{2\lceil\sqrt{p}\rceil}\}$ of cardinality $2\lceil\sqrt{p}\rceil$. Consider now a generating set S of the full group G containing:

$$\begin{aligned} v = (0, 1, -1) &\qquad (1 \text{ element}) \\ a_x = (0, x, 1), x \in F^* \setminus \{1\} &\qquad (p - 2 \text{ elements}) \\ b_x = (x, x, -1), x \in F^* &\qquad (p - 1 \text{ elements}) \\ c_i, i = 1 \ldots 2\lceil\sqrt{p}\rceil &\qquad (2\lceil\sqrt{p}\rceil \text{ elements}) \end{aligned}$$

Since $v^{-1} = v, a_x^{-1} = a_{x^{-1}}, b_x^{-1} = b_x$ and $\{c_1, c_2, \ldots, c_{2\lceil\sqrt{p}\rceil}\}$ is inverse-closed it follows that S is inverse-closed. To show that the diameter is 2, it suffices to show that every element of the group can be expressed as the product of at most two of these generators. We consider all the possible cases as follows.

If $x \neq 0, x \neq y$ then $(x, y, -1) = (0, z, 1)(x, x, -1) = a_z b_x$ where $z = yx^{-1}$.

If $x \neq 0, x = y$ then $(x, y, -1) = (x, x, -1) = b_x$.

If $x = 0, y \neq 1$ then $(x, y, -1) = (0, y, 1)(0, 1, -1) = a_y v$.

If $x = 0, y = 1$ then $(x, y, -1) = (0, 1, -1) = v$.

If $y \neq 1, x \neq 0$ then $(x, y, 1) = (z, z, -1)(t, t, -1) = b_z b_t$ where
$z = yx(y-1)^{-1}, t = x(y-1)^{-1}$.

If $y \neq 1, x = 0$ then $(x, y, 1) = (0, y, 1) = a_y$.

If $y = 1$ then $(x, y, 1) \in C$ and so is the product of at most two c_i.

Since $|S| = 2(p + \lceil \sqrt{p} \rceil - 1)$ the result follows. □

The previous result shows that $\limsup_{d \to \infty} \frac{DC(d,2)}{d^2} \geq \frac{1}{2}$. The next result shows that $1/2$ is in fact also an upper bound.

Lemma 4.3. *Let G be a generalised dihedral group of order $2n$ and let S be an inverse-closed generating set for G not containing the identity. Suppose that the Cayley graph $\mathrm{Cay}(G, S)$ has diameter 2. Then the degree d of $\mathrm{Cay}(G, S)$ satisfies $d \geq 2\sqrt{n} - 1$.*

Proof. Let $G = H \rtimes C_2$ where H is an abelian group of order n and C_2 acts on H by inversion. Let C be the index 2 subgroup of G isomorphic to H and write $S = A \cup B$ where $A \subset C$ and $B \subset G \setminus C$. Let $m_1 = |A|, m_2 = |B|$.

Consider how the n elements of $G \setminus C$ can be expressed as a product of at most two elements in S. There are m_2 possibilities from the set B itself, then $m_1 m_2$ elements of the form ab where $a \in A, b \in B$. Since $a^{-1}b = ba$ and the set A is inverse-closed the products of the form ba do not contribute any further elements. So we require:

$$m_2(m_1 + 1) \geq n$$

The degree d of the Cayley graph is $|S| = m_1 + m_2$. All numbers are inherently positive and so elementary calculus shows that the minimum possible value of $m_1 + m_2$ occurs when $m_2 = m_1 + 1 = \sqrt{n}$. So $d \geq 2\sqrt{n} - 1$. □

The bound $|G| \leq \frac{1}{2}(d+1)^2$ in Lemma 4.3 is valid for all values of d, but as it stands Lemma 4.2 only holds for a restricted set of values. We can extend the result of Lemma 4.2 by using the ideas first used in [66] to obtain a lower bound valid for all values of d.

Lemma 4.4. *Let $d \geq 6$ and let p be the largest prime satisfying $D(p) = 2(p + \lceil\sqrt{p}\rceil - 1) \leq d$. Then $DC(d, 2) \geq 2p(p-1)$.*

Proof. Let p be as in the statement, $n = 2p(p-1)$ and $G = D_n$. By Lemma 4.2 there is an inverse-closed unit-free subset $S \subset G$ with $|S| = D(p)$ such that $\mathrm{Cay}(G, S)$ has diameter 2. We can add $d - D(p)$ involutions from $G \setminus S$ to form a new inverse-closed unit-free generating set S'. The diameter of $\mathrm{Cay}(G, S')$ is still 2 and the result follows. □

Using the method of [66] and Lemma 3.3 we may use ideas from number theory to obtain a result independent of p. Specifically, from [7, Therorem 1] we know that for all sufficiently large D, there is some prime p in the range $D - D^{0.525} \leq p \leq D$.

Theorem 4.5. *For all sufficiently large d, $DC(d, 2) \geq 0.5d^2 - 1.39d^{1.525}$.*

Proof. For given d, let p be the largest prime such that $2(p + \lceil\sqrt{p}\rceil - 1) \leq d$. Then p is at least as large as the largest prime q satisfying $2(q + \sqrt{q}) \leq d$. Rearranging this we find that q is the largest prime not exceeding $D = \frac{1}{2}(d - \sqrt{2d+1} + 1)$. By [7, Therorem 1], for sufficiently large d we have $q \geq D - D^{0.525}$. So for large d we have:

$$\begin{aligned} p \geq q &\geq D - \left(\frac{d}{2}\right)^{0.525} \\ &= \frac{1}{2}\left(d - (2d+1)^{0.5} + 1 - 2^{0.475} d^{0.525}\right) \end{aligned}$$

For large d the term in $d^{0.525}$ dominates terms of lower powers of d and since $2^{0.475} \approx 1.389918$, for large d we have $p \geq \frac{1}{2}\left(d - 1.38992 d^{0.525}\right)$. For sufficiently large d we therefore have

$$2p(p-1) \geq \frac{d^2}{2} - 1.39 d^{1.525}$$

Since we can construct a Cayley graph of degree d and diameter 2 on the dihedral group of order $2p(p-1)$ by Lemma 4.4 the result follows. □

Lemma 4.3 and Theorem 4.5 allow us to determine completely the asymptotic behaviour of $DC(d, 2)$.

Corollary 4.6.

$$\lim_{d \to \infty} \frac{DC(d, 2)}{d^2} = \frac{1}{2}$$

and hence in the class of Cayley graphs of dihedral groups,

$$L^-(2) = L^+(2) = \frac{1}{2}$$

4.1.2 Remarks

It is tempting to try to extend these results to other split extensions of abelian groups where the action is via an automorphism other than the inversion map. However the counting argument of Lemma 4.3 relies on the fact that in our Cayley graph $\mathrm{Cay}(G, S)$ the generating set S has the very particular form $\{(a, 1) : a \in S_1\} \cup \{(b, -1) : b \in S_2\}$ where the set S_1 is inverse-closed and hence closed under the acting automorphism. The upper asymptotic bound $d^2/2$ would not necessarily hold for other more general semidirect products of an abelian group with C_2, although no family of such groups with a larger bound is known.

We illustrate this remark with a couple of examples. Firstly, the construction of Abas [1], uses a direct product of the form $D_{2m} \times \mathbb{Z}_n$ which we may regard as the semidirect product $(\mathbb{Z}_m \times \mathbb{Z}_n) \rtimes C_2$, where C_2 acts on $\mathbb{Z}_m$ via its inversion automorphism and on $\mathbb{Z}_n$ via the identity automorphism. In this case a generating set S which is inverse-closed is not necessarily of the form in the previous paragraph, since the action in the semidirect product is not inversion.

However, we may modify the construction to obtain a family of groups for which the counting argument of Lemma 4.3 does hold. In the first example above, if $n = 2$ then the identity automorphism coincides with the inversion automorphism and the Lemma applies. Further, if we replace $\mathbb{Z}_2$ by any elementary abelian 2-group then its inversion automorphism is the identity and the argument continues to hold. Thus the upper bound of Lemma 4.3 holds for any group of the form $D_{2n} \times H$ where H is an elementary abelian 2-group.

We note also that the argument of Lemma 4.3 holds in the case of dicyclic groups. For any $n \geq 1$, the dicyclic group of order $4n$ has presentation $\langle a, b \,|\, a^{2n} = b^4 = 1, a^n = b^2, ab = ba^{-1} \rangle$. The final relator in the presentation means that in any inverse-closed generating set, asymptotically half the possible products are duplicated as in the argument of Lemma 4.3. We therefore have the following corollary.

Corollary 4.7. *In the class of Cayley graphs of dicyclic groups,*

$$L^-(2) \leq \frac{1}{2}.$$

4.2 Bounds for larger diameters

The construction of the previous section cannot readily be extended to higher diameters. However, we can extend the logic of Lemma 4.3 to obtain an upper bound for diameters 3 and 4. We begin by recalling the method of Lemma 4.3.

Our group $G = D_{2n}$ has an index 2 cyclic subgroup C of order n. Our generating set S is decomposed into $A = S \cap C$ and $B = S \cap (G \setminus C)$. We set $m_1 = |A|, m_2 = |B|$. The argument proceeds by counting the ways in which we can multiply at most two elements of S.

If $V = S \cup SS$ is the set of elements which can be written as a product of at most two elements of S, then the proof of the lemma shows that $|V \cap (G \setminus C)| \leq m_1 m_2 + m_2$. By a similar argument we can see that $|V \cap C| \leq \frac{1}{2}m_1^2 + m_2^2 + m_1 - m_2 + 1$.

We now use this information to bound the possible order of a Cayley graph of a dihedral group of diameter 2. To ease the notation we denote the degree of the graph by d and the number m_2 of generating elements outside the cyclic subgroup by x. Then $m_1 = d - x$. The functions above can then be expressed in terms of the single variable x which makes finding the maximum easier.

To be able to generate a dihedral group of order $2n$, for a given d and x we must have $f(x) \geq n$ and $g(x) \geq n$. Thus a bound on the largest possible graph of diameter 2 and degree d is given by:

$$N = 2 \max_{0 \leq x \leq d} \min\{f(x), g(x)\}$$

The proof of Lemma 4.3 then solves this in the diameter 2 case. We can now use a similar method for diameters 3 and 4.

4.2.1 Diameter 3

With the notation of the previous section, we now set $V = S \cup SS \cup SSS$ to be the set of elements which can be written as a product of at most 3 elements of S. It is a straightforward but tedious counting exercise to show that bounds on V are given by:

$$|V \cap C| \leq \frac{1}{6}m_1^3 + m_1 m_2^2 + \frac{1}{2}m_1^2 - m_1 m_2 + m_2^2 + \frac{4}{3}m_1 - m_2 + 1$$

$$|V \cap (G \setminus C)| \leq \frac{1}{2}m_1^2 m_2 + \frac{1}{2}m_2^3 + m_1 m_2 - \frac{1}{2}m_2^2 + m_2$$

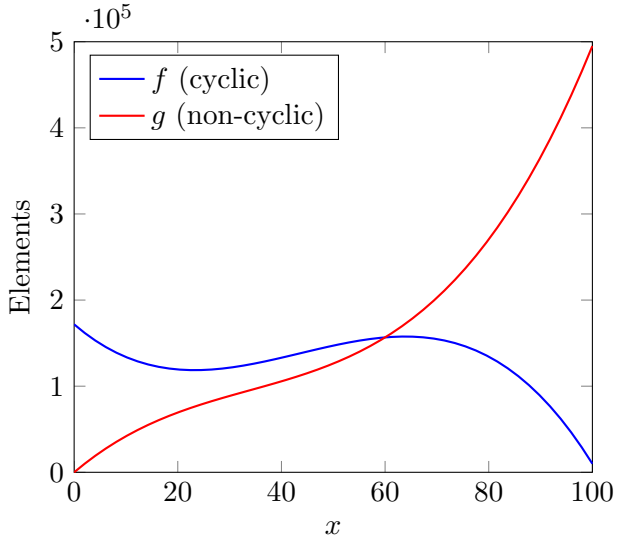

Figure 4.1: Illustrative plot of diameter 3 polynomials with $d = 100$

Once again we express the formulae for cyclic and non-cyclic elements in terms of a single variable $x = m_2$.

$$f(x) = \frac{1}{6}(d-x)^3 + (d-x)x^2 + \frac{1}{2}(d-x)^2 - (d-x)x + x^2 + \frac{4}{3}(d-x) - x + 1$$

$$g(x) = \frac{1}{2}(d-x)^2 x + \frac{1}{2}x^3 + (d-x)x - \frac{1}{2}x^2 + x$$

To be able to generate a dihedral group of order $2n$, for a given d and x we must have $f(x) \geq n$ and $g(x) \geq n$. Thus a bound on the largest possible graph of diameter 3 and degree d is given by:

$$N = 2 \max_{0 \leq x \leq d} \min\{f(x), g(x)\}$$

It remains to find this maximum. We want to find the value of x (as a proportion of d) which maximises the order, and find an expression for that order. It is helpful to view these functions graphically to see how they behave as we vary x from 0 to d in a particular case. An illustrative plot appears in Figure 4.1 in the case $d = 100$.

For our solution, we are in fact content with the leading term i.e. the d^3 term in this case. It is clear that the maximum will occur at one of the following points:

- An end point of the interval: $x = 0$ or $x = d$
- A critical point (local maximum) of either f or g
- An intersection point of f and g

In practice we can discount the end points since $x = 0$ can generate no non-cyclic

elements at all and $x = d$ leads to a bipartite graph. In the diameter 3 case we see that the maximum occurs at the local maximum of f. Elementary calculus shows that this local maximum is at $x = \frac{1}{21}\left(9d + 15 + \sqrt{18d^2 + 18d - 69}\right)$.

We want an asymptotic expression for this local maximum in terms of d, so it suffices to take $x = Kd$ where $K = \frac{3+\sqrt{2}}{7}$. At this point, the value of f is $d^3\left(\frac{1}{6}(1-K)^3 + K^2(1-K)\right) + o(d^3) = \frac{20+2\sqrt{2}}{147}d^3 + o(d^3)$.

So we have proved:

Theorem 4.8. *In the class of Cayley graphs of dihedral groups,*

$$L^+(3) \leq \frac{40+4\sqrt{2}}{147} \approx 0.31059$$

We note that the optimal asymptotic proportion of involutions in the generating set is $K \approx 0.63060$. This is in contrast to the diameter 2 case where the optimal proportion was $\frac{1}{2}$.

4.2.2 Diameter 4

We use the same method, this time setting $V = S \cup SS \cup SSS \cup SSSS$. The equations bounding the possible sizes of V are more awkward, but with the help of computer algebra packages we obtain:

$$\begin{aligned}|V \cap C| \leq{}& \frac{1}{24}m_1^4 + \frac{1}{2}m_1^2m_2^2 + \frac{1}{4}m_2^4 + \frac{1}{6}m_1^3 - \frac{1}{2}m_1^2m_2 + m_1m_2^2 \\ &- \frac{1}{2}m_2^3 + \frac{5}{6}m_1^2 - m_1m_2 + \frac{7}{4}m_2^2 + \frac{4}{3}m_1 - \frac{3}{2}m_2 + 1\end{aligned}$$

$$\begin{aligned}|V \cap (G \setminus C)| \leq{}& \frac{1}{6}m_1^3m_2 + \frac{1}{2}m_1m_2^3 + \frac{1}{2}m_1^2m_2 - \frac{1}{2}m_1m_2^2 \\ &+ \frac{1}{2}m_2^3 + \frac{4}{3}m_1m_2 - \frac{1}{2}m_2^2 + m_2\end{aligned}$$

Once again we express the formulae for cyclic and non-cyclic elements in terms of a single variable $x = m_2$.

$$\begin{aligned}f(x) = {}& \frac{1}{24}(d-x)^4 + \frac{1}{2}(d-x)^2x^2 + \frac{1}{4}x^4 + \frac{1}{6}(d-x)^3 - \frac{1}{2}(d-x)^2x + (d-x)x^2 \\ &- \frac{1}{2}x^3 + \frac{5}{6}(d-x)^2 - (d-x)x + \frac{7}{4}x^2 + \frac{4}{3}(d-x) - \frac{3}{2}x + 1\end{aligned}$$

$$\begin{aligned}g(x) = {}& \frac{1}{6}(d-x)^3x + \frac{1}{2}(d-x)x^3 + \frac{1}{2}(d-x)^2x - \frac{1}{2}(d-x)x^2 \\ &+ \frac{1}{2}x^3 + \frac{4}{3}(d-x)x - \frac{1}{2}x^2 + x\end{aligned}$$

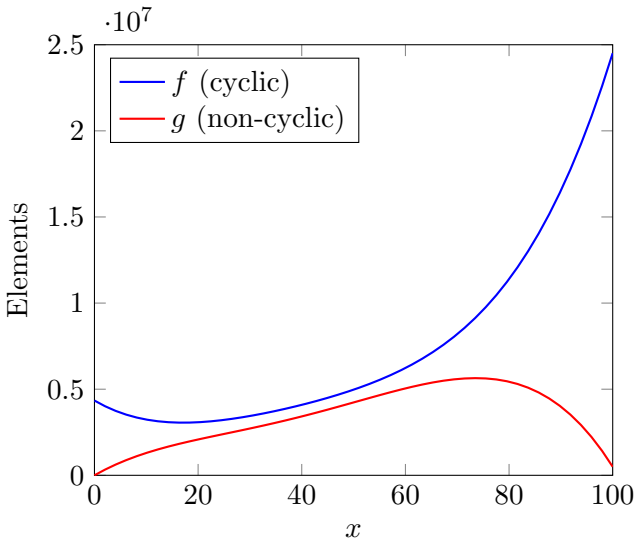

Figure 4.2: Illustrative plot of diameter 4 polynomials with $d = 100$

To be able to generate a dihedral group of order $2n$, for a given d and x we must have $f(x) \geq n$ and $g(x) \geq n$. Thus a bound on the largest possible graph of diameter 4 and degree d is given by:

$$N = 2 \max_{0 \leq x \leq d} \min\{f(x), g(x)\}$$

Again it is helpful to view the illustrative plot in Figure 4.2. Although the equations are much more complex, with the aid of computer algebra we can deduce that this time the optimal point is the local maximum of g which occurs at $x = Kd + o(d)$ where $K = \dfrac{k + \frac{1}{k} + 3}{8}, \quad k = (4\sqrt{3} + 7)^{1/3}$.

We can then show that:

Theorem 4.9. *In the class of Cayley graphs of dihedral groups,*

$$L^+(4) \leq \frac{K(1-K)(1-2K+4K^2)}{3} \approx 0.10983$$

The optimal asymptotic proportion of involutions in the generating set this time is $K \approx 0.72771$.

4.2.3 Comments

Ideally, we would have a construction for diameter 3 and 4 Cayley graphs of dihedral groups which at least approaches these bounds. Unfortunately, the construction from the diameter 2 case based on finite fields cannot readily be extended to larger diameters.

Chapter 5

Large Cayley graphs of fixed small diameter

Chapters 3 and 4 derived new asymptotic bounds in the degree-diameter problem for the classes of Cayley graphs on, respectively, cyclic groups and dihedral groups. We now turn our attention to the problem of general Cayley graphs.

Much of the existing literature is focused on the diameter 2 case. For larger diameters k, we have a useful lower bound $L^-(k) \geq k/3^k$ for any $k > 2$ by Macbeth, Šiagiová, Širáň and Vetrík [54]. For the specific cases of diameters 3, 4 and 5, Vetrík [73] has a series of constructions giving the best known asymptotic results, namely $L^-(3) \geq \frac{3}{16}$; $L^-(4) \geq \frac{32}{5^4}$ and $L^-(5) \geq \frac{25}{4^5}$. Both these papers above use variations of a semidirect product construction, and we generalise this idea in Section 5.2 to obtain improvements for diameters 3 to 7.

In the directed case, the best currently available results are by Vetrík [72] who shows that $L^-(2) \geq 8/9$ and for $k \geq 3$, $L^-(k) \geq k/2^k$. We generalise our method in Section 5.2 to the directed case to obtain new larger bounds at diameters 3, 4 and 5.

However we begin with a new construction for diameter 3, for which the result given in Vetrík is that $L^-(3) \geq \frac{3}{16}$. Our improved bound uses a construction which is, as far as we know, the first example to use matrix groups over finite fields.

5.1 Cayley graphs of matrix groups

Our strategy will be to find a suitable Cayley graph on a group based on a particular subgroup of $SL(3, p)$ for any odd prime p. However, as in the case of the diameter 2 Cayley graphs of dihedral groups from Chapter 4, we will be left with some awkward subgroups which our chosen generating set is unable to cover directly. We therefore begin with two lemmas on diameter 3 Cayley graphs of cyclic and elementary abelian groups, extending the idea noted in the dihedral proof.

Lemma 5.1. *For any $n \geq 6$ there is a subset $S \subseteq \mathbb{Z}_n$ of cardinality $6\left\lceil \frac{n^{1/3}}{2} \right\rceil$ such that $\mathrm{Cay}(\mathbb{Z}_n, S)$ has diameter at most 3.*

Proof. Let $n \geq 6$ and let $K = \lceil n^{1/3} \rceil$ and $M = \lfloor \frac{K}{2} \rfloor$. Let $S \subseteq \mathbb{Z}_n$ be the set $\{\pm 1, \pm 2, \ldots, \pm M, \pm K, \pm 2K, \ldots, \pm MK, \pm K^2, \pm 2K^2, \ldots, \pm MK^2\}$. Then it is easy to see that we can express any element of $\mathbb{Z}_n$ as a sum of at most 3 elements of S. □

Lemma 5.2. *For all large n, there is a subset $T \subseteq \mathbb{Z}_n \times \mathbb{Z}_n$ of cardinality $9n^{2/3} + o(n^{2/3})$ such that $\mathrm{Cay}(\mathbb{Z}_n \times \mathbb{Z}_n, T)$ has diameter at most 3.*

Proof. For the set T we may take the Cartesian product of two copies of the set S from Lemma 5.1. □

Given our strategy to find a family of groups based on prime numbers, we will also need to use Lemma 3.3 in Chapter 3 to extend our construction to be valid for all degrees.

Now we are ready to describe the main construction. For any odd prime p, we begin with a group H which is the unique non-abelian group of order p^3 with exponent p. This has the form $(\mathbb{Z}_p \times \mathbb{Z}_p) \rtimes \mathbb{Z}_p$. It is well known that the group H can be viewed as the *upper unitriangular* subgroup of $SL(3, p)$, i.e. the subgroup consisting of matrices of the form $\begin{pmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix}$ where a, b, c are arbitrary elements of $GF(p)$. The group G for our Cayley graph will be a direct product of this group with $\mathbb{Z}_2$.

Lemma 5.3. *Let p be an odd prime. Let H be the upper unitriangular subgroup of $SL(3, p)$ and let $G = H \times \mathbb{Z}_2$. Then there is an inverse-closed subset S of G with cardinality $2p + O(p^{2/3})$ such that the Cayley graph $\mathrm{Cay}(G, S)$ has diameter 3, and S contains neither the identity nor the unique involution of G.*

Proof. We construct our generating set S for G as follows. For each $x \in GF(p)^*$ we define the following elements of G.

$$\alpha_x = \left(\begin{pmatrix} 1 & x & x \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, 0 \right); \qquad \beta_x = \left(\begin{pmatrix} 1 & 0 & x \\ 0 & 1 & x \\ 0 & 0 & 1 \end{pmatrix}, 1 \right)$$

Let S_1 be the set consisting of α_x and β_x for all $x \in GF(p)^*$. Notice that S_1 contains neither the identity nor the involution. We now show that all elements of G of the forms $\left(\begin{pmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix}, 0 \right), a \neq 0$ and $\left(\begin{pmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix}, 1 \right), c \neq 0$ may be expressed as a product of at most 3 elements from S_1.

First consider $X = \left(\begin{pmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix}, 0 \right), a \neq 0$. There are three cases to consider. If $b = a + c$ then we choose any $u \notin \{0, c\}$ and then $X = \beta_u \beta_{c-u} \alpha_a$. Otherwise if $b = a + c + ac$ then again we choose $u \notin \{0, c\}$ and this time $X = \alpha_a \beta_u \beta_{c-u}$. Otherwise we let $x = c - (b - c)/a + 1; y = a; z = (b - c)/a - 1$ and then $X = \beta_x \alpha_y \beta_z$.

Now consider $X = \left(\begin{pmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix}, 1 \right), c \neq 0$. Let $x = (b - a)/c - 1; y = c; z = a - (b - a)/c + 1$. If $b = a + c$ then $X = \beta_y \alpha_z$. Otherwise if $b = a + c + ac$ then $X = \alpha_x \beta_y$. Otherwise $X = \alpha_x \beta_y \alpha_z$.

Now we deal with the remaining cases. The elements of the form $\left(\begin{pmatrix} 1 & 0 & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix}, 0 \right)$ form a subgroup of G isomorphic to $\mathbb{Z}_p \times \mathbb{Z}_p$. By Lemma 5.2 there is a set S_2 of size $9p^{2/3} + o(p^{2/3})$ such that each of these can be expressed as a product of at most 3 elements of S_2.

Finally, the elements of the form $\left(\begin{pmatrix} 1 & a & b \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, 1 \right)$ are contained in a subgroup of G isomorphic to $\mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_2$. In a similar way, we can find a set S_3 of size $18p^{2/3} + o(p^{2/3})$ such that each of these can be expressed as a product of at most 3 elements of S_3. Letting $S = S_1 \cup S_2 \cup S_3$ we see that $\mathrm{Cay}(G, S)$ has diameter at most 3 and $|S| = 2p + O(p^{2/3})$ as required. □

The main result now follows.

Theorem 5.4. *In the class of general Cayley graphs,*

$$L^-(3) \geq \frac{1}{4}$$

Proof. The graphs in Lemma 5.3 have order $2p^3$ and degree $2p + O(p^{2/3})$, and satisfy the conditions of Lemma 3.3. □

5.2 A semidirect product construction

5.2.1 Motivation

The results of Section 5.1 provide a useful improvement to the asymptotic bound at diameter 3. For larger diameters, it is possible that other subgroups of matrix groups might yield interesting results. However, it is unlikely that a general construction covering a range of diameters would be possible with this approach.

To find a more general construction, we are inspired by the ideas of Vetrík [73], Macbeth et al [54], Bevan [10] and others. A common strategy of such constructions for a given diameter k is to begin with a k-fold direct product of some group H, and then to permute the coordinate positions in the direct product by means of a semidirect product of H^k by some other group K.

We recall from our definition in Chapter 1 that the semidirect product $G \rtimes_\varphi K$ of two groups G and K has multiplication defined by:

$$(g_1, k_1)(g_2, k_2) = (g_1^{\varphi(k_2)} g_2, k_1 k_2)$$

where the superscript on g_1 indicates the image of g_1 under the automorphism $\varphi(k_2)$ of G.

In a k-fold direct product H^k, any permutation of the k coordinate positions is an automorphism of the group. These automorphisms of H^k form a subgroup N of its full automorphism group, with N isomorphic to the symmetric group S_k. We restrict ourselves in our semidirect products $H^k \rtimes_\varphi K$ to homomorphisms φ into this restricted subgroup N.

Our goal is again to find a lower bound on the quantity $L^-(k)$ for certain fixed values of k, as defined in Section 2.3. To achieve this we first fix a diameter k, and then try to construct an infinite sequence of Cayley graphs of degree sm and asymptotic order $m^k n$ for every $m \geq 2$ and for some fixed constants n, s.

We begin the discussion with an example at diameter 6 which should help clarify the overall method. The example is derived from an original note of Tuite [71], amended to conform to our notation.

5.2.2 Diameter 6 example

Let H be an abelian group of order m. For the purposes of our construction we take $H = \mathbb{Z}_m$ and use additive notation for the group operation. Let $k = 6$ and denote the 6-fold direct product of H by H^6. Let σ be the automorphism of H^6 which maps the element $(x_1, x_2, x_3, x_4, x_5, x_6)$ to $(x_6, x_1, x_2, x_3, x_4, x_5)$. Let K be the group $\mathbb{Z}_{36}$ and let $\varphi : K \to \mathrm{Aut}(H^k)$ be the group homomorphism given by $\varphi(r) = \sigma^r$. Let $G = H^6 \rtimes_\varphi K$.

We write the elements of G in the form $(x_1, x_2, x_3, x_4, x_5, x_6; y)$ where each $x_i \in H$ and $y \in K$. We construct our generating set as follows. For each $x \in H$ we define:

$$\begin{aligned} a(x) &= (0, 0, 0, 0, 0, x; 1) \\ A(x) &= (0, 0, 0, 0, x, 0; -1) \\ b(x) &= (0, x, x, 0, 0, x; 4) \\ B(x) &= (0, x, 0, x, x, 0; -4) \end{aligned}$$

Then the generating set is:

$$X = \bigcup_{x \in H} \{a(x), A(x), b(x), B(x)\}$$

Note that since $a(x)^{-1} = A(-x)$ and $b(x)^{-1} = B(-x)$, the set S is inverse-closed.

We claim that the graph $\mathrm{Cay}(G, X)$ has diameter 6. To do this, it suffices to show that every element of G can be expressed as a product of at most 6 elements of X. For a given $y \in K$, we can express the element $(x_1, x_2, x_3, x_4, x_5, x_6; y) \in G$ via the products in Figure 5.1. For $y = 18 \ldots 35$, we may obtain expressions simply by inverting the appropriate products. Thus $\mathrm{Cay}(G, X)$ has diameter 6 as claimed.

To illustrate the multiplication rules we show here an example from the solution. In the case of $y = 3$ we clam that

$$(x_1, x_2, x_3, x_4, x_5, x_6; 3) = a(x_2)a(x_1)b(x_6)A(x_3 - x_6)A(x_4)A(x_5 - x_6).$$

Expanding the right hand side one step at a time we get the following.

$$\begin{aligned} &a(x_2)a(x_1) \\ &\quad = (0, 0, 0, 0, 0, x_2; 1)(0, 0, 0, 0, 0, x_1; 1) \end{aligned}$$

$$
\begin{aligned}
y = 0 &: a(-x_2 + x_3 + x_5)a(-x_2 + x_3 + x_4)b(x_3) \\
&\quad A(-x_3 + x_6)A(x_1)B(x_2 - x_3) \\
y = 1 &: b(x_3)b(x_2 - x_4)A(-x_2 - x_3 + x_5) \\
&\quad A(-x_3 + x_6)A(x_1 - x_2 + x_4)B(x_4) \\
y = 2 &: b(x_6)a(x_3 - x_4 + x_6)a(x_2) \\
&\quad a(x_1 - x_4)B(x_4 - x_6)A(x_5) \\
y = 3 &: a(x_2)a(x_1)b(x_6) \\
&\quad A(x_3 - x_6)A(x_4)A(x_5 - x_6) \\
y = 4 &: a(-x_1 + x_3)a(x_1 + x_2 - x_4)B(-x_1 + x_4) \\
&\quad a(x_1 - x_4 + x_5)b(x_1)a(x_6) \\
y = 5 &: a(x_2 + x_4 + x_5 - x_6)a(-x_2 + x_3 - x_5)B(x_5) \\
&\quad A(x_1 + x_2 + x_5 - x_6)b(-x_2 - x_5 + x_6)b(x_2) \\
y = 6 &: a(x_5)a(x_4)a(x_3) \\
&\quad a(x_2)a(x_1)a(x_6) \\
y = 7 &: a(x_1 - x_3 + x_6)a(2x_1 - x_3 - x_4 + x_5)b(x_1) \\
&\quad b(-x_1 + x_3)a(x_1 + x_2 - x_4)B(-x_1 + x_4) \\
y = 8 &: a(x_1 + x_3 - x_6)b(x_3)a(x_2) \\
&\quad b(-x_3 + x_6)A(x_3 + x_4 - x_6)A(-x_3 + x_5) \\
y = 9 &: a(x_2 - x_3)a(x_1)a(-x_3 + x_6) \\
&\quad a(x_5)a(x_4)b(x_3) \\
y = 10 &: a(x_3)b(x_2 - x_4)A(-x_2 + x_5) \\
&\quad b(x_4)a(x_1 - x_2 + x_4)a(x_6) \\
y = 11 &: a(x_4)b(-x_1 + x_2)A(-x_3 + x_6) \\
&\quad b(x_1 - x_2 + x_3)b(x_1)A(-2x_1 + x_2 - x_3 + x_5) \\
y = 12 &: a(x_5)a(-x_1 + x_4)a(x_1 + x_3 - x_6) \\
&\quad a(x_1 + x_2 - x_6)b(x_1)b(-x_1 + x_6) \\
y = 13 &: a(-2x_2 + x_3 - x_4 + x_6)a(x_5)b(-x_2 + x_3) \\
&\quad A(x_1 - x_4)b(x_2 - x_3 + x_4)b(x_2) \\
y = 14 &: b(-x_3 + x_6)b(x_3)b(x_3 + x_4 - x_6) \\
&\quad A(-x_1 + x_2 - 3x_3 - x_4 + 2x_6)b(x_1 + x_3 - x_6)A(-x_1 - 2x_3 - x_4 + x_5 + 2x_6) \\
y = 15 &: a(x_2 - x_3 - x_5)a(x_1 - x_4 + x_5)a(-x_3 - x_4 + x_5 + x_6) \\
&\quad b(x_5)b(x_4 - x_5)b(x_3) \\
y = 16 &: a(-2x_2 + x_3 - x_4 + 2x_5 + x_6)b(x_5)b(x_2 + x_4 - x_5 - x_6) \\
&\quad A(x_1 - x_4 - x_5)b(-x_2 + x_5 + x_6)b(x_2 - x_5) \\
y = 17 &: B(x_1 - 2x_2 - x_3 + 2x_5 - x_6)B(-x_2 + x_5)B(x_2 + x_3 - x_5) \\
&\quad a(x_1 - 5x_2 - 2x_3 + x_4 + 4x_5 - 2x_6)B(x_2 - x_5 + x_6)B(-x_1 + 3x_2 + x_3 - 2x_5 + x_6)
\end{aligned}
$$

Figure 5.1: Solution for diameter 6

The multiplication rule is that we rotate the first 6 coordinates of the first term by the final coordinate of the second term, then add. So we get:

$$\begin{aligned} &a(x_2)a(x_1) \\ &\quad = (x_2, 0, 0, 0, 0, x_1; 2) \end{aligned}$$

We continue in this way.

$$\begin{aligned} &a(x_2)a(x_1)b(x_6) \\ &\quad = (x_2, 0, 0, 0, 0, x_1; 2)(0, x_6, x_6, 0, 0, x_6; 4) \\ &\quad = (0, x_6, x_6, x_1, x_2, x_6; 6) \end{aligned}$$

$$\begin{aligned} &a(x_2)a(x_1)b(x_6)A(x_3 - x_6) \\ &\quad = (0, x_6, x_6, x_1, x_2, x_6; 6)(0, 0, 0, 0, x_3 - x_6, 0; -1) \\ &\quad = (x_6, x_6, x_1, x_2, x_3, 0; 5) \end{aligned}$$

$$\begin{aligned} &a(x_2)a(x_1)b(x_6)A(x_3 - x_6)A(x_4) \\ &\quad = (x_6, x_6, x_1, x_2, x_3, 0; 5)(0, 0, 0, 0, x_4, 0; -1) \\ &\quad = (x_6, x_1, x_2, x_3, x_4, x_6; 4) \end{aligned}$$

$$\begin{aligned} &a(x_2)a(x_1)b(x_6)A(x_3 - x_6)A(x_4)A(x_5 - x_6) \\ &\quad = (x_6, x_1, x_2, x_3, x_4, x_6; 4)(0, 0, 0, 0, x_5 - x_6, 0; -1) \\ &\quad = (x_1, x_2, x_3, x_4, x_5, x_6; 3) \end{aligned}$$

Since $|G| = 36m^6$ and $|X| = 4m$, it follows that for every degree d of the form $4m$, there exists a Cayley graph of diameter 6 and order $36d^6/4^6$. To cover the cases $d \equiv 1, 2, 3 \pmod 4$ we may simply add one more involution from G and/or one more pair of mutually inverse elements to our set X. We have therefore proved the following result.

Proposition 5.5. *In the class of Cayley graphs,*

$$L^-(6) \geq \frac{36}{4^6} \approx 0.00878$$

This is, as far as we know, the first specific result for diameter 6 and is an improvement on the bound of $6/3^6 \approx 0.00823$ from [54]. However, the method is

capable of generalisation and we now describe the full construction.

5.2.3 The general construction

We begin by drawing out the key features of the construction in Section 5.2.2. Recall that $H = \mathbb{Z}_m$, $K = \mathbb{Z}_{36}$ and $k = 6$. We define $n = |K|$, so $n = 36$ in our example. Finally, $\varphi : K \to \mathrm{Aut}(H^k)$ is the group homomorphism given by $\varphi(r) = \sigma^r$ and $G = H^k \rtimes_\varphi K$.

Our generating set was constructed as follows. We have a set $S = \{s_1, s_2, s_3, s_4\} = \{1, -1, 4, -4\}$ which is a subset of K of cardinality 4. It can readily be checked that the set S has the property that every element of K can be expressed as a sum of exactly k elements of S. Moreover, the sums satisfy the further restriction that no element of S appears consecutively with its inverse. For example, from the table above for the case $y = 3$ we have $3 = 1 + 1 + 4 - 1 - 1 - 1$.

We have a set $V = \{v_1, v_2, v_3, v_4\} = \{000001, 000010, 011001, 010110\}$ of 4 non-zero 0/1 vectors of length $k = 6$. This set has two important properties. The first is that $v_2 = v_1^{\sigma^{-s_1}}$ and $v_4 = v_1^{\sigma^{-s_3}}$, where as before σ represents a right rotation of one place in the coordinates. This ensures that our generating set defined below will be inverse-closed. The second is that the vectors have been carefully chosen to ensure that our eventual graph will have diameter 6.

For every $x \in H$, we define $v_i(x)$ to be the element of H^k with x in every coordinate position where v_i has a 1, and 0 otherwise. We now define our generating set X to consist of four sets of elements of G as follows.

$$\begin{aligned} a(x) &= (v_1(x); s_1) \quad \text{for all } x \in H \\ A(x) &= (v_2(x); s_2) \quad \text{for all } x \in H \\ b(x) &= (v_3(x); s_3) \quad \text{for all } x \in H \\ B(x) &= (v_4(x); s_4) \quad \text{for all } x \in H \end{aligned}$$

Note that because of the forms of the vectors v_i explained above, we have $a(x)^{-1} = A(-x)$, $b(x)^{-1} = B(-x)$ and so X is an inverse-closed subset of G.

The most awkward part of the process is to determine how to express any possible element of our group G as a product of k of our generators. To determine how this can be done we proceed as follows. We must show, for each $i = 0 \ldots n-1$, that we can express any element of the form $(\mathbf{x}; i)$ as a product of k generators. Since the

generator set is inverse-closed we need only check $i \leq \lfloor \frac{n}{2} \rfloor$. To find products which work we proceed as follows for each such i.

a Find all possible ways in which i can be expressed as a sum of k elements of S (ignoring order in the sum).

b Find all unique ways to order the elements in this sum, say $T = (t_1, t_2, \ldots, t_k)$ with each $t_j \in S$ and $\sum t_j = i$. We insist also that $t_{j+1} \neq -t_j$ for $j = 1 \ldots k-1$.

c For each T, compute the vector U of k numbers chosen from $\{1, 2, 3, 4\}$ such that $t_j = s_{u_j}$ for each j. That is to say, we identify in order those elements of S involved in the sum. At this point we know our product must have the form $(v_{u_1}(y_1); s_{u_1})(v_{u_2}(y_2); s_{u_2}) \cdots (v_{u_k}(y_k); s_{u_k})$ for some $\mathbf{y} = (y_1, y_2, \ldots, y_k)$.

d To determine whether there is a solution we compute the *mapping matrix* M such that $\mathbf{y}M = \mathbf{x}$. If M is invertible over $\mathbb{Z}$ (i.e. it has determinant ± 1) we have found a solution for i, otherwise we proceed with the search.

In the final step, it is easy to see that the mapping matrix M has the following form:

$$M = \begin{pmatrix} v_{u_1}^{\sigma^{r_1}} \\ v_{u_2}^{\sigma^{r_2}} \\ \vdots \\ v_{u_k}^{\sigma^{r_k}} \end{pmatrix}^T ; \qquad r_w = \sum_{j<w} t_j$$

The elements of this construction which can be generalised are as follows.

(i) The target diameter of our Cayley graph could be any $k > 2$.

(ii) The group K could be an arbitrary group of order n rather than being restricted to cyclic groups.

(iii) The size $|S|$ of set of elements of K need not be 4.

(iv) The homomorphism φ in the semidirect product could be any non-trivial homomorphism from our group K to the group of coordinate permutations of H^k.

(v) Our set V of 0/1 vectors could be any set, provided the resulting set of generators is inverse-closed.

It is clear that with the large number of variables, and the relatively complex nature of the construction, some form of automated search for feasible solutions is essential. We outline the search algorithm below. We begin with the following inputs:

- A target diameter k.
- A set size $s = |S|$.
- A target order n for our group K.

Given these parameters, we run the search using a GAP [35] script as follows.

1. Find all groups K of order n from the small groups library.
2. For each K, find (up to conjugacy) all possible homomorphisms φ from K to S_k. (To avoid trivial cases, we consider only homomorphisms whose image has no fixed point.)
3. For each K, find all sets S of size s with the property that any element of K can be written as a product of exactly k elements of S.
4. For each combination of φ and S, find all possible sets $V = \{v_1, \ldots, v_s\}$ of 0/1 vectors of length k, such that, given the elements of S, the resulting generating set will be inverse-closed.

For each viable combination found, we then search for a solution using a modified version of the diameter 6 example. So for each element $i \in K$ we test whether the following procedure succeeds.

(a) Find all ways to express i as a product of k elements of S, say $T = (t_1, t_2, \ldots, t_k)$ with each $t_j \in S$ and $\sum t_j = i$. As before, we insist that $t_{j+1} \neq t_j^{-1}$ for $j = 1 \ldots k-1$.

(b) For each T, compute the vector U of k numbers chosen from $1 \ldots s$ such that $t_j = s_{u_j}$ for each j. So we know our product must have the form $(v_{u_1}(y_1); s_{u_1})(v_{u_2}(y_2); s_{u_2}) \cdots (v_{u_k}(y_k); s_{u_k})$ for some $\mathbf{y} = (y_1, y_2, \ldots, y_k)$.

(c) To determine whether there is a solution we again compute the mapping matrix M such that $\mathbf{y}M = \mathbf{x}$. If M is invertible over $\mathbb{Z}$ we have found a solution for i, otherwise we proceed with the search.

If this procedure finds a solution for all $i \in K$, then our search has yielded a positive result.

Set size s	Group order n	Group K	$L^-(3)$ bound
4	12	$\mathbb{Z}_{12}$	$12/4^3 \approx 0.18750$
5	24	S_4	$24/5^3 \approx 0.19200$
6	48	$(\mathbb{Z}_4 \times \mathbb{Z}_4) \rtimes \mathbb{Z}_3$	$48/6^3 \approx 0.22222$
7	72	$(\mathbb{Z}_2^2 \rtimes \mathbb{Z}_9) \rtimes \mathbb{Z}_2$	$72/7^3 \approx 0.20991$

Table 5.1: Best results for undirected graphs of diameter 3

When a solution has been found, we know that for any m, we can create a Cayley graph of diameter k, order $m^k n$ and degree sm. Thus by the same argument as in the diameter 6 example, we will have proved that in the class of Cayley graphs:

$$L^-(k) \geq \frac{n}{s^k}$$

The object now is to choose the parameters for the search in such a way that we can improve the existing asymptotic bounds. The following sections describe our best results.

5.2.4 Undirected graphs

5.2.4.1 Diameters 2 and 3

For diameter 2, our method will never produce a useful result. This is because our construction requires us to be able to express every element of our group K as a product of k elements chosen from S so that no element follows its inverse in the product. With $k = 2$ this is clearly impossible.

For diameter 3, our best results for set sizes $s = 4, 5, 6, 7$ are shown in Table 5.1. (There were no useful solutions with $s = 3$.) The best existing published result is by Vetrík [73] giving $L^-(3) \geq \frac{3}{16}$. Although our results improve on that, we are unable to do better than the specific diameter 3 construction from Section 5.1 above.

5.2.4.2 Diameter 4

For diameter 4, the increasing size of the search space means that we were only able to search for solutions with set sizes of 3, 4 and 5. The results are summarised in Table 5.2. The best existing published result is again by Vetrík [73] giving $L^-(4) \geq \frac{32}{5^4} \approx 0.05120$. For set sizes 4 and 5, we obtain results better than that bound.

We note that in contrast to the diameter 6 example construction above which used a cyclic group K, the groups found by the computer search are not at all obvious and the combination of group K, homomorphism φ, set S and vectors V lead to a solution

Set size s	Group order n	Group K	$L^-(4)$ bound
3	4	$\mathbb{Z}_4$	$4/3^4 \approx 0.04938$
4	24	S_4	$24/4^4 \approx 0.09375$
5	60	$\mathbb{Z}_{15} \rtimes \mathbb{Z}_4$	$60/5^4 \approx 0.09600$

Table 5.2: Best results for undirected graphs of diameter 4

Set size s	Group order n	Group K	$L^-(5)$ bound
3	6	S_3	$6/3^5 \approx 0.02469$
4	60	A_5	$60/4^5 \approx 0.05859$

Table 5.3: Best results for undirected graphs of diameter 5

which is complex and lengthy to tabulate. For reasons of brevity therefore we omit all the full solutions here. The simplest solution to describe, although not the one yielding the largest value, uses a set of size 4. In this case we are fortunate that the group is S_4 and the homomorphism φ is simply the identity mapping. We therefore illustrate the results by tabulating this solution below in the same format as our diameter 6 example in Figure 5.1.

5.2.4.3 Diameter 5

For diameter 5 we were able to search for solutions with set sizes of 3 and 4, with the results summarised in Table 5.3. As before, the best existing published result is by Vetrík [73] giving $L^-(5) \geq \frac{25}{4^5} \approx 0.02441$. For set size 3, we have a marginal improvement and at set size 4 our best solution is a substantial increase.

5.2.4.4 Diameters 6 and 7

For diameters 6 and 7 we were again able to search for solutions with set sizes of 3 and 4, with the results summarised in Tables 5.4 and 5.5. There are no specific published results at diameters 6 and 7. The best available published result comes from the general construction of Macbeth, Šiagiová, Širáň and Vetrík [54] which yields $L^-(6) \geq \frac{6}{3^6} \approx 0.00823$ and $L^-(7) \geq \frac{7}{3^7} \approx 0.00320$.

Recall that our diameter 6 example with a set size of 4 already yielded an improvement to ≈ 0.00878, but with the aid of the computer search we are able to more than double this figure. At diameter 7 our best result is now more than three times that obtained by the published general construction.

$$
\begin{aligned}
K &= S_4 \\
S &= \{(2\ 3\ 4), (2\ 4\ 3), (3\ 4), (1\ 2)\} \\
V &= \{1010, 1100, 0100, 1110\} \\
a(x) &= (x, 0, x, 0; (2\ 3\ 4)) \\
A(x) &= (x, x, 0, 0; (2\ 4\ 3)) \\
b(x) &= (0, x, 0, 0; (3\ 4)) \\
c(x) &= (x, x, x, 0; (1\ 2))
\end{aligned}
$$

$$
\begin{aligned}
y = ():&\quad b(x_2 - x_3 - x_4)c(x_4)b(x_1 - x_3 - x_4)c(x_3) \\
y = (1\ 2):&\quad a(x_2 - x_3 - x_4)a(x_4)a(-x_1 + x_2 - x_4)c(x_1 - x_2 + x_3 + x_4) \\
y = (1\ 3):&\quad b(x_2 - x_3)c(x_4)a(-x_1 + x_3 - x_4)c(x_1) \\
y = (1\ 4):&\quad b(-x_1 + x_2)c(x_4)A(x_1 - x_3 - x_4)c(x_3) \\
y = (2\ 3):&\quad A(x_2 - x_3 - x_4)c(x_4)b(x_1 - x_2)c(x_3) \\
y = (2\ 4):&\quad a(x_4)a(x_2)b(-x_1 + x_2 + x_3 + x_4)a(x_1 - x_2 - x_4) \\
y = (3\ 4):&\quad b(-x_1 + x_2 + x_3 + x_4)a(x_1 - x_3 - x_4)a(x_4)a(x_3) \\
y = (1\ 2)(3\ 4):&\quad A(-x_1 + x_2)b(x_1 - x_2 + x_4)A(x_1 - x_3)c(x_3) \\
y = (1\ 3)(2\ 4):&\quad a(-x_2 + x_3)c(x_2 - x_3 + x_4)a(-x_1 + x_3 - x_4)c(x_1) \\
y = (1\ 4)(2\ 3):&\quad b(x_3)A(x_1 - x_2)c(-x_1 + x_2 + x_4)A(x_1 - x_4) \\
y = (1\ 2\ 3):&\quad a(-x_1 + x_2)b(x_4)A(x_2 - x_3)c(x_1 - x_2 + x_3) \\
y = (1\ 3\ 2):&\quad b(x_1 - x_2 - x_4)c(x_3)A(-x_3 + x_4)A(x_2) \\
y = (1\ 2\ 4):&\quad b(-x_2 + x_3 + x_4)a(x_2 - x_3)a(-x_1 + x_3)c(x_1) \\
y = (1\ 4\ 2):&\quad b(x_1 - x_3 - x_4)c(x_2)a(-x_2 + x_4)a(x_3) \\
y = (1\ 3\ 4):&\quad a(x_2)b(x_1 + x_2 - x_3)c(x_4)a(-x_2 + x_3 - x_4) \\
y = (1\ 4\ 3):&\quad A(x_1 - x_2 - x_3)c(x_3)b(-x_1 + x_2 + x_4)A(x_2) \\
y = (2\ 3\ 4):&\quad b(-x_1 + x_2 + x_3)a(x_1 - x_2)b(x_4)A(x_2) \\
y = (2\ 4\ 3):&\quad b(x_4)A(x_3)A(x_1 - x_3)b(-x_1 + x_2 + x_3) \\
y = (1\ 2\ 3\ 4):&\quad b(x_3)A(x_1 - x_4)c(x_4)b(-x_1 + x_2) \\
y = (1\ 4\ 3\ 2):&\quad b(x_1 - x_2 - x_3)c(x_3)b(-x_3 + x_4)A(x_2) \\
y = (1\ 2\ 4\ 3):&\quad b(-x_2 + x_3 + x_4)a(x_2 - x_3)b(x_1 - x_3)c(x_3) \\
y = (1\ 3\ 4\ 2):&\quad c(x_4)b(x_3 - x_4)A(x_1 - x_4)b(-x_1 + x_2 + x_4) \\
y = (1\ 3\ 2\ 4):&\quad a(-x_1 + x_2 + x_4)c(x_1 - x_2 + x_3 - x_4)A(-x_3 + x_4)A(x_2) \\
y = (1\ 4\ 2\ 3):&\quad a(x_3)a(-x_1 + x_2)c(x_1 - x_2 - x_3 + x_4)A(x_2 + x_3 - x_4)
\end{aligned}
$$

Figure 5.2: Solution for diameter 4

Set size s	Group order n	Group K	$L^-(6)$ bound
3	12	A_4	$12/3^6 \approx 0.01646$
4	78	$\mathbb{Z}_2 \times (\mathbb{Z}_{13} \rtimes \mathbb{Z}_3)$	$78/4^6 \approx 0.01904$

Table 5.4: Best results for undirected graphs of diameter 6

Set size s	Group order n	Group K	$L^-(7)$ bound
3	14	D_{14}	$14/3^7 \approx 0.00640$
4	168	$\mathbb{Z}_8 \times (\mathbb{Z}_7 \rtimes \mathbb{Z}_3)$	$168/4^7 \approx 0.01025$

Table 5.5: Best results for undirected graphs of diameter 7

5.2.4.5 Summary

We collect the results above into a single theorem.

Theorem 5.6. *In the class of undirected Cayley graphs,*

$$L^-(4) \geq \frac{60}{5^4} \approx 0.09600$$
$$L^-(5) \geq \frac{60}{4^5} \approx 0.05859$$
$$L^-(6) \geq \frac{78}{4^6} \approx 0.01904$$
$$L^-(7) \geq \frac{168}{4^7} \approx 0.01025$$

5.2.5 Directed graphs

The search method we used for undirected Cayley graphs can be modified to search for Cayley digraphs. The only substantial difference is that our generating set X need not be inverse-closed. This has two major consequences for the search:

- The set S of elements of K need not be inverse-closed.
- The set V of 0/1-vectors is not restricted by the requirement that the resulting generating set be inverse-closed.

These consequences taken together result in a substantial increase in the search space for a given set of parameters. Due to this effect, we were only able to search a limited range of set sizes (2, 3 and 4) for diameters 3, 4 and 5. The best results are summarised in Table 5.6.

In general as one would expect, removing the restriction on generating sets results in bounds which are much better than the corresponding undirected bounds. As far as we know, there are no better published results.

The current table has the curious feature that the best result we were able to find for diameter 5 is better than that for diameter 4. This is counter-intuitive, but may simply be a consequence of the restricted space that we were able to search.

Diameter k	Set size s	Group order n	Group K	$L^-(k)$ bound
3	4	48	$\mathbb{Z}_2 \times S_4$	$48/4^3 \approx 0.75000$
4	3	36	$\mathbb{Z}_3 \times A_4$	$36/3^4 \approx 0.44444$
5	3	120	S_5	$120/3^5 \approx 0.49382$

Table 5.6: Best results for directed graphs

We summarise these results again into a single theorem.

Theorem 5.7. *In the class of directed Cayley graphs,*

$$L^-(3) \geq \frac{48}{4^3} \approx 0.75000$$
$$L^-(4) \geq \frac{36}{3^4} \approx 0.44444$$
$$L^-(5) \geq \frac{120}{3^5} \approx 0.49382$$

Our new constructions are able to better the directed graph bounds of Vetrík [72] at diameters 3, 4 and 5.

5.3 Diameter two revisited

Although our main focus in this chapter has been on Cayley graphs of diameters 3 and above, we conclude with a small incremental improvement to the asymptotic bound for general Cayley graphs of diameter 2. Abas shows in [2] that $L^-(2) \geq 0.684$. To do this, he uses the "prime gaps" technique using a method based on Cullinan and Hajir [25] and Ramaré and Rumely [63] to show for degrees $d > 360756$, his construction yields a lower bound of 0.684.

Because in our version we are interested in the asymptotic version of the bounds, we are able to tolerate a much larger validity bound on d. We can then use the tables in [63] to get to a slightly larger bound.

This leads to the following small improvement on the bound of Abas.

Theorem 5.8. *In the class of undirected Cayley graphs, for all sufficiently large d we have $n(d, 2) \geq 0.68762d^2$. Thus*

$$L^-(2) \geq 0.68762.$$

Proof. We recall Abas' proof in [2]. His construction works for values of d of the form $17p - 1$, where p is a prime such that $p \equiv 1 \pmod{10}$. He shows that for $n > 10^{10}$

there is a prime p in this congruence class in the interval $(n, (1+\delta)n)$ where $\delta = 2\epsilon/(1-\epsilon)$ and the value of $\epsilon = 0.002785$ is taken from [63, Table 1] in the row for $k = 10$ and column for 10^{10}. This yields $\delta = 0.00558556$ and since his construction gives an asymptotic graph order of $\frac{200}{289}d^2$ for degrees d of the form $17p - 1$, he concludes that an asymptotic order of $\frac{200}{289(1+\delta)^2}d^2$ is valid for all sufficiently large d.

We now replace the value of ϵ used by Abas with the one from the 10^{100} column, giving $\epsilon = 0.001606$ and hence $\delta \leq 0.003207$. This means that for d sufficiently large, $n(d, 2) \geq \frac{200}{289(1+\delta)^2}d^2 \geq 0.68762d^2$. □

Chapter 6

The degree-diameter problem for mixed graphs

As we noted in Chapter 2, the degree-diameter problem for mixed graphs has received much attention in recent years. Despite this interest, there remain a large number of open questions which appear to be more resistant to progress than their counterparts in the undirected and undirected problems. In Section 6.2 we will discuss the surprisingly large number of open cases in which the existence or otherwise of a Moore graph is unknown. However, we begin with one of the most basic ideas in the degree-diameter problem: the Moore bound itself.

6.1 The mixed Moore bound

6.1.1 Introduction

In this chapter we are concerned with the *mixed* or *partially directed* degree-diameter problem, in which we allow some of the edges in our graph to be directed and some undirected. We conform to the most usual notation in the literature, so that the maximum undirected degree of a vertex (the number of undirected edges incident to it) is denoted by r. The maximum directed degree is taken to mean the maximum number of out-arcs from any vertex and is denoted by z. As usual we denote the diameter of a graph by k.

To bound the maximum possible number of vertices, the approach is to consider a spanning tree rooted at some arbitrary vertex. It is not difficult to see that maximality is only achieved when each vertex has a unique parent at the previous level in the tree, and the maximum possible number of neighbours at the next level. Figure 6.1 shows such a tree for the case $z = 3, r = 3, k = 2$.

Note that this Moore bound is only attained in a very small number of known cases. Nguyen, Miller and Gimbert [60] show that no graphs attaining the bound exist if the diameter $k \geq 3$. For $k = 2$, the known examples [57] are a family of Kautz graphs with $r = 1, z \geq 1$ and a graph of Bosák with $r = 3, z = 1$. Recently, Jørgensen [41] has discovered a pair of graphs with $r = 3, z = 7$.

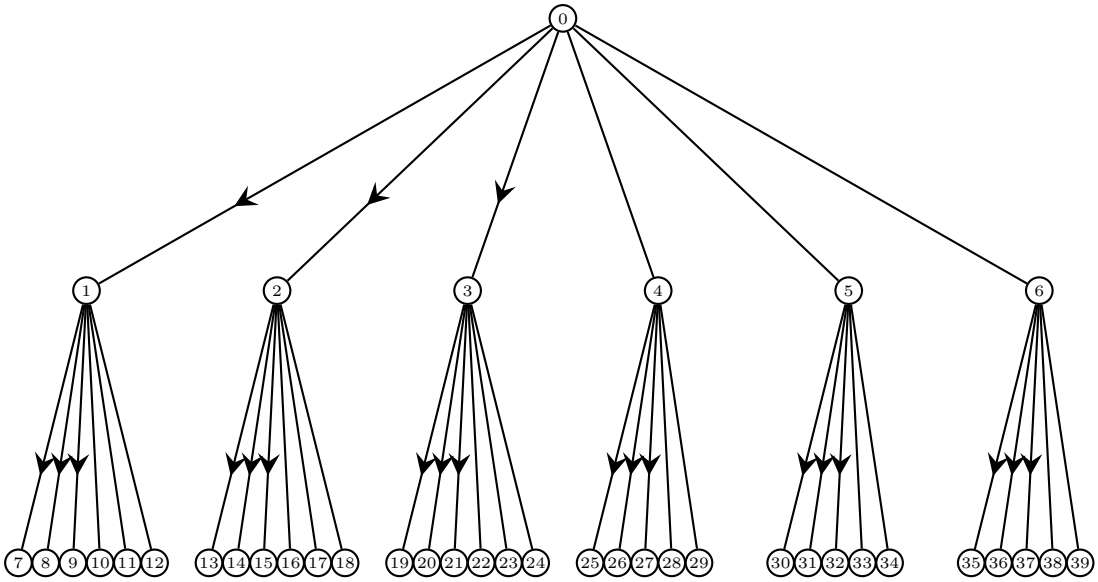

Figure 6.1: The labelled Moore tree for $z = 3, r = 3, k = 2$

In [59] the general Moore bound for the largest possible order of a graph with parameters z, r, k is given as

$$M_{z,r,k} = 1 + (z+r) + z(z+r) + r(z+r-1) + \ldots + z(z+r)^{k-1} + r(z+r-1)^{k-1} \quad (6.1)$$

It seems that this formula may have been extrapolated from the expressions for graphs of small diameter. However, it turns out that for diameters greater than 3 the formula in 6.1 is not correct. In this section we develop a corrected formula for the Moore bound, and show that for all diameters greater than 3 this is strictly smaller than the bound stated in [59]. We follow the structure of our published paper [19].

6.1.2 Revised Moore Bound

Theorem 6.1. *Let $M_{z,r,k}$ denote the largest possible number of vertices in a mixed graph of diameter k, maximum directed degree z and maximum undirected degree r. Then:*

$$M_{z,r,k} = A\frac{u_1^{k+1} - 1}{u_1 - 1} + B\frac{u_2^{k+1} - 1}{u_2 - 1} \quad (6.2)$$

where:

$$
\begin{aligned}
v &= (z+r)^2 + 2(z-r) + 1 \\
u_1 &= \frac{z+r-1-\sqrt{v}}{2} \\
u_2 &= \frac{z+r-1+\sqrt{v}}{2} \\
A &= \frac{\sqrt{v}-(z+r+1)}{2\sqrt{v}} \\
B &= \frac{\sqrt{v}+(z+r+1)}{2\sqrt{v}}
\end{aligned}
$$

To prove the formula, we count vertices in the spanning tree by fixing an arbitrary vertex w and consider the distance partition from w. Denote by L_j the maximum possible number of vertices in the graph at distance j from w.

Lemma 6.2. *The maximum number of vertices in the distance partition satisfies the recurrence*

$$L_j = (z+r-1)L_{j-1} + zL_{j-2}; \quad L_0 = 1; \quad L_1 = z+r$$

Proof. Clearly $L_0 = 1, L_1 = z+r$. For $j \geq 2$ we proceed inductively. The key observation is that in a maximal graph, a vertex at level $j-1$ has exactly one parent at level $j-2$, but the number of its children at level j depends on whether the edge from its parent is undirected or directed. If the edge is undirected then the vertex has at most $z+r-1$ children, and if it was directed then the vertex has at most $z+r$ children, i.e. one more. Since the number of vertices at level $j-1$ with a directed edge from their parent is at most zL_{j-2}, the recurrence follows. □

Proof of Theorem 6.1. Clearly $M_{z,r,k} = \sum_{j=0}^{k} L_j$. To find an explicit form for L_j we solve the second-order homogeneous recurrence defined by Lemma 6.2. The characteristic equation of the recurrence system is $u^2 + (1-z-r)u - z = 0$. This has roots u_1 and u_2 as defined in the theorem. Since $v = (z+r-1)^2 + 4z$, in all non-degenerate cases $v > 0$ and so the roots are real and distinct. From the general theory of second-order recurrences, the general solution of the system is $L_j = Au_1^j + Bu_2^j$ where A, B are constants defined by the initial conditions. Elementary algebraic manipulation gives the values of A, B as defined in the theorem.

Since $M_{z,r,k} = \sum_{j=0}^{k} L_j$, summing the geometric series for the L_j gives the result

$$M_{z,r,k} = A\frac{u_1^{k+1}-1}{u_1-1} + B\frac{u_2^{k+1}-1}{u_2-1}$$

□

Although the closed form solution (6.2) looks rather more complex than the old version, we can show via straightforward algebraic manipulation that it generalises both the undirected and directed formulae.

Proposition 6.3.

(a) Setting $z = 0$ in Equation (6.2) recovers the undirected Moore bound (2.1).

(b) Setting $r = 0$ in Equation (6.2) recovers the directed Moore bound (2.2).

At first glance the formula (6.2) offers little insight into the behaviour of the bound as k increases. However we can obtain a relatively straightforward estimate of its asymptotic behaviour.

Proposition 6.4. *Suppose $r > 0$. In the notation of Theorem 6.1, for sufficiently large k, $M_{z,r,k}$ is the nearest integer to $B\frac{u_2^{k+1}-1}{u_2-1} - \frac{A}{u_1-1}$.*

Proof. It suffices to show that $|u_1| < 1$.

Now $2\frac{\partial u_1}{\partial z} = 1 - \frac{z+r+1}{\sqrt{(z+r)^2+2(z-r)+1}}$. Since $r > 0$, it follows that $(z+r+1)^2 > (z+r)^2 - 2(z+r) + 1$ and so $\frac{\partial u_1}{\partial z} < 0$. So for any fixed $r > 0$, u_1 is strictly decreasing as z increases. When $z = 0, u_1 = 0$ and for any z we have $u_1 > -1$ since $z + r + 1 - \sqrt{(z+r)^2+2(z-r)+1} > 0$. Thus $0 \geq u_1 > -1$ for any $r > 0$ and any $z \geq 0$. □

6.2 Mixed Cayley Moore graphs

6.2.1 Introduction

In this section we discuss the possible existence of Moore graphs in the mixed problem. As noted above, Nguyen, Miller and Gimbert [60] showed in 2007 that no

mixed Moore graph can exist for diameters greater than 2. The validity of this result is unaffected by the correction to the Moore bound in the previous section.

In the diameter 2 case, the Moore bound for graphs with undirected degree r and directed degree z is $(r+z)^2+z+1$. In 1979, Bosák [15] derived (using a modification of the spectral method used by Hoffman and Singleton in the undirected case) a numerical constraint on the sets of parameters r, z for which a mixed Moore graph of diameter 2 can exist. Bosák's condition is that $r = (c^2+3)/4$ for some odd integer c dividing $(4z-3)(4z+5)$.

We can see (Table 6.1) that the range of feasible pairs (r, z) for which a Moore graph can exist is quite limited. Firstly, by Bosák's condition, r is always odd. In the case $r = 1$, it is immediate that any positive integer z yields a feasible pair. For $r = 3$ we have $c = 3$ and so we must have $z \equiv 0, 1 \pmod 3$. We have no solution with $r = 5$, then for $r = 7$ we have $c = 5$ so $z \equiv 0, 2 \pmod 5$. In a similar way, the next possible values of r are $13, 21, 31, 43, 57, \ldots$ and for each value of r, z is constrained to two congruence classes modulo c.

Existence or otherwise of these graphs has only been determined in some special cases. For $r = 1$, Moore graphs always exist by the following construction.

The *Kautz digraphs* [42] $Ka(d, 2)$ are a family of mixed Moore graphs of diameter 2, directed degree $z = d-1$ and undirected degree $r = 1$. The vertices are the words ab of length 2 over an alphabet of $d+1$ letters where we insist $a \neq b$. So there are $d(d+1) = (r+z)^2+z+1$ vertices. There is a directed edge from ab to bc for all of the d eligible values of c. The edge from ab to ba can be considered as the undirected edge since the reverse edge also exists. All other edges from ab are purely directed.

The graph has diameter 2 since there is a path $ab \to xy$ of length 1 if $x = b$ and $ab \to bx \to xy$ of length 2 if $x \neq b$. An example in the case $d = 2$ is shown in Figure 6.2.

The Kautz digraphs $Ka(d, 2)$ are not Cayley graphs for all values of d, and in fact they turn out to be Cayley graphs precisely when $d+1$ is a prime power (see for example [18]). Until very recently, these graphs and a single further example of Bosák with parameters $r = 3, z = 1$ (and hence order 18) were the only known mixed Moore graphs. (See the survey paper [57] for more on these known graphs.)

Recently, Jørgensen [41] has reported a pair of graphs with $r = 3, z = 7$ and hence order 108. These graphs are interesting because they are Cayley graphs (as indeed is Bosák's graph of order 18). The two graphs are in fact a transpose pair, where one

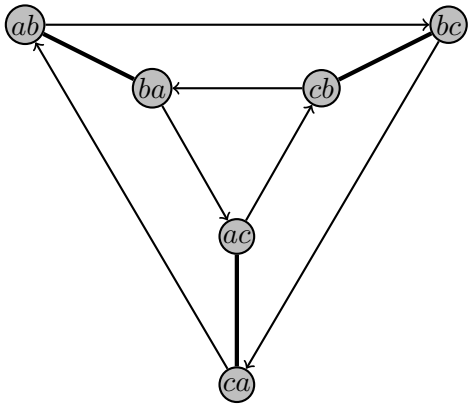

Figure 6.2: The Kautz digraph $Ka(2,2)$

Undirected degree r	Directed degree z	Order n	Existence
1	any	z^2+3z+2	Yes [42]
3	1	18	Yes [15]
	3	40	No [50]
	4	54	No [50]
	6	88	Unknown
	7	108	Yes [41]
	$0,1 \pmod 3$	$z^2+7z+10$	Unknown
7	2	84	No [50]
	5	150	Unknown
	7	204	Unknown
	$0,2 \pmod 5$	$z^2+15z+50$	Unknown
13	$4,6 \pmod 7$	$z^2+27z+170$	Unknown
21	$1,3 \pmod 9$	$z^2+43z+442$	Unknown

Table 6.1: Feasible values for mixed Moore graphs up to $r = 21$

graph is obtained from the other by reversing the direction of the directed arcs.

On the negative side, no simple combinatorial argument has yet been found to rule out any feasible parameter pairs satisfying Basák's condition. For small graphs, an exhaustive computational approach is now becoming feasible with advances in CPU power and algorithms. Very recently, López, Miret and Fernández [50] have used computational techniques to show that there are no mixed Moore graphs at orders 40, 54 or 84.

It seems unlikely that brute-force algorithms will take us much further in the table. Inspired by Jørgensen's result and the fact that the Bosák graph of order 18 is also a Cayley graph, we describe a search algorithm for further examples of mixed Moore Cayley graphs.

6.2.2 The algorithm

Given a feasible pair r, z, we wish to find a group G and a set $S \subseteq G$ such that the graph $\mathrm{Cay}(G, S)$ has order $n = (r+z)^2 + z + 1$, undirected degree r, directed degree z and diameter 2. For ease of explanation we split S into the undirected generators S_1 and the directed generators S_2. Then $|S_1| = r, |S_2| = z, S_1 = S_1^{-1}, S_2 \cap S_2^{-1} = \emptyset$.

As we have seen before, the naive approach of simply testing all possible sets S very quickly becomes computationally infeasible. Our strategy therefore is to look for properties of Moore graphs and corresponding properties of Cayley graphs which will allow us to reduce the search space. We begin with some elementary yet useful properties of mixed Moore graphs.

Proposition 6.5. *Let Γ be a mixed Moore graph of diameter 2, undirected degree r and directed degree z.*

(i) If $u, v \in V(\Gamma)$ are distinct vertices then there is one and only one path of length 1 or 2 from u to v.

(ii) Γ contains no undirected cycle of length 3 or 4.

(iii) Γ is totally regular.

(iv) Every arc in Γ is contained in exactly one directed 3-cycle.

Proof. Item *(i)* follows immediately from the counting argument deriving the Moore bound by considering the spanning tree of Γ rooted at u. Item *(ii)* is a consequence of *(i)*. Item *(iii)* was proved by Bosák [15].

To see why *(iv)* is true, consider a vertex $u \in V(\Gamma)$. Then u has z directed out-neighbours $v_1, \ldots, v_z$. Since Γ is totally regular, u must have z directed in-neighbours $w_1, \ldots w_z$. These cannot be at distance 1 from u, so each w_i is reached by a path of length 2 from u. There can be no undirected edges in any of these paths, since that would lead to the end vertices of such an edge violating *(i)*. So each w_i is reached by a directed path of length 2 from u passing through some v_j. These v_j must be distinct, since if any were repeated it would have two paths of length 2 to u. Thus every arc $u \to v_j$ emanating from u lies in the unique directed triangle $u \to v_j \to w_i \to u$. □

Now we can use these properties to develop constraints on our generating set $S = S_1 \cup S_2$ to narrow the search for mixed Moore Cayley graphs.

Proposition 6.6. *Let Γ be a mixed Moore graph of diameter 2, undirected degree r and directed degree z. Suppose that $\Gamma \cong \mathrm{Cay}(G, S)$ where G is a group of order $n = (r+z)^2 + z + 1$ and the generating set S consists of undirected generators S_1 and directed generators S_2. Then:*

(i) No element of S_1 has order 3 or 4.

(ii) No element of S_2 is an involution.

(iii) No pair of elements in S_1 has a product of order 2.

(iv) No two distinct elements of S commute, apart from the inverse pairs in S_1.

(v) S is product-free (that is, $S \cap SS = \emptyset$).

(vi) All non-identity products of two elements of S are unique.

(vii) The elements of S_2 are of two types:

1. *Elements of order 3*
2. *Triples of distinct elements $\{a, b, c\}$, each of order at least 4, such that $(ab)^{-1} = c$*

Proof. These facts follow immediately from the properties of the graph and Proposition 6.5. □

We note that the conditions of Proposition 6.6(v) and (vi) must also hold for any subset of S. This motivates the following definition.

Let $T \subseteq G$ with $T = T_1 \cup T_2, T_1 = T_1^{-1}, T_2 \cap T_2^{-1} = \emptyset, |T_1| = r', |T_2| = z'$. Define $P(T) = |\{1\} \cup T \cup TT|$. We say T is a *feasible* subset of generators if $P(T) = (z' + r')^2 + z' + 1$.

We have two further ways to reduce the number of sets S we need to search for a given group G. The first is the same idea as we used in Section 3.7.1 which is that if ϕ is an automorphism of the group G, then $\mathrm{Cay}(G,S) \cong \mathrm{Cay}(G,\phi(S))$. So we need not consider all possible sets – only orbit representatives under the action of $\mathrm{Aut}(G)$.

The second idea is to exploit the fact that all mixed Moore graphs must have even order (a consequence of Bosák's condition). So a suitable group G for a Cayley graph must have even order, and may in many cases have an index 2 subgroup.

Proposition 6.7. *Let Γ be a mixed Moore graph of diameter 2, undirected degree r and directed degree z. Suppose that $\Gamma \cong \mathrm{Cay}(G,S)$ where G is a group of order $n = (r+z)^2 + z + 1$ and the generating set S consists of undirected generators S_1 and directed generators S_2. Suppose further that G admits an index 2 subgroup H and that $|S_1 \cap H| = s_1$ and $|S_2 \cap H| = s_2$. Then:*

$$s_1 + s_2 = \frac{2(z+r) - 1 \pm \sqrt{4r-3}}{4}$$

Proof. We know each non-identity element of H can be expressed uniquely as a product of 1 or 2 elements of S. We count these products. Firstly, there are $s_1 + s_2$ elements of $S \cap H$. Any other element is either a product of 2 elements of $S \cap H$ or 2 elements of $S \cap (G \setminus H)$. In the first case there are $s_1(s_1-1) + 2s_1s_2 + s_2^2$ possibilities. In the second case there are $(r-s_1)(r-s_1-1) + 2(r-s_1)(z-s_2) + (z-s_2)^2$. Writing $s = s_1 + s_2$ we see following some manipulation that the total number of elements of H which we can write as a product of 0, 1 or 2 elements of S is $2s^2 + s(1-2(r+z)) + (r+z)^2 - r + 1$. But H is an index 2 subgroup and so contains exactly $((r+z)^2 + z + 1)/2$ elements. Solving this quadratic equation for s yields the stated result. □

It might be thought that this provides a very strong condition, since the expression for $s_1 + s_2$ must clearly give an integer result. However, it is interesting that Bosák's condition on allowable values of r, z means that this expression always gives one integer solution for $s_1 + s_2$. Nevertheless, the condition does give a useful way to cut down the search space when we have an index 2 subgroup H, since it precisely determines the overall split of generators between H and $G \setminus H$. In addition, we have a useful corollary allowing us to exclude some groups from consideration entirely.

Corollary 6.8. *Suppose Γ and G are as in the statement of Proposition 6.7. Then if $2(z+r) - \sqrt{4r-3} > 9$ then G cannot contain an index 2 abelian subgroup H.*

Proof. If H is an index 2 abelian subgroup of G, then if $2(z+r) - \sqrt{4r-3} > 9$, by

Proposition 6.7 the generating set S contains more than 2 elements of H. This is contrary to Proposition 6.6(iv). □

We can now describe the search algorithm. Given a feasible pair z, r we use a `GAP` [35] script.

1. Find all groups G of order $n = (z + r)^2 + z + 1$.
2. If G has an abelian index 2 subgroup, ignore it.
3. Compute the list U of orbit representatives of all inverse-closed sets A of size r such that $|AA| = r(r - 1) + 1$.
4. If G admits an index 2 subgroup H, delete any infeasible sets from U.
5. Compute the list D of all inverse-free sets $B = \{a, b, (ab)^{-1}\}$ such that $|B \cup BB| = 12$.
6. Try to extend each $S \in U$ by adding directed generators of order 3 or triples from D until we have added z generators.

6.2.3 Search results

Results of the search on feasible orders up to 200 are in Table 6.2. For completeness the case $r = 1$ is included. As explained above, we know there is a unique Moore graph with $r = 1$ for every $z \geq 1$, but these are Cayley only if $q = z + 2$ is a prime power. The algorithm reproduces all the known Cayley Moore graphs and confirms that there are no more examples below order 200.

We then continued the search for feasible orders up to 500. The results are in Table 6.3. The algorithm was unable to complete the search at order 486 due to the large numbers of groups and the increasing search space. However, there are no more examples at any of the other feasible orders up to 485.

We summarise these results as tabulated in Tables 6.2 and 6.3 in a theorem.

Theorem 6.9. *Up to order 485, the only mixed Moore Cayley graphs of undirected degree r, directed degree z and diameter 2 are as follows.*

- $r = 1$ *and* $z \leq 20$ *where* $z + 2$ *is a prime power (Kautz graphs).*
- $r = 3$ *and* $z = 1$ *(Bosák's graph).*
- $r = 3$ *and* $z = 7$ *(the two graphs of Jørgensen).*

n	r	z	Graphs
18	3	1	1
40	3	3	0
54	3	4	0
84	7	2	0
88	3	6	0
108	3	7	2
150	7	5	0
154	3	9	0
180	3	10	0

n	r	z	Graphs
6	1	1	1
12	1	2	1
20	1	3	1
30	1	4	0
42	1	5	1
56	1	6	1
72	1	7	1
90	1	8	0
110	1	9	1
132	1	10	0
156	1	11	1
182	1	12	0

Table 6.2: Cayley Moore graphs up to order 200

n	r	z	Graphs
204	7	7	0
238	3	12	0
270	3	13	0
294	13	4	0
300	7	10	0
340	3	15	0
368	13	6	0
374	7	12	0
378	3	16	0
460	3	18	0
486	21	1	?

n	r	z	Graphs
210	1	13	0
240	1	14	1
272	1	15	1
306	1	16	0
342	1	17	1
380	1	18	0
420	1	19	0
462	1	20	0

Table 6.3: Cayley Moore graphs from order 200 to 500

Chapter 7

The degree-girth problem

7.1 Background

In the degree-girth problem, the objective is to find the smallest possible order of a graph with given girth g and having each vertex of degree d. To avoid trivialities, we restrict attention to the case $d \geq 3$. In a similar way to the Moore bound in the degree-diameter problem, a counting argument based on a spanning tree rooted at a given vertex yields a lower bound (also called the Moore bound) on the order of a d-regular graph of girth g:

$$M(d,g) = \begin{cases} \dfrac{d(d-1)^{(g-1)/2}-2}{d-2} & \text{if } g \text{ is even} \\ \dfrac{2(d-1)^{g/2}-2}{d-2} & \text{if } g \text{ is odd} \end{cases}$$

Graphs attaining the lower bound above are also called Moore graphs. While they are not quite as rare as their counterparts in the diameter problem, the values of the parameters d, g for which they may exist are still very restricted and the following are the only possibilities:

- For $d = 2$, the cycle graph C_g is a Moore graph for any $g \geq 3$.
- For $g = 3$ and $g = 4$, the complete graphs K_{d+1} and complete bipartite graphs $K_{d,d}$ are respectively d-regular Moore graphs for any $d \geq 3$.
- For $g = 5$, the Moore graphs correspond exactly to the Moore graphs of diameter 2 in the diameter problem. Thus the known graphs are the 5-cycle, the Petersen graph, the Hoffman-Singleton graph and possibly an unknown graph of degree 57.
- For $g = 6$, $g = 8$ and $g = 12$, the Moore graphs are precisely the incidence graphs of certain generalised polygons — projective planes for girth 6, generalised quadrangles for girth 8 and generalised hexagons for girth 12. In all three cases, these graphs are known to exist for degrees d such that $d - 1$ is a prime power.

Given the scarcity of Moore graphs, progress can still be made in a number of directions, for example:

- Fix a particular girth g and degree d and try to find the smallest possible d-regular graph of girth g.
- Fix a degree d and find an infinite family of graphs with "good" asymptotic order as we increase the girth g.

We focus here on the second approach. In this problem, we seek a family of graphs $\mathcal{G}$ with the property that there exists some $\gamma > 0$ such that for any $G \in \mathcal{G}$,

$$\text{girth}(G) \geq \gamma \log_{d-1}(|G|)$$

From the Moore bound we know that γ is at most 2, but there is no known family of graphs attaining the bound. The best asymptotic families to date are based on constructions of Lubotzky, Phillips and Sarnak [51], and of Lazebnik, Ustimenko and Woldar [45, 46, 43, 47]. Both of these achieve a value of $\gamma = 4/3$ for certain values of d. Despite the ages of these constructions, no better families have been found which more closely approach the upper bound of 2. For more information on families of graph of large girth, see the survey paper by Exoo and Jajcay [31] and the paper of Biggs [13].

We concentrate here on the construction of Lazebnik, Ustimenko and Woldar. These graphs first appeared in a series of papers [45, 46, 43, 47] between 1995 and 1997. The graphs as constructed in those papers have their origins in Lie algebras, and use a notation which is somewhat awkward. Partly for this reason, few other authors have attempted a systematic analysis of the properties of the graphs.

More recently, a survey paper by Lazebnik and Sun [44] recasts the notation of these graphs in a more accessible way, and summarises their known properties. Nevertheless, there is still no complete published account of the automorphisms of these graphs. In an attempt to better understand this important family of graphs, we derive a group of automorphisms which arise in a natural way from the recasting of the description of the graphs in a more accessible format. As a consequence of this derivation, we prove that these graphs have a higher level of symmetry than had previously been known, being 3-arc transitive in 75% of cases.

7.2 The graphs of Lazebnik, Ustimenko and Woldar

7.2.1 Construction and properties

We begin by defining the graphs using a notation similar to that used in [44].

Let q be a prime power and let $n \geq 3$. Let P, L be two copies of the vector space of dimension n over $GF(q)$. For convenience we denote a vector $p \in P$ by $(p_1, p_2, \ldots, p_n)$ and a vector $\ell \in L$ by $[\ell_1, \ell_2, \ldots, \ell_n]$. We define a bipartite graph $D(n, q)$ to have vertex partitions P and L and an edge between p and ℓ if and only if the following $n-1$ identities for $\ell_2, \ell_3, \ldots, \ell_n$ are simultaneously satisfied:

$$\begin{aligned}
\ell_2 &= p_2 + p_1\ell_1 \\
\ell_3 &= p_3 + p_1\ell_2 \\
\ell_i &= \begin{cases} p_i + p_{i-2}\ell_1 & \text{if } i \equiv 0, 1 \pmod 4,\ 4 \leq i \leq n \\ p_i + p_1\ell_{i-2} & \text{if } i \equiv 2, 3 \pmod 4,\ 6 \leq i \leq n \end{cases}
\end{aligned}$$

An example graph with $q = 3, n = 3$ is shown in Figure 7.1.

We note that the paper [44] in fact expresses the adjacency equations in the form $p_2 + \ell_2 = p_1\ell_1$ and so on. By Proposition 2 of that paper however, the two forms yield isomorphic graphs and so are equivalent.

It can be seen that given any p, the first coordinate ℓ_1 may be chosen freely and then $\ell_2, \ldots, \ell_n$ are determined. Thus any vertex in P has exactly q neighbours. By reversing the equations it is readily seen that any $\ell \in L$ also has exactly q neighbours. Thus $D(n, q)$ is a q-regular bipartite graph of order $2q^n$. The basic properties of $D(n, q)$ were explored in the original papers and summarised in [44]. The crucial points are as follows.

1. For odd n, the graph $D(n, q)$ has girth at least $n + 5$.
2. For even n, the graph $D(n, q)$ has girth at least $n + 4$.
3. For $n \geq 6$ and odd q, the graph $D(n, q)$ is disconnected and consists of q^t mutually isomorphic components, where $t = \lfloor \frac{n-2}{4} \rfloor$.

Because the graphs are disconnected, we focus attention on the connected components which we denote $CD(n, q)$. These graphs remain asymptotically the best general construction in the girth problem. Thus their properties are of interest.

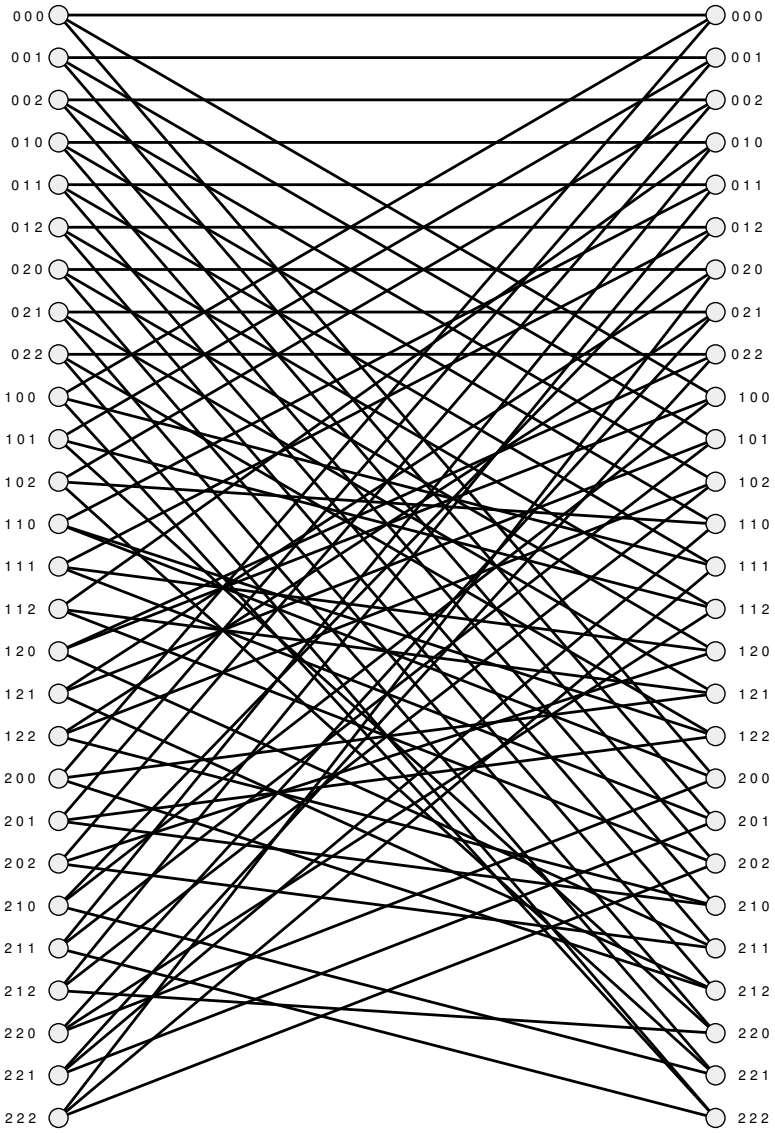

Figure 7.1: The graph $D(3,3)$

A number of known automorphisms of the graphs $D(n,q)$ are summarised in [44]. It is known that the graphs are edge-transitive, and hence transitive on each of the vertex partition sets P, L. In addition, if q is even then the graphs are vertex-transitive. However, there is as far as we know no more complete investigation of the automorphism groups of these graphs. We begin with a more detailed treatment of those automorphisms which preserve the partition sets P, L.

7.2.2 Partition-preserving automorphisms

Because $D(n,q)$ contains a large number of mutually isomorphic connected components as n grows, its automorphism group becomes very unwieldy. We will therefore concentrate on the connected components $CD(n,q)$. Our main result towards a classification of the partition-preserving automorphisms of $CD(n,q)$ is as follows.

Theorem 7.1. *Let $n \geq 2$ and let $m = n - \lfloor \frac{n-2}{4} \rfloor$. Let q be any prime power. Then there exists a group of automorphisms of $CD(n,q)$ of order $q^{m+1}(q-1)^2$ which preserves the vertex partition.*

It will be convenient for the proof of this result to consider the vertices of $CD(n,q)$ as $(n+1)$-dimensional vectors over $GF(q)$ where we simply add a 1 in the final coordinate position. The strategy of the proof is to show that for all n and q, there exists the following:

- a group of $(n+1) \times (n+1)$ matrices G_P over $GF(q)$
- a function ϕ mapping each element of G_P onto another $(n+1) \times (n+1)$ matrix

such that for every $M_P \in G_P$ the following map ψ is an automorphism of the component $CD(n,q)$ of $D(n,q)$ containing the vertex $(0,0,\ldots,0)$:

$$\psi(v) = \begin{cases} vM_P & \text{if } v \in P \\ v\phi(M_P) & \text{if } v \in L \end{cases}$$

We will see that the matrices M_P will have an "affine" block form $\begin{pmatrix} M & 0 \\ F & 1 \end{pmatrix}$ where M is an upper triangular $n \times n$ matrix and F is a $1 \times n$ block. It will turn out that m of the n entries in F are free parameters in $GF(q)$, and the matrices M have the form

$\begin{pmatrix} a & b & y_{13} & \dots & y_{1n} \\ 0 & c & y_{23} & \dots & y_{2n} \\ 0 & 0 & y_{33} & \dots & y_{3n} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & y_{nn} \end{pmatrix}$. In M, a and c can be chosen freely from $GF(q)^*$, and b from $GF(q)$. All the y_{ij} will be completely determined by the freely chosen variables.

The set of such matrices forms a group of order $q^{m+1}(q-1)^2$.

To illustrate the method we begin with the case $n = 3$.

Lemma 7.2. *Let q be a prime power, let $a, c \in GF(q)^*$ and let $b, d, e, f \in GF(q)$. Let M_P be the matrix*

$$M_P = \begin{pmatrix} a & b & db - ea & 0 \\ 0 & c & dc & 0 \\ 0 & 0 & ac & 0 \\ d & e & f & 1 \end{pmatrix}$$

Let $\phi(M_P) = M_L$ be the matrix

$$M_L = \begin{pmatrix} c/a & dc/a & d^2c/a & 0 \\ 0 & c & 2dc & 0 \\ 0 & 0 & ac & 0 \\ -b/a & e - db/a & f + de - d^2b/a & 1 \end{pmatrix}$$

Then the mapping

$$\psi(v) = \begin{cases} vM_P & \text{if } v \in P \\ vM_L & \text{if } v \in L \end{cases}$$

is an automorphism of $D(3, q)$.

Proof. Let $p = (p_1, p_2, p_3, 1)$ and $\ell = [\ell_1, \ell_2, \ell_3, 1]$. We must show that $\psi(p)$ and $\psi(\ell)$ satisfy the adjacency equations if and only if p and ℓ do. So suppose p and ℓ are adjacent. Then we know $\ell_2 = p_2 + p_1\ell_1$ and $\ell_3 = p_3 + p_1\ell_2$. Now:

$$\psi(p) = \begin{pmatrix} a\,p_1 + d \\ b\,p_1 + c\,p_2 + e \\ (db - ea)\,p_1 + dc\,p_2 + ac\,p_3 + f \\ 1 \end{pmatrix}^T$$

$$\psi(\ell) = \begin{bmatrix} c/a\,\ell_1 - b/a \\ dc/a\,\ell_1 + c\,\ell_2 + e - db/a \\ d^2c/a\,\ell_1 + 2dc\,\ell_2 + ac\,\ell_3 + f + de - d^2b/a \\ 1 \end{bmatrix}^T$$

Elementary manipulation now shows that, in an obvious notation, $\psi(\ell)_2 = \psi(p)_2 + \psi(p)_1\psi(\ell)_1$ and $\psi(\ell)_3 = \psi(p)_3 + \psi(p)_1\psi(\ell)_2$ as required. So $\psi(p)$ and $\psi(\ell)$ are adjacent.

It is easy to see that this argument is reversible so that $\psi(p)$ and $\psi(\ell)$ are adjacent if and only if p and ℓ are adjacent. □

To extend the idea to arbitrary n, we will derive the form of the matrices in Lemma 7.2 in a way which will point towards an inductive approach for larger dimensions. But before that, we need to understand more about the structure of the components $CD(n, q)$ of our graphs. These components were investigated by Lazebnik, Ustimenko and Woldar in a follow-up paper [43] and we present the results of that investigation in a form suitable for our needs.

Lemma 7.3. *Let q be a prime power and let $p = (p_1, \ldots, p_6)$ and $\ell = [\ell_1, \ldots, \ell_6]$ be vertices of $D(n, q)$. Define $f : D(6, q) \to GF(q)$ by*

$$f(p) = p_1p_4 - p_2^2 - p_5 + p_6$$
$$f(\ell) = \ell_1\ell_3 - \ell_2^2 - \ell_5 + \ell_6$$

Then f is constant on a connected component of $D(6, q)$.

Proof. Using the adjacency relations it follows by elementary algebraic manipulation that if p is adjacent to ℓ then $f(p) = f(\ell)$. The result follows immediately. □

The corollary of this result is that since f can take any value in $GF(q)$, there are at least q components in $D(6, q)$. In fact the authors go further and show that there are exactly q components when q is odd (the behaviour of the graphs when q is even is slightly different). Moreover, this same splitting happens again at $n = 10$ where they show the existence of a second function which is constant on connected components, leading to a total of q^2 components of the full graph. In general, for any $n = 4k + 2$ where $k \geq 1$, another such function appears. This is essentially due to the periodic structure of the adjacency relations, which repeat their form with a cycle of length 4.

Now we are ready to extend Lemma 7.2 to arbitrary n. To illustrate the inductive

approach we will begin with the simpler case $n = 2$, and then extend this to the case $n = 3$ to re-derive the form of the matrices in Lemma 7.2. We amend our matrix notation slightly by including the dimension n of the space under consideration as a superscript.

Let q be a prime power and let $a, c \in GF(q)^*$. Let $b, d, e \in GF(q)$. Let $M_P^{(2)} = \begin{pmatrix} a & b & 0 \\ 0 & c & 0 \\ d & e & 1 \end{pmatrix}$. We wish to find a matrix $M_L^{(2)} = \begin{pmatrix} x & y & 0 \\ 0 & z & 0 \\ t & u & 1 \end{pmatrix}$ such that the mapping

$$\psi(v) = \begin{cases} vM_P^{(2)} & \text{if } v \in P \\ vM_L^{(2)} & \text{if } v \in L \end{cases}$$

is an automorphism of $D(2, q)$. Our goal then is to find values of x, y, z, t, u in terms of a, b, c, d, e to make this work, that is to say that $p = (p_1, p_2)$ is adjacent to $\ell = [\ell_1, \ell_2]$ if and only if $\psi(p) = (p'_1, p'_2)$ is adjacent to $\psi(\ell) = [\ell'_1, \ell'_2]$.

We begin with a simple set of equations given by the form of ψ:

$$\begin{aligned} p'_1 &= ap_1 + d \\ p'_2 &= bp_1 + cp_2 + e \\ \ell'_1 &= x\ell_1 + t \\ \ell'_2 &= y\ell_1 + z\ell_2 + u \end{aligned}$$

From the adjacency relations, if $p \sim \ell$ we must have $\ell_2 = p_2 + p_1\ell_1$ so that:

$$\ell'_2 = y\ell_1 + z(p_2 + p_1\ell_1) + u \tag{7.1}$$

If $p \sim \ell$ we must have $\psi(p) \sim \psi(\ell)$ so that $\ell'_2 = p'_2 + p'_1\ell'_1$ or:

$$\ell'_2 = bp_1 + cp_2 + e + (ap_1 + d)(x\ell_1 + t) \tag{7.2}$$

We need these equations to be satisfied for all possible values of p_1, p_2, ℓ_1. So we

equate coefficients in 7.1 and 7.2 to get:

$$\begin{aligned} y &= xd \\ z &= c \\ z &= ax \\ u &= dt + e \\ b + at &= 0 \end{aligned}$$

This leads to $M_L^{(2)} = \begin{pmatrix} c/a & dc/a & 0 \\ 0 & c & 0 \\ -b/a & e - db/a & 1 \end{pmatrix}$.

Our goal now is to extend these matrices for $n = 2$ to the case $n = 3$. To do this, notice that $M_P^{(2)}$ has the block form $\begin{pmatrix} M^{(2)} & 0 \\ F^{(2)} & 1 \end{pmatrix}$ where $M^{(2)}$ is a 2×2 matrix, and $F^{(2)}$ is a 1×2 row. The matrix $M_L^{(2)}$ has a similar form $\begin{pmatrix} N^{(2)} & 0 \\ G^{(2)} & 1 \end{pmatrix}$.

The incidence rules for $D(n, q)$ mean that the first $n - 1$ coordinates satisfy the same equations as in the graph $D(n - 1, q)$. Thus if we have matrices

$$M_P^{(n-1)} = \begin{pmatrix} M^{(n-1)} & 0 \\ F^{(n-1)} & 1 \end{pmatrix}; \qquad M_L^{(n-1)} = \begin{pmatrix} N^{(n-1)} & 0 \\ G^{(n-1)} & 1 \end{pmatrix}$$

then we expect the extended matrices to have the form

$$M_P^{(n)} = \begin{pmatrix} M^{(n-1)} & X^{(n-1)} & 0 \\ 0 & x_n & 0 \\ F^{(n-1)} & f_n & 1 \end{pmatrix}; \qquad M_L^{(n)} = \begin{pmatrix} N^{(n-1)} & Y^{(n-1)} & 0 \\ 0 & y_n & 0 \\ G^{(n-1)} & g_n & 1 \end{pmatrix}.$$

In the above, $X^{(n-1)} = (x_1, x_2, \ldots, x_{n-1})^T$ and x_n are entries which will be determined by the adjacency rules, and f_n may be chosen freely. In $M_L^{(n)}$, all the new matrix entries are determined by the adjacency rules and our choice of f_n.

To see how this works we complete the extension to the case $n = 3$. Since $M_P^{(2)} = \begin{pmatrix} a & b & 0 \\ 0 & c & 0 \\ d & e & 1 \end{pmatrix}$ and $M_L^{(2)} = \begin{pmatrix} c/a & dc/a & 0 \\ 0 & c & 0 \\ -b/a & e - db/a & 1 \end{pmatrix}$, we expect our matrices to

have the following forms:

$$M_P^{(3)} = \begin{pmatrix} a & b & x & 0 \\ 0 & c & y & 0 \\ 0 & 0 & z & 0 \\ d & e & f & 1 \end{pmatrix}; \qquad M_L^{(3)} = \begin{pmatrix} c/a & dc/a & t & 0 \\ 0 & c & u & 0 \\ 0 & 0 & v & 0 \\ -b/a & e - db/a & w & 1 \end{pmatrix}$$

where the values of x, y, z, t, u, v, w need to be determined.

As before, the form of ψ gives us:

$$\begin{aligned} p_3' &= xp_1 + yp_2 + zp_3 + f \\ \ell_3' &= t\ell_1 + u\ell_2 + v\ell_3 + w \end{aligned}$$

If $p \sim \ell$ then the second equation expands to:

$$\ell_3' = t\ell_1 + u(p_2 + p_1\ell_1) + v(p_3 + p_1(p_2 + p_1\ell_1)) + w \tag{7.3}$$

If $p \sim \ell$ then we must have $\psi(p) \sim \psi(\ell)$ and so:

$$\begin{aligned} \ell_3' &= p_3' + p_1'\ell_2' \\ &= xp_1 + yp_2 + zp_3 + f + (ap_1 + d)(dc/a\ell_1 + c(p_2 + p_1\ell_1) + e - db/a) \end{aligned} \tag{7.4}$$

Since these equations must hold for all possible values of p_1, p_2, p_3, ℓ_1, we equate coefficients in 7.3 and 7.4 to get $x = db - ea$, $y = dc$, $z = ac$, $t = d^2c/a$, $u = 2dc$, $v = ac$ and $w = f + de - d^2b/a$. Notice that as expected, the value of f can be chosen freely but all other matrix entries were completely determined by the adjacency rules. We have therefore recovered the form of the matrices from Lemma 7.2:

$$M_P^{(3)} = \begin{pmatrix} a & b & db - ea & 0 \\ 0 & c & dc & 0 \\ 0 & 0 & ac & 0 \\ d & e & f & 1 \end{pmatrix}; \qquad M_L^{(3)} = \begin{pmatrix} c/a & dc/a & d^2c/a & 0 \\ 0 & c & 2dc & 0 \\ 0 & 0 & ac & 0 \\ -b/a & e - db/a & f + de - d^2b/a & 1 \end{pmatrix}.$$

We now proceed inductively. To move from the case $n = 3$ to $n = 4$, we set up the equations in the same way as before, but applying the adjacency rules in the fourth coordinate and equating coefficients to determine the required matrix entries. The

details are tedious but routine, and the resulting matrices are:

$$M_P^{(4)} = \begin{pmatrix} a & b & db-ea & b^2/a & 0 \\ 0 & c & dc & 2bc/a & 0 \\ 0 & 0 & ac & 0 & 0 \\ 0 & 0 & 0 & c^2/a & 0 \\ d & e & f & g & 1 \end{pmatrix}$$

$$M_L^{(4)} = \begin{pmatrix} c/a & dc/a & d^2c/a & ec/a & 0 \\ 0 & c & 2dc & bc/a & 0 \\ 0 & 0 & ac & 0 & 0 \\ 0 & 0 & 0 & c^2/a & 0 \\ -b/a & e-db/a & f+de-d^2b/a & g-eb/a & 1 \end{pmatrix}$$

When $n = 5$, the same process yields:

$$M_P^{(5)} = \begin{pmatrix} a & b & db-ea & b^2/a & b^2d/a-be & 0 \\ 0 & c & dc & 2bc/a & 2bcd/a-ce & 0 \\ 0 & 0 & ac & 0 & bc & 0 \\ 0 & 0 & 0 & c^2/a & c^2d/a & 0 \\ 0 & 0 & 0 & 0 & c^2 & 0 \\ d & e & f & g & h & 1 \end{pmatrix}$$

$$M_L^{(5)} = \begin{pmatrix} c/a & dc/a & d^2c/a & ec/a & cf/a & 0 \\ 0 & c & 2dc & bc/a & bcd/a-ce & 0 \\ 0 & 0 & ac & 0 & 0 & 0 \\ 0 & 0 & 0 & c^2/a & c^2d/a & 0 \\ 0 & 0 & 0 & 0 & c^2 & 0 \\ -b/a & e-db/a & f+de-d^2b/a & g-eb/a & h-bf/a & 1 \end{pmatrix}$$

Moving to the case $n = 6$ is similar, but we need to take care because in this case the

graph $D(6,q)$ becomes disconnected. Applying the same process as before gives:

$$M_P^{(6)} = \begin{pmatrix} a & b & db-ea & \frac{b^2}{a} & \frac{b^2d}{a}-be & be-ag & 0 \\ 0 & c & dc & \frac{2bc}{a} & \frac{2bcd}{a}-ce & ce & 0 \\ 0 & 0 & ac & 0 & bc & bc & 0 \\ 0 & 0 & 0 & c^2/a & \frac{c^2d}{a} & 0 & 0 \\ 0 & 0 & 0 & 0 & c^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & c^2 & 0 \\ d & e & f & g & h & i & 1 \end{pmatrix}$$

$$M_L^{(6)} = \begin{pmatrix} \frac{c}{a} & \frac{dc}{a} & \frac{d^2c}{a} & \frac{ec}{a} & \frac{cf}{a} & \frac{cde}{a} & 0 \\ 0 & c & 2dc & \frac{bc}{a} & \frac{bcd}{a}-ce & \frac{bcd}{a}+ce & 0 \\ 0 & 0 & ac & 0 & 0 & bc & 0 \\ 0 & 0 & 0 & \frac{c^2}{a} & \frac{c^2d}{a} & \frac{c^2d}{a} & 0 \\ 0 & 0 & 0 & 0 & c^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & c^2 & 0 \\ -\frac{b}{a} & e-\frac{db}{a} & f+de-\frac{d^2b}{a} & g-\frac{eb}{a} & h-\frac{bf}{a} & i+dg-\frac{bde}{a} & 1 \end{pmatrix}$$

These matrices certainly represent an automorphism of $D(6,q)$. But this automorphism may send a vertex out of its connected component, so it may not be an automorphism of $CD(n,q)$. We want to understand the automorphisms of $CD(n,q)$, so we need to prevent this.

Recall from Lemma 7.3 that two vertices p, p' of $D(n,q)$ are in the same connected component if and only if $f(p) = f(p')$ as defined in the lemma. The form of f means that given p, we can express p'_6 as a function of $p'_1, \ldots, p'_5$ to ensure that p and p' are in the same component. In terms of our matrices, this is equivalent to saying that the entry i in the matrix $M_P^{(6)}$ can no longer be chosen freely from $GF(q)$, but must be constrained to be the unique value which will keep p' in the same component as p.

Finally we apply the same process to the case $n = 7$ and calculate the matrices $M_P^{(7)}$

and $M_L^{(7)}$:

$$\begin{pmatrix} a & b & db-ea & \frac{b^2}{a} & \frac{b^2d}{a}-be & be-ag & bf-ah & 0 \\ 0 & c & dc & \frac{2bc}{a} & \frac{2bcd}{a}-ce & ce & cf & 0 \\ 0 & 0 & ac & 0 & bc & bc & bcd-ace & 0 \\ 0 & 0 & 0 & c^2/a & \frac{c^2d}{a} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & c^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & c^2 & c^2d & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & ac^2 & 0 \\ d & e & f & g & h & i & j & 1 \end{pmatrix}$$

$$\begin{pmatrix} \frac{c}{a} & \frac{dc}{a} & \frac{d^2c}{a} & \frac{ec}{a} & \frac{cf}{a} & \frac{cde}{a} & \frac{cdf}{a} & 0 \\ 0 & c & 2dc & \frac{bc}{a} & \frac{bcd}{a}-ce & \frac{bcd}{a}+ce & cf-cde+\frac{bcd^2}{a} & 0 \\ 0 & 0 & ac & 0 & 0 & bc & bcd-ace & 0 \\ 0 & 0 & 0 & \frac{c^2}{a} & \frac{c^2d}{a} & \frac{c^2d}{a} & \frac{c^2d^2}{a} & 0 \\ 0 & 0 & 0 & 0 & c^2 & 0 & c^2d & 0 \\ 0 & 0 & 0 & 0 & 0 & c^2 & c^2d & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & ac^2 & 0 \\ -\frac{b}{a} & e-\frac{db}{a} & f+de-\frac{d^2b}{a} & g-\frac{eb}{a} & h-\frac{bf}{a} & i+dg-\frac{bde}{a} & j+dh-\frac{bdf}{a} & 1 \end{pmatrix}$$

At this point we have constructed pairs of matrices yielding automorphisms of $CD(n,q)$ in the cases $n = 3, 4, 5, 6, 7$ having started from $n = 2$. Given that the adjacency equations for our graphs repeat with a cycle of period 4 as from $n = 3$ onwards, it is clear that our inductive construction approach can be continued up to any arbitrary dimension n. If $n \not\equiv 2 \pmod 4$, then we add one more free variable into our matrix $M_P^{(n)}$. If $n \equiv 2 \pmod 4$, then the graph $D(n,q)$ splits again into another q components as discussed following Lemma 7.3 and the new variable in the matrix must be determined so as to keep a vertex within its connected component. This completes the proof of Theorem 7.1.

7.2.3 Vertex transitivity of $D(n,q)$

As noted in Section 7.2.1, it is known that the automorphism group of $D(n,q)$ acts transitively on each of the partition sets P, L. (In fact, we will see in the next section that this result can be strengthened.) For the graphs to be vertex-transitive, all that is required is that there should exist some automorphism which exchanges P and L. So far, this has been shown [45, Theorem 3.2] to be true for even q or n, although the proof of this result is in a form and uses a notation which is awkward for our use. Our main result in this section shows that such a "swapping" automorphism exists for any

odd q unless $n \equiv 3 \pmod 4$.

Theorem 7.4. *For any prime power q and for any $n \not\equiv 3 \pmod 4$ the graph $D(n,q)$ is vertex-transitive.*

The proof of this result will use induction on n and we begin with the smallest non-trivial case.

Lemma 7.5. *For any prime power q, the graph $D(4,q)$ is vertex-transitive.*

Proof. Let q be any prime power and let P, L be the two partitions of $D(4,q)$. For clarity, we denote an element of P by $p = (a,b,c,d)$ and of L by $\ell = [x,y,z,t]$. The adjacency relations are that $p \sim \ell$ if and only if:

$$y = b + ax$$
$$z = c + ay$$
$$t = d + bx$$

For each $\beta \in GF(q)$, every vertex v in P or L is adjacent to exactly one vertex $N_\beta(v)$ with first coordinate β.

$$N_\beta(p) = [\beta, b + \beta a, c + ab + \beta a^2, d + \beta b]$$

$$N_\beta(\ell) = (\beta, y - \beta x, z - \beta y, t - xy + \beta x^2)$$

From [45, Thm 3.2] we know that the graph is edge transitive and hence transitive on each of P and L. To show vertex transitivity we need only exhibit a "swapping" automorphism ϕ which interchanges the vertices of P and L. Define ϕ by:

$$\phi = \begin{cases} (a,b,c,d) & \mapsto [a,-b,d,c] \\ [x,y,z,t] & \mapsto (x,-y,t,z) \end{cases}$$

We need to show that $p \sim \ell$ if and only if $\phi(p) \sim \phi(\ell)$. Since ϕ fixes the first coordinate it suffices to show for each β that $N_\beta(\phi(p)) = \phi(N_\beta(p))$ and $N_\beta(\phi(\ell)) = \phi(N_\beta(l))$. This follows from the calculations below.

$$N_\beta(\phi(p)) = N_\beta([a,-b,d,c]) = (\beta, -b - \beta a, d + \beta b, c + ab + \beta a^2)$$
$$\phi(N_\beta(p)) = [\beta, b + \beta a, c + ab + \beta a^2, d + \beta b]^\phi = (\beta, -b - \beta a, d + \beta b, c + ab + \beta a^2)$$

$$N_\beta(\phi(\ell)) = N_\beta((x,-y,t,z)) = [\beta, -y + \beta x, t - xy + \beta x^2, z - \beta y]$$
$$\phi(N_\beta(l)) = \phi((\beta, y - \beta x, z - \beta y, t - xy + \beta x^2)) = [\beta, -y + \beta x, t - xy + \beta x^2, z - \beta y]$$

So ϕ is an automorphism. □

Our strategy now is somewhat similar to our strategy for the proof of Theorem 7.1. We complete some small cases to see how the process works, and then find an inductive argument based on the period four cycle in the adjacency relations. The proof for the case $n = 5$ is along similar lines:

Lemma 7.6. *For any prime power q, the graph $D(5, q)$ is vertex-transitive.*

Proof. Let q be any prime power and let P, L be the two partitions of $D(5, q)$. We denote an element of P by $p = (a, b, c, d, e)$ and of L by $l = [x, y, z, t, u]$. The adjacency relations are that $p \sim l$ if and only if:

$$y = b + ax$$
$$z = c + ay$$
$$t = d + bx$$
$$u = e + cx$$

For each $\beta \in GF(q)$, every vertex v in P or L is adjacent to exactly one vertex $N_\beta(v)$ with first coordinate β.

$$N_\beta(p) = [\beta, b + \beta a, c + ab + \beta a^2, d + \beta b, e + \beta c]$$

$$N_\beta(l) = (\beta, y - \beta x, z - \beta y, t - xy + \beta x^2, u - xz + \beta xy)$$

Again, we need only show the existence of an automorphism ϕ swapping P and L. Define ϕ by:

$$\phi = \begin{cases} (a, b, c, d, e) & \mapsto [a, -b, d, c, ad - b^2 - e] \\ [x, y, z, t, u] & \mapsto (x, -y, t, z, xz - y^2 - u) \end{cases}$$

Since ϕ fixes the first coordinate it suffices to show for each β that

$N_\beta(\phi(p)) = \phi(N_\beta(p))$ and $N_\beta(\phi(\ell)) = \phi(N_\beta(\ell))$:

$$\begin{aligned} N_\beta(\phi(p)) &= N_\beta([a,-b,d,c,ad-b^2-e]) \\ &= (\beta,-b-\beta a,d+\beta b,c+ab+\beta a^2,-b^2-e-\beta ab) \\ \phi(N_\beta(p)) &= \phi([\beta,b+\beta a,c+ab+\beta a^2,d+\beta b]) \\ &= (\beta,-b-\beta a,d+\beta b,c+ab+\beta a^2,-b^2-e-\beta ab) \end{aligned}$$

$$\begin{aligned} N_\beta(\phi(\ell)) &= N_\beta((x,-y,t,z,xz-y^2-u)) \\ &= [\beta,-y+\beta x,t-xy+\beta x^2,z-\beta y,xz-y^2-u+\beta t] \\ \phi(N_\beta(\ell)) &= \phi((\beta,y-\beta x,z-\beta y,t-xy+\beta x^2)) \\ &= [\beta,-y+\beta x,t-xy+\beta x^2,z-\beta y,xz-y^2-u+\beta t] \end{aligned}$$

So ϕ is an automorphism. □

For the case $n = 6$, we know from [43, Prop 5.1] that the graph $D(6,q)$ is disconnected, consisting of q connected components each isomorphic to $D(5,q)$. So the following is immediate:

Lemma 7.7. *For any prime power q, the graph $D(6,q)$ is vertex-transitive.*

We are now ready to prove Theorem 7.4. The strategy is to show that for any $k \geq 1$, the existence of a swapping automorphism for $D(4k,q)$ implies one for $D(4k+2,q)$ and $D(4k+4,q)$. Further, the one for $D(4k+2,q)$ implies one for $D(4k+1,q)$. Since Lemmas 7.5, 7.6 and 7.7 have proved the result for $n = 4,5,6$ the result follows.

Proof of Theorem 7.4. Let $k \geq 1$, let q be a prime power and suppose there is an automorphism ϕ of $D(4k,q)$ which interchanges the sets P and L and has the following form:

$$\begin{cases} p = (p_1,p_2,\ldots,p_{4k}) & \mapsto [p_1,-p_2,p_4,p_3,\ldots,-p_{4k-2},-p_{4k-3},p_{4k},p_{4k-1}] \\ \ell = [\ell_1,\ell_2,\ldots,\ell_{4k}] & \mapsto (\ell_1,-\ell_2,\ell_4,\ell_3,\ldots,-\ell_{4k-2},-\ell_{4k-3},\ell_{4k},\ell_{4k-1}) \end{cases}$$

We deal first with $4k+2$. We write p' for the vector p extended by two coordinates and similarly for ℓ and ϕ. In the usual notation we need to show that $\phi'(N_\beta(p')) = N_\beta(\phi'(p'))$. The adjacency rules give:

$$N_\beta(p') = [N_\beta(p), p_{4k+1} + \beta p_{4k-1}, p_{4k+2} + p_1p_{4k} + \beta p_1 p_{4k-2}]$$

$$\phi'(N_\beta(p')) = [\phi(N_\beta(p)), -p_{4k+2} - p_1p_{4k} - \beta p_1p_{4k-2}, -p_{4k+1} - \beta p_{4k-1}]$$

If we write $\phi'(p') = [\ell_1, \ell_2, \ldots, \ell_{4k}, \ell_{4k+1}, \ell_{4k+2}]$ then by the operation of ϕ':

$$\begin{aligned}
\ell_1 &= a_1 \\
\ell_{4k-3} &= -p_{4k-2} \\
\ell_{4k-2} &= -p_{4k-3} \\
\ell_{4k-1} &= p_{4k} \\
\ell_{4k} &= p_{4k-1} \\
\ell_{4k+1} &= -p_{4k+2} \\
\ell_{4k+2} &= -p_{4k+1}
\end{aligned}$$

If we write $N_\beta(\phi'(p')) = [b_1, b_2, \ldots, b_{4k}, b_{4k+1}, b_{4k+2}]$ then by the adjacency rules:

$$\begin{aligned}
b_{4k+1} &= \ell_{4k+1} - x_1\ell_{4k-1} + \beta\ell_1 x_{4k-3} = -p_{4k+2} - p_1 p_{4k} - \beta p_1 p_{4k-2} \\
b_{4k+2} &= \ell_{4k+2} - \beta\ell_{4k} = -p_{4k+1} - \beta p_{4k-1}
\end{aligned}$$

Since $\phi(N_\beta(p)) = N_\beta(\phi(p))$ by the induction hypothesis, the result follows.

A similar argument shows that $\phi'(N_\beta(\ell')) = N_\beta(\phi'(\ell'))$.

We deal with the case $4k+4$ in exactly the same way, and omit the details for brevity.

For $4k+1$, by [43, Prop 5.1] each component of $D(4k+1, q)$ is isomorphic to a component of $D(4k+2, q)$ and so the result follows immediately. □

7.2.4 Arc transitivity of $D(n, q)$

Our last result on the automorphisms of $D(n, q)$ shows that in fact these graphs have a very high degree of symmetry.

Theorem 7.8. *Let q be any prime power and let $n \geq 3$. Let P and L be the vertex partitions of the component $CD(n, q)$ of $D(n, q)$ containing the vector $(0, 0, 0, \ldots, 0)$. Then the automorphism group of $CD(n, q)$ acts transitively on the set of paths of length 3 with initial vertex in P.*

Proof. We begin the proof in the case $n = 3$. We know that the matrices for

automorphisms of $D(3,q)$ have the following form:

$$M_P^{(3)} = \begin{pmatrix} a & b & db-ea & 0 \\ 0 & c & dc & 0 \\ 0 & 0 & ac & 0 \\ d & e & f & 1 \end{pmatrix} \qquad M_L^{(3)} = \begin{pmatrix} c/a & dc/a & d^2c/a & 0 \\ 0 & c & 2dc & 0 \\ 0 & 0 & ac & 0 \\ -b/a & e-db/a & f+de-d^2b/a & 1 \end{pmatrix}$$

As before, we let $p \in P, \ell \in L$ and consider the vectors defining the vertices to be extended by a 1 in the final coordinate position. An automorphism ψ is determined by these matrices by defining $\psi(p) = pM_P^{(3)}, \psi(\ell) = \ell M_L^{(3)}$.

We want to find an automorphism ψ which maps the path $u_0 \to u_1 \to u_2 \to u_3$ to the path $v_0 \to v_1 \to v_2 \to v_3$ where both u_0 and v_0 are in P. Because we know the automorphism group acts transitively on P it is sufficient to consider the case $u_0 = (0,0,0,1)$. If $v_0 = (p_1, p_2, p_3, 1)$ then because $\psi(u_0) = v_0$ we can simply fill in the bottom row of the matrix A so that $d = p_1, e = p_2, f = p_3$.

Now u_1 is a neighbour of u_0 with first coordinate β_1 say, and β_1 uniquely determines this neighbour. Similarly, u_2 is the unique β_2-neighbour of u_1 and u_3 is the unique β_3-neighbour of u_2. In the transformed path we denote the first coordinate of successor neighbour of each v_i by $\alpha_{i+1}, i = 0, 1, 2$.

Notice that because we are considering paths we disallow backtracking walks, so we have that $\beta_2 \neq 0$ and $\beta_1 \neq \beta_3$. In the same way, $\alpha_2 \neq p_1$ and $\alpha_1 \neq \alpha_3$.

Now we use the equations:

$$\begin{aligned} \psi(u_1) &= u_1 M_L^{(3)} = v_1 \\ \psi(u_2) &= u_2 M_P^{(3)} = v_2 \\ \psi(u_3) &= u_3 M_L^{(3)} = v_3 \end{aligned}$$

Working this through leads to the following solution for a, b, c:

$$\begin{aligned} a &= \frac{\alpha_2 - p_1}{\beta_2} \\ c &= \frac{(\alpha_1 - \alpha_3)a}{\beta_1 - \beta_3} \\ b &= \beta_1 c - \alpha_1 a \end{aligned}$$

Since we have a solution, we conclude that there are matrices $M_P^{(3)}$ and $M_L^{(3)}$ defining an automorphism ψ mapping the path $u_0 \to u_1 \to u_2 \to u_3$ to the path

$v_0 \to v_1 \to v_2 \to v_3$. Thus $\mathrm{Aut}(D(3,q))$ acts transitively on paths of length 3 with initial vertex in P.

To deal with the cases where $n > 3$, we note that the only additional free variables in the matrix $M_P^{(n)}$ occur in the last row. These are immediately determined by the coordinates of v_0 and then finding values for a, b, c works exactly as in the $n = 3$ case. □

Since there exists a swapping automorphism in the case where $n \not\equiv 3 \pmod 4$, in that case we conclude that any path of length 3 starting in P may be mapped to any path of length 3 starting in L. Additionally, $D(n,q)$ consists of a number of mutually isomorphic copies of $CD(n,q)$. We therefore have the following corollary.

Corollary 7.9. *Let $n > 3$ with $n \not\equiv 3 \pmod 4$ and let q be any prime power. Then the graph $D(n,q)$ is 3-arc transitive.*

7.2.5 Summary of automorphisms of $D(n,q)$

Theorems 7.4 and 7.8 together imply that if $n \not\equiv 3 \pmod 4$, then there is a group of automorphisms of $CD(n,q)$ of order $2q^{m+1}(q-1)^2$ where $m = n - \lfloor \frac{n-2}{4} \rfloor$. However, these may not be the only automorphisms of $CD(n,q)$.

It is easy to see that any automorphism of the field $GF(q)$ induces an automorphism of $CD(n,q)$ by acting on P, L in the natural way. Such an automorphism will of course preserve the vertex partition, but any non-trivial field automorphism induces a graph automorphism which is not in our matrix form from Theorem 7.1. To see why, let $q = p^k$ where p is prime and consider the subset P^* of P consisting of those vectors where each component $p_1, \ldots, p_n$ is contained in the prime subfield $GF(p)$. Then any non-trivial automorphism σ of $GF(q)$ fixes the prime subfield, and hence fixes all elements of the set P^*. But an examination of the form of the matrices $M_P^{(n)}$ from Section 7.2.2 shows that no non-identity matrix of this form can fix all elements of P^*.

If $q = p^k$ where p is prime, then there are k such field automorphisms, and so we have a group of order $2kq^{m+1}(q-1)^2$.

To investigate the full automorphism group of $CD(n,q)$ for small values of n and q we used the `GRAPE` [69] package within `GAP` [35] to construct the graphs. This package uses `nauty` [55] to compute the automorphism group. The results are tabulated in Table 7.1.

For cubic graphs ($q = 3$) we may verify our results for small values of n by comparing with the Foster census [22] of cubic arc-transitive graphs. For $n = 4$, the graph

$CD(4,3)$ has order 162 and appears as entry `F162C` in the census, which gives the order of its automorphism group as 1944 in agreement with our table. Isomorphism testing with `GRAPE` [69] confirms that $CD(4,3)$ and `F162C` are isomorphic. Similarly, $CD(5,3)$ appears as entry `F486C` with automorphism group of order 5832. (Since $CD(3,3)$ is not arc-transitive it does not appear in the Foster census.)

In most cases, the full automorphism group has order equal to the subgroup we computed above. The exceptions are for even q where it is known [43] that the split of $D(n,q)$ into connected components behaves slightly differently, and for $q = 3$ where the properties of $D(n,3)$ for some values of n are not fully understood [44]. In addition, for odd q it appears from the computational experiments that there is no swapping automorphism if $n \equiv 3 \pmod 4$. We therefore make the following conjecture.

Conjecture 7.10. *Let $n \geq 3$, let $m = n - \lfloor \frac{n-2}{4} \rfloor$ and let $q = p^k$ be an odd prime power larger than 3. Then the automorphism group of a connected component $CD(n,q)$ of the graph $D(n,q)$ has exact order:*

$$\begin{cases} kq^{m+1}(q-1)^2 & \text{if } q \equiv 3 \pmod 4 \\ 2kq^{m+1}(q-1)^2 & \text{if } q \equiv 0,1,2 \pmod 4 \end{cases}$$

In particular, if $q \not\equiv 3 \pmod 4$ then $\mathrm{Aut}(CD(n,q))$ acts regularly on the 3-arcs of $CD(n,q)$.

7.2.6 Extension to other degrees

Recall that in the diameter problem, we had a number of constructions of Cayley graphs which were valid only for degrees related to prime powers. Our strategy in that case was to add edges to our graphs by adding generators to our set, to cover remaining degrees.

In a similar way, it is immediate from the definition of the graphs $D(n,q)$ that the construction is only valid for degrees which are prime powers. However, we can amend the graphs by removing vertices in a controlled way to produce graphs of any required degree. We do this by finding a *perfect dominating set* in our graph: that is, a set of vertices S such that any vertex not in S is adjacent to precisely one vertex in S. Removal of such a vertex set results in a regular graph of degree one less than the original, and of course this process cannot decrease the girth.

Lemma 7.11. *Let q be a prime power and let $g \geq 8$ be an even number. Let d be any desired degree with $d \leq q$. Then there exists a d-regular graph of girth at least g and*

n	$q = p^k$	m	$\lvert \mathrm{Aut}(CD(n,q)) \rvert$	Interpretation
3	3	3	1296	$kq^{m+1}(q-1)^2 \times 4$
3	4	3	18432	$2kq^{m+1}(q-1)^2 \times 2$
3	5	3	10000	$kq^{m+1}(q-1)^2$
3	7	3	86436	$kq^{m+1}(q-1)^2$
3	8	3	1204224	$2kq^{m+1}(q-1)^2$
3	9	3	839808	$kq^{m+1}(q-1)^2$
3	11	3	1464100	$kq^{m+1}(q-1)^2$
3	13	3	4112784	$kq^{m+1}(q-1)^2$
3	16	3	117964800	$2kq^{m+1}(q-1)^2$
3	17	3	21381376	$kq^{m+1}(q-1)^2$
3	19	3	42224004	$kq^{m+1}(q-1)^2$
4	3	4	1944	$2kq^{m+1}(q-1)^2$
4	4	4	18432	$2kq^{m+1}(q-1)^2 \times 2$
4	5	4	100000	$2kq^{m+1}(q-1)^2$
4	7	4	1210104	$2kq^{m+1}(q-1)^2$
4	8	4	1849688064	$2kq^{m+1}(q-1)^2 \times 2$
4	9	4	15116544	$2kq^{m+1}(q-1)^2$
4	11	4	32210200	$2kq^{m+1}(q-1)^2$
4	13	4	106932384	$2kq^{m+1}(q-1)^2$
5	3	5	5832	$2kq^{m+1}(q-1)^2$
5	4	5	36864	$2kq^{m+1}(q-1)^2 \times 4$
5	5	5	500000	$2kq^{m+1}(q-1)^2$
5	7	5	8470728	$2kq^{m+1}(q-1)^2$
5	8	5	77070336	$2kq^{m+1}(q-1)^2$
6	3	5	5832	$2kq^{m+1}(q-1)^2$
6	4	5	36864	$2kq^{m+1}(q-1)^2 \times 4$
6	5	5	500000	$2kq^{m+1}(q-1)^2$
6	7	5	8470728	$2kq^{m+1}(q-1)^2$
6	8	5	77070336	$2kq^{m+1}(q-1)^2$
7	3	6	34992	$kq^{m+1}(q-1)^2 \times 4$
7	4	6	73728	$2kq^{m+1}(q-1)^2 \times 8$
8	3	7	52488	$2kq^{m+1}(q-1)^2$
9	3	8	157464	$2kq^{m+1}(q-1)^2$
10	3	8	157464	$2kq^{m+1}(q-1)^2$
11	3	9	236196	$kq^{m+1}(q-1)^2$

Table 7.1: Automorphisms of $CD(n,q)$

order $2q^{g-5} - 2q(q-d)$.

Proof. Let $n = g - 5$. Then the graph $D(n, q)$ has girth at least g and is q-regular. Let P, L be the vertex partitions of $D(n, q)$. Recall that for any $\beta \in GF(q)$, a vertex $p \in P$ has precisely one neighbour $N_\beta(p) \in L$ with first coordinate β. Likewise, a vertex $\ell \in L$ has precisely one neighbour $N_\beta(\ell) \in P$ with first coordinate β. Thus the set of vertices with first coordinate β forms a perfect dominating set for $D(n, q)$.

This vertex set of size $2q$ may be removed from the graph, leaving a $(q-1)$-regular graph of order $2q^n - 2q$ and girth at least g. This process can be repeated $q - d$ times to yield the desired graph. □

We note that a similar construction is described in the recent preprint of Lazebnik and Sun [44].

7.2.7 Voltage lifts

We conclude our discussion of the properties of the graphs $CD(n, q)$ by describing another method for their construction by means of iterated voltage lifts. We begin with some background and notation.

Let Γ be an undirected graph which we call a *base* graph. In contrast to our convention to date, we allow Γ to have loops and multiple edges. Although Γ is undirected, we think of its edges as being formed by pairs of oppositely directed arcs which in this context we call *darts*. If e is a dart then e^{-1} will denote its reverse. If $D(\Gamma)$ is the dart set of Γ, then $|D(\Gamma)| = 2|E(\Gamma)|$.

For a finite group G, a mapping $\alpha : G \to D(\Gamma)$ is called a *voltage assignment* if $\alpha(e^{-1}) = (\alpha(e))^{-1}$ for all $e \in D(\Gamma)$. Given a voltage assignment α, we define the *voltage lift* Γ^α as follows.

The vertex set $V(\Gamma^\alpha)$ is the Cartesian product $V(\Gamma) \times G$, and the dart set is $D(\Gamma) \times G$. Let e be a dart in Γ from vertex u to v. We define the dart (e, g) to have initial vertex (u, g) and terminal vertex $(v, g\alpha(e))$. Note that by the definition of voltage assignments, Γ^α is an undirected graph.

An example base graph is shown in Figure 7.2. It consists of a pair of vertices with three edges (pairs of darts) between them. The darts from u to v are assigned voltages $0, 1, 3$ from the group $\mathbb{Z}_7$. (Naturally, their respective inverses are assigned $0, -1, -3$.) The lifted graph is depicted in Figure 7.3.

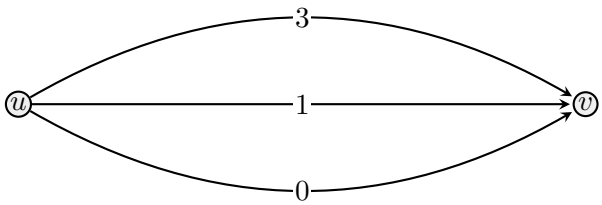

Figure 7.2: Base graph Γ

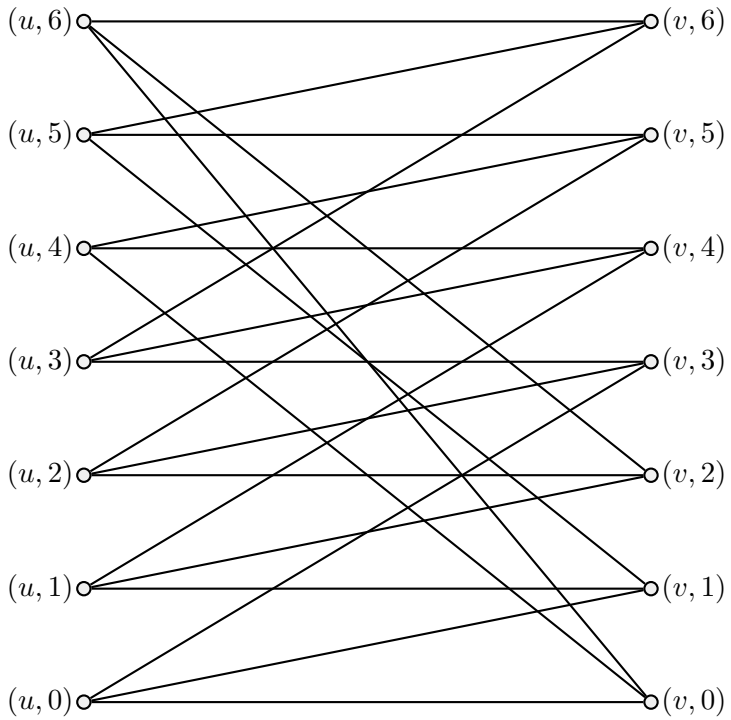

Figure 7.3: Lifted graph Γ^{α}

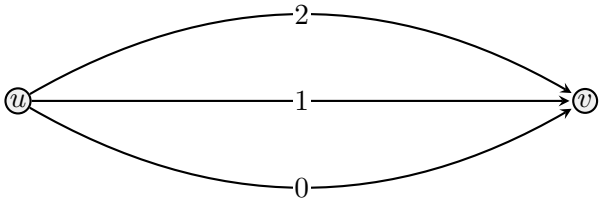

Figure 7.4: Base dipole graph Γ

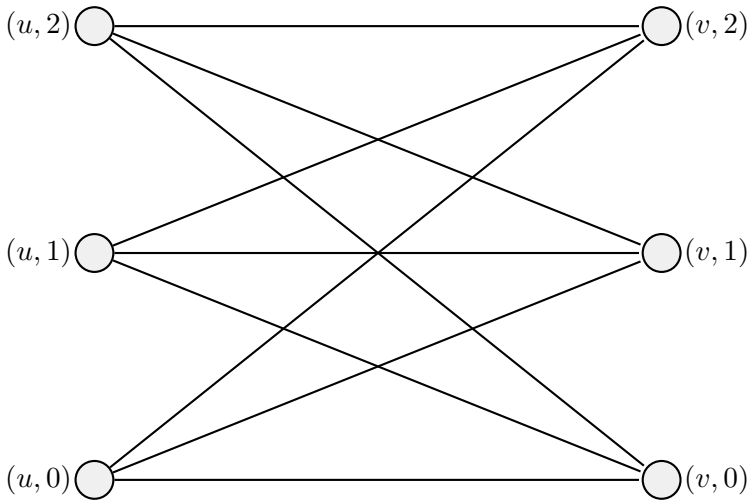

Figure 7.5: Lifted graph $\Gamma^\alpha \cong K_{3,3} \cong CD(1,3)$

Notice that the lifted graph is regular of degree 3, as was the base graph. In fact this graph turns out to be the Heawood graph, which is the unique graph attaining the Moore bound at degree 3 and girth 6.

Our aim now is to construct our graphs $CD(n,q)$ via an iterated sequence of voltage lifts. Our base graph Γ will be a graph with two vertices u, v and q pairs of darts between them. The voltage assignment α assigns each element of $GF(q)$ to exactly one of the darts from u to v.

It is not hard to see that the lifted graph Γ^α is isomorphic to the complete bipartite graph $K_{q,q}$. We may equivalently view this graph as $D(1,q)$ or $CD(1,q)$ since the operation of the adjacency rules places no restriction on the first (in this case, only) coordinate of an adjacent vertex. The base and lifted graphs in the case $q = 3$ are shown in Figures 7.4 and 7.5.

Now we proceed to construct $CD(2,q)$ as a lift from a base graph of $CD(1,q)$. In $CD(1,q)$ we identify the vertex $p = (p_1)$ with the vertex (u, p_1) in our first lifted graph, and similarly $\ell = [\ell_1]$ with (v, ℓ_1).

Our second lifted graph will have vertices of the forms (u, p_1, p_2) and (v, ℓ_1, ℓ_2) which we will identify with vertices (p_1, p_2) and $[\ell_1, \ell_2]$ of $CD(2,q)$ in the obvious way. It

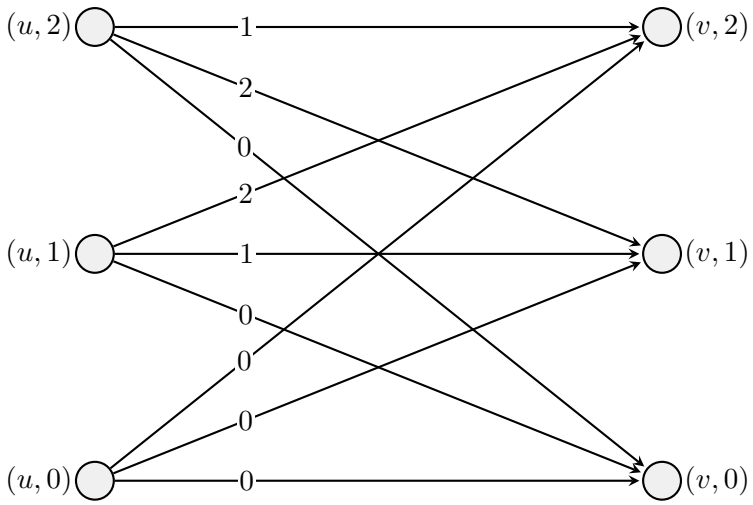

Figure 7.6: The second base graph $CD(1,3)$

remains now to find a voltage assignment α so that the lifted graph will be isomorphic to $CD(2,q)$.

Recall that the adjacency rule for $CD(2,q)$ is that $\ell_2 = p_2 + p_1\ell_1$. Suppose we have a dart $e = (u,p_1) \to (v,\ell_1)$ in the base graph $CD(1,q)$. All that is required is to assign a voltage $\alpha(e) = p_1\ell_1$ in the additive group $GF(q)^+$ to this dart, and then the lifted graph will have the correct adjacencies. The second base graph with voltage assignments is shown in Figure 7.6, and the resulting lift in Figure 7.7.

It is now clear how to proceed. To lift $CD(2,q)$ to $CD(3,q)$ we notice that the second adjacency rule in $CD(3,q)$ is $\ell_3 = p_3 + p_1\ell_2$. Thus the voltage on a dart $(u,p_1,p_2) \to (v,\ell_1,\ell_2)$ is $p_1\ell_2$. We would then recover the graph of Figure 7.1.

In general, to move from $CD(n,q)$ to $CD(n+1,q)$ via a voltage lift we need only to notice that the final adjacency relation is of the form $\ell_{n+1} = p_{n+1} + p_i\ell_j$ for some $1 \le i,j \le n$ and so we assign a voltage $p_i\ell_j$ on the dart $(u,p_1,\ldots,p_n) \to (v,\ell_1,\ldots,\ell_n)$.

The exception is if n is of the form $4k+1$ for $k \ge 1$, where we know that in fact $CD(n+1,q)$ is isomorphic to $CD(n,q)$ since the graph splits into further connected components as in Lemma 7.3. We therefore have the following result.

Proposition 7.12. *Let n be a positive integer not of the form $4k+1$ for $k \ge 1$. Let q be a prime power. Let $\Gamma = CD(n,q)$. Then there is a voltage assignment $\alpha : D(\Gamma) \to GF(q)^+$ such that the lifted graph Γ^α is isomorphic to $CD(n+1,q)$.*

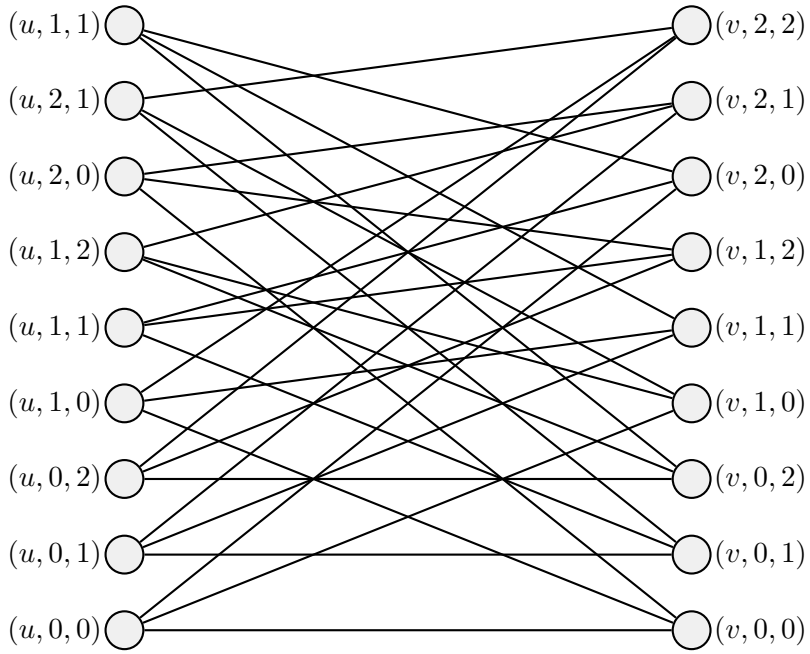

Figure 7.7: The second lifted graph graph $D(2,3)$

Chapter 8

Filled groups

We turn now to a different though somewhat related problem in group theory. The idea is to investigate product-free sets S in a group G, particularly in the case where the set and its products cover the whole group, that is to say $G \setminus \{1\} \subseteq S \cup SS$. Since this condition is equivalent to saying that the digraph $\mathrm{Cay}(G, S)$ has diameter two, we can see that there is a link to the degree-diameter problem. Indeed in Section 8.4 we will use the ideas we develop to derive a new bound on the asymptotic order of a certain family of Cayley graphs.

8.1 Preliminaries

Let S be a non-empty subset of a group G. We say S is *product-free* if $S \cap SS = \emptyset$, where $SS = \{ab : a, b \in S\}$. Note that we do not require a and b to be distinct. A product-free set S is said to be *locally maximal* if whenever Σ is product-free in G and $S \subseteq \Sigma$, then $S = \Sigma$. A product-free set S *fills* G if $G^* \subseteq S \cup SS$ (where G^* is the set of all non-identity elements of G). Product-free sets that fill G are also called *complete* sum-free sets, for example in [20]. We say G is a *filled* group if every locally maximal product-free set in G fills G. This definition, due to Street and Whitehead [70], was motivated by the observation that a product-free set in an elementary abelian 2-group A is locally maximal if and only if it fills A, and hence the elementary abelian 2-groups are filled groups. They asked which other groups, if any, are filled. They classified the filled abelian groups and some small dihedral groups. In this chapter, we classify filled groups of various kinds. In Section 8.2.1 we deal with dihedral groups. Section 8.2.2 covers nilpotent groups. Section 8.2.3 looks at groups of order $2^n p$ where p is an odd prime and n is a positive integer. Finally in Section 8.3 we describe an algorithm which we have implemented in GAP [35], which allows us to check for filled non-nilpotent groups of all orders up to 2000. In the rest of this section we establish notation and state some known results.

Throughout this chapter, we write C_n for the cyclic group of order n and D_{2n} for the dihedral group of order $2n$. Let S be a subset of a group G. We define S^{-1}, $T(S)$ and

$\sqrt{S}$ as follows.

$$S^{-1} = \{s^{-1} : s \in S\}$$
$$T(S) = S \cup SS \cup SS^{-1} \cup S^{-1}S;$$
$$\sqrt{S} = \{x \in G : x^2 \in S\}.$$

We end this preliminary section with the following result, which gathers together some useful facts that we will need.

Theorem 8.1. *(i) [37, Lemma 3.1] Let S be a product-free set in a finite group G. Then S is locally maximal if and only if $G = T(S) \cup \sqrt{S}$.*

(ii) [70, Lemma 1] If G is a filled group, and N is a normal subgroup of G, then G/N is filled.

(iii) [70, Theorem 2] A finite abelian group is filled if and only if it is C_3, C_5 or an elementary abelian 2-group.

(iv) [4, Lemma 2.3] The only filled group with a normal subgroup of index 3 is C_3.

(v) [4, Lemma 2.5] If G is a filled group with a normal subgroup N of index 5 such that not every element of order 5 is contained in N, then $G \cong C_5$.

(vi) [4, Theorem 2.6] The only filled groups of odd order are C_3 and C_5.

(vii) [4, Prop 2.8] For $n \geq 2$, the dicyclic group of order $4n$ is not filled.

For convenience and to give a flavour of the techniques used in analysis of this problem, we repeat below brief proofs of some of the above results.

The proof of *(i)* is immediate by noting that $T(S)$ and $\sqrt{S}$ represent the elements of G which, if added to S, would cause it to cease to be product-free.

Item *(ii)* is crucial to our later analysis and classification of filled groups. To prove it, we suppose that we have some non-filling locally maximal product-free set in the quotient G/N. Then it is easy to see that the corresponding union (in S) of cosets of N would be product-free, locally maximal and non-filling in S, contrary to our assumption that G is filled.

For *(iii)*, we prove the result in one direction which is that any elementary abelian 2-group G must be filled. To see this, note that if S is a locally maximal product-free set in G, then since all elements of G have order 2 we must have $S^{-1} = S$ and $\sqrt{S} = \emptyset$. Thus $T(S) \cup \sqrt{S} = S \cup SS$ and so $G = S \cup SS$ by part *(i)*.

We prove *(vii)* by constructing a non-filling locally maximal product-free set. Let G be the dicyclic group of order $4n$ with presentation $\langle a, b \,|\, a^{2n} = b^4 = 1, a^n = b^2, ab = ba^{-1}\rangle$. Then G has an index 2 cyclic subgroup $C = \langle a\rangle$ and every element $x \in bC$ has order 4 and satisfies $x^2 = a^n \in C$. We choose a locally maximal product-free set $S \subseteq C$ containing the element a^n. (It is immediate that any element of a group is contained in some locally maximal product-free set.) Then, since $bC \subseteq \sqrt{S}$, it follows that S is also a locally maximal product-free subset of G. But $S \subseteq C$ and so clearly does not fill G.

8.2 Classification of filled groups

8.2.1 Dihedral groups

A list of non-abelian filled groups of order less than or equal to 32 was given in [4]. There are eight such groups: six are dihedral, and the remaining two are 2-groups. The dihedral groups on the list are those of order 6, 8, 10, 12, 14 and 22. Our aim in this section is to show that these are in fact the only filled dihedral groups. The arguments in this section are originally by Anabanti and Hart, and are contained in our joint paper [3]. We include this short section here for completeness, and to illustrate the fact the the argument for dihedral groups is somewhat more awkward than that for the dicyclic groups.

We write $D_{2n} = \langle a, b \,|\, a^n = b^2 = 1, ab = ba^{-1}\rangle$ for the dihedral group of order $2n$ (where $n > 2$). In D_{2n}, the elements of $\langle a\rangle$ are called *rotations* and the elements of $\langle a\rangle b$ are called *reflections*. For any subset S of D_{2n}, we write $A(S)$ for $S \cap \langle a\rangle$, the set of rotations of S, and $B(S)$ for $S \cap \langle a\rangle b$, the set of reflections of S.

Observation 8.2. *Suppose S is a subset of D_{2n}. Let $A = A(S)$ and $B = B(S)$. Then, because of the relations in the dihedral group, we have $AA^{-1} = A^{-1}A$, $AB = BA^{-1}$ and $B^{-1} = B$. Therefore*

$$\begin{aligned} SS &= AA \cup BB \cup AB \cup BA; \\ SS^{-1} &= AA^{-1} \cup BB \cup AB; \\ S^{-1}S &= AA^{-1} \cup BB \cup BA; \\ T(S) &= A \cup B \cup AA \cup AA^{-1} \cup BB \cup AB \cup BA \\ &= S \cup SS \cup AA^{-1}. \end{aligned}$$

We also note that $\sqrt{S} = \sqrt{A} \subseteq \langle a\rangle$.

Proposition 8.3. *Let n be an odd integer, with $n \geq 13$. Then D_{2n} is not filled.*

Proof (Anabanti). Let n be an odd integer with $n \geq 13$. Then there is an odd number k for which n is either $5k-6$, $5k-4$, $5k-2$, $5k$ or $5k+2$.

Suppose first that n is $5k-2$ for an odd integer k. Since $n \geq 13$, we note that $k \geq 3$. Now consider the following set S:

$$S := \{a^k, a^{k+2}, \cdots, a^{3k-2}; b, ab, \cdots, a^{k-1}b\}$$

We calculate that
$A(SS) = \{a^{2k}, a^{2k+2}, \cdots, a^{5k-3}\} \cup \{1, a, \cdots, a^{k-1}\} \cup \{a^{4k-1}, a^{4k}, \cdots, 1\}$ and
$B(SS) = \langle a \rangle b - B(S)$. Observe that $a^{3k} \notin S \cup SS$; so S does not fill G.

Let $A = A(S)$. Then $AA^{-1} = \{1, a^2, a^4, \cdots, a^{2k-2}\} \cup \{a^{3k}, a^{3k+2}, \cdots, a^{5k-4}, 1\}$. Thus $T(S) = G$. By Theorem 8.1(i) therefore, S is locally maximal product-free in G, but we have noted that S does not fill G.

Next we suppose $n = 5k$ for an odd integer k, and again since $n \geq 13$, we have $k \geq 3$. Taking the same set $S = \{a^k, a^{k+2}, \cdots, a^{3k-2}; b, ab, \cdots, a^{k-1}b\}$ we find that S is locally maximal product-free but does not fill G.

Now suppose $n = 5k+2$ for $k \geq 3$ and odd. The set U given by

$$U = \{a^{k-2}, a^k, \cdots, a^{3k-2}; b, ab, \cdots, a^{k-3}b\}$$

is locally maximal product-free in G (again using Theorem 8.1(i)), but does not fill G since for example $a^{3k} \notin U \cup UU$.

Next suppose $n = 5k-6$ for $k \geq 5$ and odd. Then the set U given by

$$V = \{a^k, a^{k+2} \cdots, a^{3k-2}; b, ab, \cdots, a^{k-1}b\}$$

is a locally maximal product-free set in G that does not fill G.

Finally, consider the case $n = 5k-4$ for $k \geq 5$ and odd. The set W given by

$$W = \{a^{k-2}, a^k, \cdots, a^{3k-4}; b, ab, \cdots, a^{k-3}b\}$$

is a locally maximal product-free set in G which does not fill G. We have now covered all possibilities for n, and have shown that in each case D_{2n} is not filled. □

Theorem 8.4. *The only filled dihedral groups are D_6, D_8, D_{10}, D_{12}, D_{14} and D_{22}.*

Proof (Anabanti). Let G be dihedral of order $2n$. The filled groups of order up to 32

were classified in [4]. The only filled dihedral groups of order up to 32 are D_6, D_8, D_{10}, D_{12}, D_{14} and D_{22}. It remains to show that if $n > 16$, then D_{2n} is not filled. Suppose $n > 16$. By Proposition 8.3, if n is odd then D_{2n} is not filled, so we can assume n is even. We will show by induction on m, that if $G \cong D_{4m}$ for some integer m greater than 3, then G is not filled. Note that by [4] D_{16}, D_{20}, D_{24}, D_{28} and D_{32} are not filled. So we can assume $m > 8$.

Observe that the quotient of G by its centre is dihedral of order $2m$. If G is filled, then by Theorem 8.1(ii), this dihedral group of order $2m$ must be filled. If m is odd, then by Proposition 8.3, and our assumption that $m > 8$, we have $m = 9$ or $m = 11$. We know that D_{18} is not filled, so $m = 11$, meaning G is D_{44}. However a straightforward calculation shows that $\{a^2, a^5, a^8, a^{18}, a^{21}, a^5y, a^{16}b\}$ is locally maximal product-free in D_{44}, but does not fill D_{44}. Thus if m is odd, then G is not filled. Suppose m is even, so $m = 2t$ for some t with $t > 4$. Inductively D_{4t} is not filled, so G is not filled. This completes the proof. □

8.2.2 Nilpotent Groups

In this section we classify the filled nilpotent groups. The bulk of the work involved here is in determining the filled 2-groups, as it will turn out that there are only two filled nilpotent groups that are not 2-groups.

We briefly recap some notation and basic results. Recall first that a group is nilpotent if and only if it can be expressed as the internal direct product of its Sylow subgroups. (Put more simply, a group is nilpotent if and only if it is a direct product of p-groups.)

For a group G we write G' for the derived subgroup (so $G' = [G, G]$) and $\Phi(G)$ for the Frattini subgroup (the intersection of the maximal subgroups of G). An *extraspecial* group is a non-abelian p-group G with the property that $Z(G) \cong C_p$ and $G/Z(G)$ is elementary abelian. We will be concerned only with the extraspecial 2-groups. Standard results from group theory tell us that G' is the smallest normal subgroup of G with abelian quotient, and that, since G is a p-group, $\Phi(G)$ is the smallest normal subgroup with elementary abelian quotient. It follows that a 2-group G is extraspecial if and only if $Z(G) = G' = \Phi(G) \cong C_2$.

It is known (see for example [9]) that the order of any extraspecial 2-group is an odd power of 2, and there are exactly two nonisomorphic extraspecial 2-groups for any given odd power of 2. To describe these, recall the construction of a *central product*. A central product $A * B$ is the quotient of the direct product $A \times B$ by a central subgroup of A and B. If G is isomorphic to $A * B$, then it has normal subgroups

which we may identify with A and B, such that $[A, B] = 1$ and $A \cap B \leq Z(G)$. The extraspecial groups of order 8 are D_8 and Q_8. If E_1 and E_2 are the extraspecial groups of order 2^{2n-1}, for $n \geq 1$, then the extraspecial groups of order 2^{2n+1} are isomorphic to $E_1 * Q_8$ and $E_2 * Q_8$.

Our first result classifies the 2-groups all of whose quotients are elementary abelian. This is relevant because every quotient of a filled group must be filled, and it will turn out that all but finitely many filled 2-groups are elementary abelian.

Theorem 8.5. *Suppose every proper nontrivial quotient of a finite nontrivial 2-group G is elementary abelian. Then G is either elementary abelian, extraspecial, C_4 or of the form $E * C_4$ where E is extraspecial and $|G| = 2|E|$.*

The proof of this result by Hart is elementary and we refer the interested reader to our joint paper [3] for details. We continue with a few technical lemmas, and again we refer to the joint paper for proofs.

Lemma 8.6. *Let G be a group of the form $E * C_4$ where E is extraspecial and $|G| = 2|E|$. Then G is not filled.*

A group G of order p^m is said to be of *maximal class* if $m > 2$ and the nilpotence class of G is $m - 1$. It is well known (for example see Theorem 1.2 and Corollary 1.7 of [9]) that the 2-groups of maximal class are dihedral, semidihedral and generalised quaternion (dicyclic). Moreover [9, Theorem 1.2] if G is a 2-group of maximal class of order at least 16, then $G/Z(G)$ is dihedral of order $\frac{1}{2}|G|$.

Lemma 8.7. *The only filled 2-group of maximal class is D_8.*

For a p-group G, we define $c_n(G)$ to be the number of subgroups of G of order p^n.

Theorem 8.8 (Theorem 1.17 of [9])**.** *Suppose a 2-group G is neither cyclic nor of maximal class. Then $c_1(G) \equiv 3 \mod 4$ and for $n > 1$, $c_n(G)$ is even.*

We are now ready to state a crucial corollary which will allow us to classify the filled 2-groups.

Corollary 8.9. *Suppose G is a filled group of order 2^n, where $n > 1$. If the only filled groups of order 2^{n-1} are elementary abelian or extraspecial, then G is either elementary abelian, extraspecial or the direct product of a filled extraspecial group of order 2^{n-1} with a cyclic group of order 2. If the only filled groups of order 2^{n-1} are elementary abelian, then G is either elementary abelian or extraspecial.*

Proof (Hart). Note first that if G is a filled 2-group, then $G' = \Phi(G)$. Suppose the only filled groups of order 2^{n-1} are elementary abelian or extraspecial. If n is 2 or 3, then the result holds, so we may assume $n \geq 4$. Now G is clearly not cyclic. Moreover, by Lemma 8.7, G is not of maximal class. Therefore G has an even number of subgroups of order 4. The length of any conjugacy class of subgroups of order 4 is either 1 or even. The composition factors of any 2-group are cyclic of order 2, and hence G has at least one normal subgroup of order 4. Therefore G has at least two normal subgroups, H and K say, of order 4. Any nontrivial normal subgroup intersects the centre of G nontrivially, and so H contains a central involution z. The quotient $G/\langle z\rangle$ is filled of order 2^{n-1} and so, by hypothesis, either elementary abelian or extraspecial. Hence G/H, which is isomorphic to $\frac{G/\langle z\rangle}{H/\langle z\rangle}$, is a nontrivial quotient of an extraspecial or elementary abelian 2-group, and is therefore elementary abelian. Similarly G/K is elementary abelian. This implies that $G' = \Phi(G) \leq H \cap K$. If G is abelian, then G is elementary abelian.

If G is non-abelian, then $G' = \Phi(G) = \langle z\rangle$, where z is a central involution. If $Z(G)$ contains an involution t other than z, then since t is not contained in $\Phi(G)$, there is a maximal subgroup N which does not contain t. Thus $G \cong N \times \langle t\rangle$. Now $G/\langle t\rangle \cong N$, which forces N to be filled of order 2^{n-1}. So G is elementary abelian unless there is a filled extraspecial group E of order 2^{n-1}, in which case we also have the possibility that $G \cong E \times C_2$. We now deal with the case that z is the only central involution. In that case, since every nontrivial normal subgroup intersects $Z(G)$ nontrivially, every nontrivial normal subgroup contains z and hence every proper quotient is elementary abelian. Therefore, by Theorem 8.5 and Lemma 8.6, G is either elementary abelian or extraspecial.

We have shown that G is either elementary abelian or extraspecial, except in the case where there is a filled extraspecial group E of order 2^{n-1}, in which case we have the further possibility that $G \cong E \times C_2$. □

The strategy now is to understand which extraspecial 2-groups are filled. We begin with a computer-assisted classification of small 2-groups, aided by the following lemmas.

Lemma 8.10. *Let S be a locally maximal product-free set in a group G. If $a \in S$ but $a^{-1} \notin S$, then $a^{-1} \in SS \cup \sqrt{S}$.*

Lemma 8.11. *Suppose G is a group of exponent 4 all of whose elements of order 4 square to the same central involution z. If S is a locally maximal product-free set that does not fill G, then S contains z and every element of S is an involution.*

Proposition 8.12. *If G is a non-abelian filled 2-group of order up to 128, then G is either D_8, $D_8 \times C_2$, $D_8 * Q_8$ or $(D_8 * Q_8) \times C_2$.*

Proof. Computer search allows us to show that the only non-abelian filled 2-groups of order up to 32 are D_8, $D_8 \times C_2$ (fitting in with Corollary 8.9) and $D_8 * Q_8$.
Corollary 8.9 tells us that the only candidates for filled groups of order 64 are C_2^6 and $(D_8 * Q_8) \times C_2$. Lemma 8.11 allows us to reduce the work involved in checking that $(D_8 * Q_8) \times C_2$ is filled, by checking only sets of involutions. By this means, it is then possible to check by machine that $(D_8 * Q_8) \times C_2$ is indeed the only filled non-abelian group of order 64. By restricting the search to non-abelian groups whose quotients are filled and looking only at product-free sets consisting of involutions, computer search also confirmed that there are no non-abelian filled groups of order 128. See Section 8.3 for more detail on the algorithms used. □

The remaining argument centres on the following result.

Theorem 8.13. *If G is an extraspecial group of order greater than 128, then G is not filled.*

The proof of this result by Hart [3] relies on construction of a particular non-filling locally maximal product-free set. The details are lengthy, though elementary, so we refer the interested reader to the joint paper.

Now we are in a position to classify all the filled 2-groups.

Corollary 8.14. *Let G be a 2-group. Then G is filled if and only if G is either elementary abelian, or one of D_8, $D_8 \times C_2$, $D_8 * Q_8$ or $(D_8 * Q_8) \times C_2$.*

Proof. The proof is immediate from Corollary 8.9, Proposition 8.12 and Theorem 8.13. □

A complete classification of the filled 2-groups now allows us to extend the result to all nilpotent groups.

Theorem 8.15. *Let G be a finite nilpotent group. Then G is filled if and only if G is either an elementary abelian 2-group or one of C_3, C_5, D_8, $D_8 \times C_2$, $D_8 * Q_8$ or $(D_8 * Q_8) \times C_2$.*

Proof (Hart). Suppose G is filled and nilpotent. Then G is the direct product of its Sylow subgroups. Therefore for any prime p dividing $|G|$, G has a normal subgroup N

of index p. Hence, by Theorem 8.1(ii) and (iii), p is one of $2, 3$ or 5. If $p = 3$, then by Theorem 8.1(iv), G must be cyclic of order 3. So we can assume the only primes dividing $|G|$ are 2 and 5. If $p = 5$ and 25 divides $|G|$ then G has a normal subgroup of index 25, but by Theorem 8.1*(iii)* there are no filled groups of order 25, a contradiction. Therefore the normal subgroup N of index 5 in G is either trivial or a 2-group. Either way, N contains no elements of order 5. Hence, by Theorem 8.1*(v)*, G must be cyclic of order 5. The only remaining possibility is that G is a 2-group. Theorem 8.15 now follows from Corollary 8.14. □

8.2.3 Groups of order $2^n p$

In this section we show that if G is a group of order $2^n p$, where n is a positive integer and p is an odd prime, then G is filled if and only if G is D_6, D_{10}, D_{12}, D_{14} or D_{22}.

Lemma 8.16. *Let p be an odd prime and let k be an integer satisfying $k > \sum_{r=1}^{\infty} \left\lfloor \frac{p}{2^r} \right\rfloor$. Let G be a group of order $2^k p$. Then G contains a non-trivial normal elementary abelian 2-subgroup of order no greater than 2^p.*

Proof. We show first that G contains some non-trivial normal 2-subgroup N. Consider the set S_2 of Sylow 2-subgroups of G. By the Sylow theorems, either $|S_2| = 1$ or $|S_2| = p$. If $|S_2| = 1$ we take N to be the unique Sylow 2-subgroup. If $|S_2| = p$ then G acts transitively by conjugation on the set S_2, and so the kernel N of this action is a normal subgroup of G which is a 2-group. The condition on k ensures that G is sufficiently large that N must be non-trivial.

It is a fundamental result that a minimal normal subgroup of a solvable group is elementary abelian. Thus N contains some non-trivial elementary abelian 2-subgroup K which is normal in G. Now K is a union of conjugacy classes of G. Since $|K|$ is even and contains the conjugacy class $\{1\}$, it must contain some other conjugacy class T of odd length. Since $|T|$ must divide $|G|$, we conclude that either $|T| = 1$ or $|T| = p$. In either case, $\langle T \rangle$ is a normal 2-subgroup of G of order at most 2^p, as required. □

Corollary 8.17. *For any $k \geq 3$, there is no filled group of order 3×2^k.*

Proof. We proceed by induction. By computer search, we know there are no filled groups of order 24, 48 or 96. So the statement is true for $k = 3, 4, 5$. Suppose the statement is true up to $k \geq 5$ and consider the case $k + 1$. If G is a group of order $3 \times 2^{k+1}$, then by Lemma 8.16 it contains a normal subgroup H of order 2, 4 or 8. Then G/H has order $3 \times 2^{k-2}$, $3 \times 2^{k-1}$ or 3×2^k and so is not filled by the induction hypothesis. Thus G is not filled. □

Corollary 8.18. *For any $k \geq 2$, there is no filled group of order 5×2^k.*

Proof. By computer search (see Section 8.3 for details) we know there are no filled groups of order 20, 40, 80, 160 or 320. So the statement is true for $k = 2, 3, 4, 5, 6$. Suppose the statement is true up to $k \geq 6$ and consider the case $k+1$. If G is a group of order $5 \times 2^{k+1}$, then by Lemma 8.16 it contains a normal subgroup H of order 2, 4, 8, 16 or 32. Then G/H has order $5 \times 2^{k-4}$, $5 \times 2^{k-3}$, $5 \times 2^{k-2}$, $5 \times 2^{k-1}$ or 5×2^k and so is not filled by the induction hypothesis. Thus G is not filled. □

Corollary 8.19. *For any $k \geq 2$, there is no filled group of order 7×2^k.*

Proof. By computer search, we know there are no filled groups of order 28, 56, 112, 224, 448, 896 or 1792. So the statement is true for $k = 2, 3, 4, 5, 6, 7, 8$. Suppose the statement is true up to $k \geq 8$ and consider the case $k+1$. If G is a group of order $7 \times 2^{k+1}$, then by Lemma 8.16 it contains a normal subgroup H of order 2, 4, 8, 16, 32, 64 or 128. Then G/H has order $7 \times 2^{k-6}$, $7 \times 2^{k-5}$, $7 \times 2^{k-4}$, $7 \times 2^{k-3}$, $7 \times 2^{k-2}$, $7 \times 2^{k-1}$ or 7×2^k and so is not filled by the induction hypothesis. Thus G is not filled. □

Lemma 8.20. *Suppose G is a filled group of order $2^n p$, where $n \geq 2$ and p is an odd prime. If G has a normal subgroup of order p, then G contains a central involution.*

Proof. Suppose N is normal of order p in G. Then $G = NH$ where H is any Sylow 2-subgroup of G. This means $G/N \cong H$. Since G is filled, G/N must be filled. By Corollary 8.14 H is either an elementary abelian 2-group, or D_8, $D_8 \times C_2$, $D_8 * Q_8$ or $D_8 * Q_8 \times C_2$. Since H has order at least 4, it follows that either H contains a Klein 4-group $K = \langle a, b \rangle$ such that K is central in H, or H contains a subgroup D which is dihedral of order 8, whose centre is also the centre of H. In the first scenario, consider the action of H on N by conjugation. Write $N = \langle x \rangle$. Now $axa^{-1} = x^i$ for some i, and $x = a^2xa^{-2} = x^{i^2}$. Thus $i = \pm 1$ (because in the cyclic group of units of $\mathbb{Z}/p\mathbb{Z}$ the element 1 has exactly 2 square roots). If $axa^{-1} = x^{-1}$ and $bxb^{-1} = x^{-1}$, then $(ab)x(ab)^{-1} = x$. Therefore at least one involution g in K centralises x. This means we have $g \in Z(H) \cap C_G(N)$. Thus $g \in Z(G)$. Now consider the second situation, where H contains a subgroup D which is dihedral of order 8 whose centre is also the centre of H. We have $D = \langle r, s : r^2 = s^2 = (rs)^4 = 1 \rangle$. Again looking at the action on N by conjugation, we have that $rxr^{-1} = x^{\pm 1}$ and $sxs^{-1} = x^{\pm 1}$, which implies $(rs)x(sr)^{-1} = x^{\pm 1}$. Let $g = (rs)^2$. Then $g \in Z(H)$ and $gxg^{-1} = x$, so $g \in C_G(N)$. Hence again G contains a central involution. □

Proposition 8.21. *For any $k \geq 2$, there is no filled group of order 11×2^k.*

Proof. We proceed by induction on $k \geq 2$. Computer search shows there is no filled group of order 44. Let G be a group of order 11×2^k for $k > 2$ and suppose for a contradiction that G is filled. If G has a normal Sylow 2-subgroup then the quotient of G by this subgroup would be filled of order 11, which is impossible. So we can assume G does not have a normal Sylow 2-subgroup. If G has a normal Sylow 11-subgroup N, then by Lemma 8.20 G contains a central involution g. The quotient $G/\langle g \rangle$ is filled of order $11 \times 2^{k-1}$. By induction $G/\langle g \rangle$ is not filled, and so G cannot be filled. Suppose then that the Sylow subgroups are not normal. The number of Sylow 11-subgroups divides 2^k and is congruent to 1 modulo 11. So the first time this can arise is when $k = 10$. A simple counting argument shows that any group of order 11×2^{10} has either a normal Sylow 11-subgroup or a normal Sylow 2-subgroup, so there is nothing to check here. There is one group of order $2^1 1 \times 11$ with non-normal Sylow subgroups, and four such groups of order $2^{12} \times 11$. The package GrpConst in GAP [35] allows the user to construct all solvable groups of given order, and the function FrattiniExtensionMethod restricts to those groups with only non-normal Sylow subgroups. Thus, even though these five groups are not contained in the Small Groups library of GAP [35], they can be constructed and tested using the methods described in Section 8.3. The upshot is that no group of order 11×2^{10}, 11×2^{11} or 11×2^{12} is filled. We may therefore assume $k \geq 13$. By Lemma 8.16 there is a normal elementary abelian 2-subgroup N of G with order at most 2^{11}. Thus G/N is filled of order 11×2^m where $m \geq 2$. Hence G is not filled. The result now follows by induction. □

Theorem 8.22. *Let G be a group of order $2^n p$ where $n \geq 1$ and p is an odd prime. If G is filled, then G is one of D_6, D_{10}, D_{14} or D_{22}.*

Proof. We have dealt with $p = 3, 5, 7, 11$. It only remains to show that if $p > 11$ there are no filled groups of order $2^n p$. We proceed by induction on n. If $n = 1$ the result holds by Theorem 8.4. Suppose $n \geq 2$. Let N be a minimal normal subgroup of G. Then N is either cyclic of order p or an elementary abelian 2-group. If N is cyclic of order p then by Lemma 8.20, G has a central involution g. Now $G/\langle g \rangle$ has order $2^{n-1}p$, so by inductive hypothesis is not filled. Hence G is not filled. So assume N is an elementary abelian 2-group. Then G/N is either cyclic of order p or has order $2^m p$ where $1 \leq m < n$. In either case, since $p > 11$, we know that G/N is not filled. Therefore G is not filled. By induction no group of order $2^n p$ is filled, when $p > 11$. This completes the proof. □

8.3 Groups of order up to 2000

In this section we describe the computer algorithms used to determine the filled status of a group. These algorithms are implemented in GAP [35] and allow us to test all groups in the library of small groups up to order 2000.

The first algorithm attempts to find a locally maximal product-free set in a given group G which does not fill G. The strategy is to repeatedly add elements at random to a product-free set S until S is maximal. At each stage we keep track of the set F of elements which could be added to S to keep it product-free. If our maximal set S fills G we discard it and start again, returning the first set S found which does not fill G. Note that by Lemma 8.11, if G is an extraspecial 2-group we may begin each search by placing the unique central involution in S. In practice if this algorithm fails to return a result in a reasonable time we abort and use the exhaustive search method of Algorithm 8.3.2.

Algorithm 8.3.1 Find a non-filling locally maximal product-free set for a group G

```
function NFS(G)
    repeat
        if G is an extraspecial 2-group then
            S ← Z(G) \ {1}
        else
            S ← ∅
        end if
        F ← G \ ({1} ∪ S ∪ √S)
        repeat
            x ← Random(F)
            S ← S ∪ {x}
            F ← F \ (S ∪ SS ∪ SS^-1 ∪ S^-1S ∪ √S)
        until F = ∅
    until {1} ∪ S ∪ SS ≠ G
    return S
end function
```

The second algorithm performs an exhaustive search of maximal product-free sets S in a group G and tests whether any fails to fill G. This algorithm is very expensive, and is only required when the random method of Algorithm 8.3.1 has failed to return a result in a reasonable time. The key to making this algorithm run efficiently is the observation that if ϕ is an automorphism of G, then S is a locally maximal product-free subset of G if and only if $\phi(S)$ is locally maximal product-free. Thus the problem of testing all possible sets S is reduced to testing only orbit representatives under the action of the automorphism group of G.

While these orbits can be readily found using GAP, in practice computing orbits of all

possible subsets of G is still prohibitively expensive. To get round this problem, we begin by computing orbits of all product-free sets S of size 3. From each orbit we choose the minimal representative set (with respect to some arbitrary ordering of the elements of G). For each such representative set S, we then try to extend S in all possible ways to obtain a locally maximal product-free set and test whether each possible extension fills our group G. We need only consider extensions using the set F of elements larger than any currently in our set S and which keep S product-free, so again we keep track of this set. Each time we add a new element x to S, we test whether S is locally maximal and if not, we (recursively) extend this new set. The algorithm terminates when either a non-filling locally maximal set has been found, or all possible sets have been examined.

Algorithm 8.3.2 Exhaustive search for locally maximal product-free sets

```
function ExhaustiveSearch(G)
    O ← set of orbit representatives of product-free sets of 3 elements of G under
action of Aut(G)
    for each S ∈ O do
        F ← G \ (S ∪ SS ∪ SS^{-1} ∪ S^{-1}S ∪ √S)
        if not ExtendPFS(G,S,F) then
            return false
        end if
    end for
    return true
end function

function ExtendPFS(G,S,F)
    if F = ∅ then
        if S ∪ SS ∪ SS^{-1} ∪ S^{-1}S ∪ √S = G then
            if {1} ∪ S ∪ SS ≠ G then
                return false
            end if
        end if
    else
        for each x ∈ F do
            S' ← S ∪ {x}
            F' ← {f ∈ F | f > x} \ (S' ∪ S'S' ∪ S'S'^{-1} ∪ S'^{-1}S' ∪ √S')
            if not ExtendPFS(G,S',F') then
                return false
            end if
        end for
    end if
    return true
end function
```

Our final algorithm is used to determine whether a given group G is filled. It uses the results from previous sections to exclude most groups without the need to resort to construction of non-filling sets. For those groups which cannot be excluded in this

way, we use the random method of Algorithm 8.3.1 to find a non-filling set. If all else fails, we resort to the exhaustive method of Algorithm 8.3.2.

We begin by defining the set $\mathcal{G}$ of filled groups of order at most 32, as given in [4, Table 1]. For larger groups we then apply the simple tests using Theorems 8.1(iii), 8.1(vi), 8.1(vii), 8.4 and 8.13. If these are not sufficient to determine the status of our group we examine its normal subgroups and invoke Theorems 8.1(ii), 8.1(iv) and 8.1(v). Finally, if the status of the group is still not resolved we use Algorithms 8.3.1 then 8.3.2 to search for non-filling sets.

Using this method we have examined all groups in the small groups library in GAP up to order 2000. The only filled groups are those noted in [4, Table 1] plus the group $(D_8 * Q_8) \times C_2$ of order 64 and the elementary abelian 2-groups. We conclude this section with the following conjecture.

Conjecture 8.23. *Let G be a finite group. Then G is filled if and only if G is either an elementary abelian 2-group or one of C_3, C_5, D_6, D_8, D_{10}, D_{12}, D_{14}, $D_8 \times C_2$, D_{22}, $D_8 * Q_8$ or $(D_8 * Q_8) \times C_2$.*

Algorithm 8.3.3 Test whether a group G is filled

```
function Filled(G)
    n ← |G|
    if n ≤ 32 then
        if G ∈ 𝒢 then
            return true
        else
            return false
        end if
    else if n is odd then
        return false
    else if G is elementary abelian then
        return true
    else if n = 2^k where k > 7 then
        return false
    else if n = 2^k p where k > 0 and p is an odd prime then
        return false
    else if G is abelian, dihedral or generalised quaternion then
        return false
    else
        for each proper non-trivial normal subgroup N ◁ G do
            if [G : N] = 3 or [G : N] = 5 and not all elements of order 5 are in N
then
                return false
            end if
            if not Filled(G/N) then
                return false
            end if
        end for
    end if
    if NFS(G) succeeds then
        return false
    else
        return ExhaustiveSearch(G)
    end if
end function
```

8.4 Application to the degree-diameter problem

We end this chapter with an application of our filled group investigations to the asymptotics of the degree-diameter problem. In particular, we consider diameter two Cayley graphs of elementary abelian 2-groups. In the class of such graphs, a folklore result yields the bound $L^+(2) \geq \frac{1}{4}$.

To see why this holds, we let $F = GF(2^m)$ and let $G = F^+ \times F^+$. Then G is an elementary abelian 2-group of order 2^{2m}. For our connection set S we take $S = \{(a, 0) : a \in F^*\} \cup \{(0, a) : a \in F^*\}$. Clearly any element of G can be expressed as the sum of at most 2 elements of S. Thus $\mathrm{Cay}(G, S)$ has diameter 2, order 2^{2m} and degree $|S| = 2(2^m - 1)$.

Although this construction is entirely elementary, there is as far as we know no better published result for this class of graphs. From our knowledge of filled groups, we know that $\mathrm{Cay}(G, S)$ has diameter 2 if and only if S fills G. We also know from Theorem 8.1*(iii)* that any locally maximal sum-free set of an elementary abelian 2-group is filling. The question now is whether we can find a family of "small" locally maximal sum-free sets in an elementary abelian 2-group to improve the bound. Our main result will be the following.

Theorem 8.24. *In the class of Cayley graphs of elementary abelian 2-groups,*

$$L^+(2) \geq \frac{64}{225}$$

To prove our result we first need to relate our locally maximal sum-free sets to certain sets in projective space. Let q be a prime power and let $F = GF(q)$. Recall that the projective space $PG(n, q)$ can be thought of as consisting of a set of points which we identify with the 1-dimensional subspaces of F^{n+1} and a set of lines which we identify with the 2-dimensional subspaces. Incidence of points and lines is defined by subspace inclusion in the natural way.

In projective space $PG(n, q)$, a set of k points is called a *k-cap* if no 3 points are collinear. It is called a *complete* k-cap if it is contained in no $k + 1$-cap. Our key observation is the following.

Observation 8.25. *A locally maximal sum-free set in $\mathbb{Z}_2^n$ corresponds to a complete cap in $PG(n - 1, 2)$.*

Proof. Let $F = GF(2)$ and let a, b be two points of $PG(n - 1, 2)$. We think of a and b as 1-dimensional subspaces of F^n and may identify them with the unique non-zero

vector in each subspace. The only other point collinear with a and b is the subspace identified with the vector $a + b$. Since F^n can be thought of as an elementary abelian 2-group of order 2^n, the result follows. □

The question now is how small can a complete k-cap be in $PG(n-1, 2)$? At first sight we have done nothing more than restate our problem in different terms. However, complete caps in projective space are studied in their own right and we have the following result.

Theorem 8.26 (Gabidulin, Davydov, Tombak [34])**.** *For $m \geq 5$, there exists a complete k-cap in $PG(2m-1, 2)$ with $k = 15 \times 2^{m-3} - 3$.*

This immediately allows us to prove our main result.

Proof of Theorem 8.24. Let $m \geq 5$. By Theorem 8.26 and Observation 8.25, there exists a locally maximal sum-free set S in $G = \mathbb{Z}_2^{2m}$ of size $15 \times 2^{m-3} - 3$. Since S is filling by Theorem 8.1*(iii)*, $\mathrm{Cay}(G, S)$ has diameter 2. Since every element of G has order 2, $S = S^{-1}$ and so $\mathrm{Cay}(G, S)$ is an undirected graph. The graph has order 2^{2m} and asymptotic degree $15 \times 2^{m-3}$. Since we can construct such a graph for infinitely many degrees, the result follows. □

We note that in the paper of Gabidulin, Davydov and Tombak [34], the complete caps emerge as a by-product of the creation of interesting linear codes of covering radius 2. For example with $m = 5$ the authors define the following matrix:

$$H = \begin{bmatrix}
00000000000000000000000000000000000000000001111111111111111 \\
00011111111000000000000000000000000000000000000011111111 \\
00000001111111111111111111111111111111111110000111100001111 \\
01100110011011001100110000011001100110000001100110011001 1 \\
10101010101101010101010000101010101010000010101010101010 1 \\
11111110000011010101100011100101010011100000000000000000 \\
11100001111110001101010101001110010101010111111111111111 \\
11111110000000000011111111111111100000000000000000000000 \\
11111110000000111000011111110000111100000000000000000000 \\
11111110000000000000000001111111111111110000000000000000
\end{bmatrix}$$

They then show that H is the check matrix of a linear code of length 57, dimension 47 and minimum distance 4. The columns of H form the complete cap which we interpret as a maximal sum-free set in $\mathbb{Z}_2^{10}$ of size 57.

It is interesting that the solution of this problem in coding theory turns out to have an application in the degree-diameter problem, via the seemingly only loosely related topics of projective geometry and sum-free sets in groups.

Our result in Theorem 8.24 is as far as we know the first specific bound for the class of Cayley graphs of elementary abelian 2-groups. However we can use similar ideas to those used in our investigations of circulant graphs in Chapter 3 to obtain a result for diameter 3 graphs in this class.

Theorem 8.27. *In the class of Cayley graphs of elementary abelian 2-groups,*

$$L^+(3) \geq \frac{1}{16}.$$

Proof. Let $H = \mathbb{Z}_2^m$ for some m. (We view H equivalently as either a vector space of dimension m over $GF(2)$ or as the additive group of $GF(2^m)$.) Let $G = (H \times H \times H) \times (\mathbb{Z}_2 \times \mathbb{Z}_2)$. We define our generating set S to contain the following.

$$\begin{array}{ll} (x,0,0;0,0) & x \in H, x \neq 0 \\ (0,x,0;0,1) & x \in H \\ (0,0,x;1,0) & x \in H \\ (x,x,x;1,1) & x \in H \end{array}$$

All elements of G are self-inverse so S is an inverse-closed set of cardinality $2^{m+2} - 1$. All elements of G can be expressed as a sum of 3 elements from S, so $\mathrm{Cay}(G, S)$ has diameter 3. So for any $d = 2^{m+2} - 1$ we have $n(d, 3) \geq 2^{3m+2}$ and so $L^+(3) \geq \frac{1}{16} = 0.06250$. □

Chapter 9

Arc-transitive graphs in the degree-diameter problem

Chapters 3, 4, 5 and 6 explored the degree-diameter problem for various classes of undirected, directed and mixed graphs. We conclude our investigations in the degree-diameter problem by turning to a class of graphs which has been much less explored in the literature, and consider the problem for undirected arc-transitive graphs.

The only published result in the literature appears to be a recent result of Zhou [76] for arc-transitive graphs of diameter 2. This construction yields an infinite family of such graphs of diameter 2, but for degree d the exponent in the asymptotic order is $5/3$ rather than 2. Thus, in our usual notation, this construction does not yield a useful lower bound for $L^+(2)$ in the class of arc-transitive graphs. To obtain such a bound we would need to find a construction with an exponent of 2 in the diameter 2 case. We now describe a construction which in fact yields an exponent of k in the general case of diameter $k \geq 2$.

Consider a graph Γ with vertex set equal to the set of all words of length k over an alphabet of size $n > 1$. Two vertices of Γ are adjacent if they differ in exactly one coordinate position (that is, the *Hamming distance* between words is exactly 1). Clearly this graph has diameter k. Less immediately clear is that it is arc-transitive.

To see why, note first that any permutation of the n symbols in a given coordinate position induces an automorphism of the graph. Also, permuting the coordinate positions gives an automorphism. (In fact the wreath product $S_n \wr S_k$ is a subgroup of $\mathrm{Aut}(\Gamma)$.)

Using the numbers $1..n$ as the symbols of the alphabet, it suffices to show that there is an automorphism mapping the arc from $11\ldots11$ to $11\ldots12$ to any other arc x. If the initial vertex of x is $x_1x_2\ldots x_k$ then the terminal vertex differs from this in exactly one position.

Assume first that the terminal vertex of x has the form $x_1x_2\ldots y$ where $y \neq x_k$. We create an automorphism of the graph by using a product of the following k

automorphisms:

1. For $i < k$, any permutation in coordinate position i sending 1 to x_i
2. Any permutation in coordinate position k sending 1 to x_k and 2 to y (since $x_k \neq y$ we can always do this)

The other case to consider is where $k > 2$ and the vertices of x differ in a different coordinate position, say the second. So the terminal vertex of x has the form $x_1 y x_3 \dots x_k$ where $y \neq x_2$. This time our automorphism is a product of:

1. Any permutation in coordinate position 1 sending 1 to x_1
2. Any permutation in coordinate position 2 sending 1 to x_3
3. Any permutation in coordinate position 3 sending 1 to x_2 and 2 to y (since $x_2 \neq y$ we can always do this)
4. Swapping coordinate positions 2 and 3

This completes the proof that Γ is arc-transitive. So for all degrees $d = k(n-1)$, we can construct an arc-transitive graph of order n^k and diameter k. This immediately yields the main result of this section.

Theorem 9.1. *Let $k \geq 2$. Then in the class of undirected arc-transitive graphs,*

$$L^+(k) \geq \frac{1}{k^k}$$

Alternatively, we can view this graph as a Cayley graph of $G = \mathbb{Z}_n^k$. Our generating set is:

$$S = \bigcup_{x \neq 0} \{(x, 0, 0, \dots), (0, x, 0, \dots), \dots\}$$

It is clear that if $s \in S$ and $g \in G$, then $g + s$ differs from g in exactly one coordinate position. Thus $\mathrm{Cay}(G, S)$ is isomorphic to our graph Γ above. This yields the following corollary.

Corollary 9.2. *Let $k \geq 2$. Then in the class of undirected arc-transitive Cayley graphs,*

$$L^+(k) \geq \frac{1}{k^k}$$

We remark here that in the case of arc-transitive graphs, it is more difficult to convert a result on L^+ into one for L^-. In the case of Cayley graphs, our usual trick is to add additional generators so that we obtain a Cayley graph of the same diameter but larger degree. Here we would need to do this in such a way as to preserve arc-transitivity, which is a rather stronger condition.

Chapter 10

Regular maps

Our final main chapter diverges somewhat from the problems of diameters and girths of graphs. We turn to the study of graphs embedded on surfaces, and in particular we focus on regular maps where the embedding in some orientable surface has a particular automorphism group.

10.1 Introduction

We begin with some definitions and basic concepts. An *embedding* of a graph Γ in a surface S is a continuous one-one mapping ϑ from (a topological representation of) Γ to S. A *face* or *region* of the embedding is a connected component of $S \setminus \vartheta(\Gamma)$. By a *map* we mean an embedding of a finite connected graph in some compact and connected surface in which each face is homeomorphic to an open disc in $\mathbb{R}^2$, and we may refer to this as a *cellular* or 2-*cell* embedding.

A vertex-edge-face incident triple in a map $\mathcal{M}$ is called a *flag* of $\mathcal{M}$ (we ignore here some degenerate cases). An automorphism of $\mathcal{M}$ is a permutation of the flags which preserves the incidences between them. The automorphisms form a group $\mathrm{Aut}(\mathcal{M})$ in the natural way. The embedding results in a natural cyclic order of the neighbours of a vertex v around the vertex. Any automorphism ϕ maps these neighbours, in order, to the neighbours of $\phi(v)$, and similar arguments apply to the vertices of a face. It follows by connectedness that an automorphism of $\mathcal{M}$ is completely determined by its action on a single flag of $\mathcal{M}$, and hence we have the well-known fact:

Proposition 10.1. *The automorphism group* $\mathrm{Aut}(\mathcal{M})$ *of a map* $\mathcal{M}$ *acts semi-regularly on the flags of* $\mathcal{M}$*.*

If there is a single orbit of the flags of the map under the action of the automorphism group, we say the map is *regular*. Since each edge of the graph is contained in 4 flags, the order of the automorphism group of a regular map is 4 times the number of edges.

We call a map *orientable* if the surface S is orientable, otherwise *non-orientable*. In the orientable case, an automorphism may either preserve or reverse the orientation of the embedding. If an orientable map admits an orientation-reversing automorphism we say the map is *reflexible*; otherwise *chiral*. The orientation-preserving

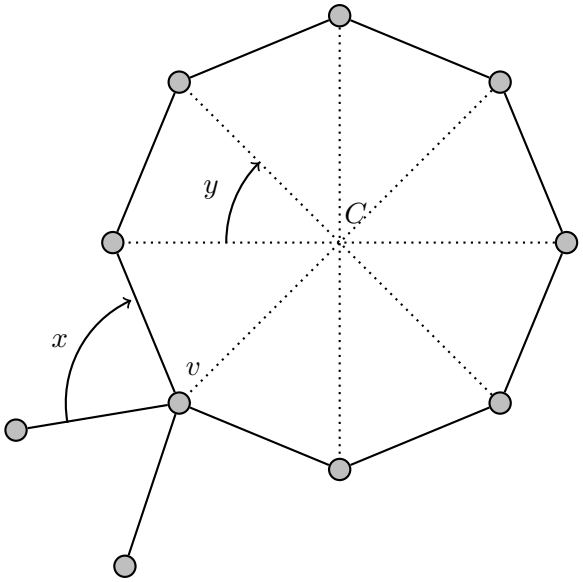

Figure 10.1: Automorphisms of a map

automorphisms form a (not necessarily proper) subgroup $\mathrm{Aut}^+(\mathcal{M})$ of the full group. In the orientation-preserving case, we may think of the flags of the map simply as arcs (ordered pairs of adjacent vertices) since the orientation of the incident faces is fixed relative to an arc. An *orientably-regular* map is a map such that $\mathrm{Aut}^+(\mathcal{M})$ is transitive, and hence regular, on arcs of the embedded graph.

In a regular map, it follows from transitivity that vertices have the same valency, say k, and all faces are bounded by closed walks of the same length, say ℓ; the map is then said to be of type (k, ℓ). The group $\mathcal{A} = \mathrm{Aut}^+(\mathcal{M})$ is generated by two elements x and y of order k and ℓ such that x acts as a clockwise rotation of $\mathcal{M}$ about a vertex v by $2\pi/k$ and y acts as a clockwise rotation by $2\pi/\ell$ about the centre C of a face incident with v (see Figure 10.1). The product xy is then a rotation of $\mathcal{M}$ about the centre of an edge that is incident to both the vertex and the face.

Orientably-regular maps can be viewed as maps having the 'highest level' of orientation-preserving symmetry among general maps. Regularity of $\mathcal{A}$ on the arc set of the embedded graph enables one to identify the map $\mathcal{M}$ with the triple $(\mathcal{A}, x, y)$ in such a way that arcs, edges, vertices and faces correspond to (say, left) cosets of the trivial group (that is, to elements of $\mathcal{A}$) and of the subgroups $\langle xy \rangle$, $\langle x \rangle$ and $\langle y \rangle$ of $\mathcal{A}$. Incidence between arcs, edges, vertices and faces is given by non-empty intersection of the corresponding cosets, and the action of $\mathcal{A}$ on the cosets is simply given by left multiplication.

It follows that orientably-regular maps are, up to isomorphism, in a one-to-one correspondence with equivalence classes of triples (G, x, y), where G is a finite group

admitting a presentation of the form $G = \langle x, y \mid x^k = y^\ell = (xy)^2 = \ldots = 1 \rangle$, with two triples (G_1, x_1, y_1) and (G_2, x_2, y_2) being equivalent if there is a group isomorphism $G_1 \to G_2$ taking x_1 onto x_2 and y_1 onto y_2. This way, investigation of orientably-regular maps can be reduced to purely group-theoretic considerations. The corresponding algebraic theory has been developed in depth by Jones and Singerman in [40].

Since the concept of an orientably-regular map includes the underlying graph, the carrier surface and the supporting automorphism group, classification attempts for such maps in most cases follow one of these three directions. A number of influential results have been obtained in classification of orientably-regular maps in the first two directions; we refer to the recent survey by Širáň [67] for details. Here we focus on the third direction, that is, classification of orientably-regular maps by their automorphism groups, in which results are much less abundant.

The only family of *simple* groups for which a classification of the corresponding orientably-regular maps is known are the groups $PSL(2, q)$ for any prime power $q > 3$ [52, 64, 23]. Classification of orientably-regular maps with automorphism group $PGL(2, q)$, the obvious degree-two extension of $PSL(2, q)$, can be extracted from the corresponding classification for $PSL(2, q^2)$ through the well-understood inclusion $PGL(2, q) < PSL(2, q^2)$, see [64, 23].

For odd q, however, the simple group $PSL(2, q^2)$ admits another interesting extension of degree two, namely, the group $M(q^2)$, also known as a *twisted* linear fractional group. By a classical result of Zassenhaus [75], the groups $PGL(2, q)$ and $M(q^2)$ are the only finite sharply 3-transitive groups (of degree $q + 1$ and $q^2 + 1$, respectively). This motivates the question of classification of orientably-regular maps with automorphism group isomorphic to $M(q^2)$.

In this chapter we present a complete enumeration of (isomorphism classes of) orientably-regular maps with automorphism group isomorphic to $M(q^2)$. The results are strikingly different from those for the groups $PGL(2, q)$ in many ways. To give three examples, we note that (a) all the orientably-regular maps for $PGL(2, q)$ are reflexible, while this is not the case for $M(q^2)$; (b) the groups $PGL(2, q)$ are also automorphism groups of non-orientable regular maps while the groups $M(q^2)$ are not; and (c) for any even $k, \ell \geq 4$ not both equal to 4 there are orientably-regular maps of type (k, ℓ) with automorphism group $PGL(2, q)$ for infinitely many values of q, while for infinitely many such pairs (k, ℓ) there are no orientably-regular maps for $M(q^2)$ of that type for any q.

By our outline and the algebraic theory of [40], enumeration of orientably-regular maps with a given automorphism group G reduces to enumeration of all triples (G, x, y) with $G = \langle x, y;\ x^k = y^\ell = (xy)^2 = \ldots = 1\rangle$ up to conjugation by elements of $\mathrm{Aut}(G)$, that is, by considering triples (G, x, y) and (G, x', y') equivalent if there is an automorphism of G taking (x, y) onto (x', y'). We do this systematically for the twisted linear groups $G = M(q^2)$. In Sections 10.2 and 10.3 we introduce the group $M(q^2)$ and study its subgroups. Sections 10.4, 10.5 and 10.6 deal with identifying 'canonical' forms of elements of G and study their conjugacy in depth. In Sections 10.7, 10.8 and 10.9 we develop arguments for counting 'canonical' pairs of elements of G. All the auxiliary facts are then processed in Section 10.10 to produce our main result:

Theorem. *Let $q = p^f$ be an odd prime power, with $f = 2^\alpha o$ where o is odd. The number of orientably-regular maps $\mathcal{M}$ with $\mathrm{Aut}^+(\mathcal{M}) \cong M(q^2)$ is, up to isomorphism, equal to*

$$\frac{1}{f}\sum_{d|o} \mu(o/d)h(2^\alpha d)\ ,$$

where $h(x) = (p^{2x} - 1)(p^{2x} - 2)/8$ and μ is the Möbius function.

The exposition in this chapter follows closely the structure of our joint paper with Hriňáková and Širáň [30]. However, for brevity we will omit some detailed computations and give only outline proofs of some intermediate results where these were completed principally by the other authors. Full details can be found in [30].

10.2 The twisted linear groups $M(q^2)$

For a finite field F let $S(F)$ and $N(F)$ be the set of non-zero squares and non-squares of F. The general linear group $GL(2, F)$ is the group of all non-singular 2×2 matrices with entries in F; restriction to matrices with determinant 1 gives the special linear group $SL(2, F)$. The groups $PGL(2, F)$ and $PSL(2, F)$, the quotients of $GL(2, F)$ and $SL(2, F)$ by the corresponding centres, are known as the linear fractional groups. They can equivalently be described as groups of all transformations $z \mapsto (az + b)/(cz + d)$ of the set $F \cup \{\infty\}$ (with the obvious rules for calculations with ∞), with $ad - bc \neq 0$ and $ad - bc \in S(F)$ for $PGL(2, F)$ and $PSL(2, F)$, respectively. The group $PSL(2, F)$ is an index 2 subgroup of $PGL(2, F)$ unless F has characteristic 2, in which case the two groups are the same.

Suppose now that F admits an automorphism σ of order 2, which happens if and only

if $|F| = q^2$ for some prime power q, and σ is then given by $x \mapsto x^q$ for every $x \in F$. If, in addition, q is odd, then one may 'twist' the transformations described above by considering the permutations of $F \cup \{\infty\}$ defined by $z \mapsto (az+b)/(zc+d)$ if $ad-bc \in S(F)$ and $z \mapsto (az^\sigma+b)/(cz^\sigma+d)$ if $ad-bc \in N(F)$. These transformations form a group under composition, denoted $M(F)$ or $M(q^2)$, and called the *twisted fractional linear group*. Observe that $PSL(2,F)$ is a subgroup of $M(F)$ of index two, again. By a well-known result due to Zassenhaus [75], the groups $PGL(2,F)$ for an arbitrary finite field F, and $M(F)$ for fields of order q^2 for an odd prime power q, are precisely the finite, sharply 3-transitive permutation groups (on the set $F \cup \{\infty\}$ in both cases).

In this chapter we will focus on the twisted fractional linear groups, with the goal to classify the orientably-regular maps they support. For our purposes, however, it will be useful to work with a different representation of these groups. From this point on, let $F = GF(q^2)$ for some odd prime power q and let $F_0 \cong GF(q)$ be its unique subfield of order q; let F^* and F_0^* be the corresponding multiplicative groups. Further, let σ be the unique automorphism of F of order 2; we have $x^\sigma = x^q$ for any $x \in F$, and $x^\sigma = x$ if and only if $x \in F_0$. If $A \in GL(2,F)$, by A^σ we denote the matrix in $GL(2,F)$ obtained by applying σ to every entry of A.

Let $J = GL(2,F) \rtimes Z_2$, where multiplication in the semidirect product is defined by $(A,i)(B,j) = (AB^{\sigma^i}, i+j)$; equivalently, J is an extension of $GL(2,F)$ by the automorphism σ. To introduce a 'twisted' subgroup of J, for every $A \in GL(2,F)$ we define $\iota_A \in Z_2 = \{0,1\}$ by letting $\iota_A = 0$ if $\det(A) \in S(F)$ and $\iota_A = 1$ if $\det(A) \in N(F)$. We now let $K = \{(A, \iota_A);\ A \in GL(2,F)\}$; multiplication in K is, of course, given by $(A,\iota_A)(B,\iota_B) = (AB^{\sigma^{\iota_A}}, \iota_A + \iota_B)$ for any $A, B \in GL(2,F)$. The group K and its quotient groups will be of principal importance in what follows.

Let $K_0 = \{(A,0);\ A \in GL(2,F),\ \iota_A = 0\}$ be the subgroup of K index 2 of K. The centre L of K_0 consists of elements of the form $(D,0)$, where $D \in GL(F)$ is a scalar matrix; obviously L is also a normal subgroup of both K and J. It can be checked that the factor group $G = K/L$ is isomorphic to $M(q^2)$, and since K has index 2 in J, the group $G = M(q^2)$ is (isomorphic to) a subgroup of index 2 of $\overline{G} = J/L$. The group $\overline{G}$ can alternatively be described as $G\langle\sigma\rangle$, the split extension of G by $\langle\sigma\rangle \simeq Z_2$. Observe also that the factor group $G_0 = K_0/L$ is isomorphic to $PSL(2,F)$, and if q is a prime, the group J/L is isomorphic to $P\Gamma L(2,q^2)$.

Elements $(A,i)L$, that is, cosets $\{(\delta A, i);\ \delta \in F^*\}$, of the factor groups $G = K/L$ and $\overline{G} = J/L$ will throughout be denoted $[A,i]$; they will be called *untwisted* if $i = 0$ and *twisted* if $i = 1$.

For our final enumeration it will be necessary to determine the automorphism group of $M(q^2)$. While the result appears to be 'obvious' we provide a simple proof based on a fact which may be folklore to group-theorists.

Lemma 10.2. *Let U be a characteristic subgroup of a group $\tilde{U}$ of index 2. Suppose that the centre of U is trivial and every automorphism of U extends to an automorphism of $\tilde{U}$. Then* $\mathrm{Aut}(U) \cong Aut(\tilde{U})$.

Proof (Širáň). The assumption of U being characteristic in $\tilde{U}$ implies that every $h \in \mathrm{Aut}(\tilde{U})$ restricts to an $h_U \in \mathrm{Aut}(U)$. Since each automorphism of U extends to an automorphism of $\tilde{U}$, the assignment $h \mapsto h_U$ is a group epimorphism $\vartheta :\ \mathrm{Aut}(\tilde{U}) \to \mathrm{Aut}(U)$. Suppose that $h \in \mathrm{Aut}(\tilde{U})$ is in the kernel of ϑ, so that h_U is the identity mapping on U. For every $x \in U$ and every $y \in \tilde{U}\backslash U$ we have $y^{-1}xy \in U$ and hence $y^{-1}xy = h(y^{-1}xy) = h(y)^{-1}xh(y)$, which implies that $h(y)y^{-1}$ commutes with x for all $x \in U$. Observe that $h(y)y^{-1} \in U$, since U was assumed to have index 2 in $\tilde{U}$. By triviality of the centre of U we have $h(y)y^{-1} = 1$ and as this is valid for all $y \in \tilde{U}\backslash U$ we conclude that h is the identity on $\tilde{U}$. It follows that the kernel of ϑ is trivial and so $\mathrm{Aut}(U) \cong Aut(\tilde{U})$. □

We now apply Lemma 10.2 to $U = G_0 = PSL(2, q^2)$ and $\tilde{U} = G = M(q^2)$ for $q = p^f$, where p is an odd prime and f a positive integer. Being a simple subgroup of $M(q^2)$, the group G_0 is characteristic (and of index two) in G. It is well known (see e.g. [36]) that $\mathrm{Aut}(G_0) \cong P\Gamma L(2, q^2) \simeq PGL(2, q^2) \rtimes \mathbb{Z}_{2f}$, with an element $(C, \varphi) \in PGL(2, q^2) \rtimes \mathbb{Z}_{2f}$ acting on G_0 by $X \mapsto (C^{-1}XC)^{\varphi}$. Now, any (C, φ) is easily seen to extend to G by $[X, \iota_X] \mapsto ([C, 0]^{-1}[X, \iota_X][C, 0])^{\varphi}$. By Lemma 10.2 we now obtain:

Proposition 10.3. *The automorphism group of $M(q^2)$ is isomorphic to $P\Gamma L(2, q^2)$.*

10.3 Twisted subgroups of $M(q^2)$

Let $q = p^f$ for an odd prime p and a positive integer f; these will be fixed throughout. In this section we will focus on the *twisted* subgroups of $M(p^{2f})$, that is, those isomorphic to $M(p^{2e})$ for suitable $e \leq f$. From now on we will use the notation $F_m = GF(p^m)$ for a Galois field of order p^m for $m \leq f$ but keep letting $F = GF(p^{2f})$. We begin by identifying the possible values of e.

Lemma 10.4. *A group $M(p^{2e})$ is isomorphic to a subgroup of $M(p^{2f})$ if and only if e is a divisor of f such that f/e is odd.*

Outline proof. If f/e is odd, we consider the automorphism $x \mapsto x^{p^f}$ of $GF(p^{2f})$ onto $GF(p^{2e})$ and show that its restriction to the subfield $GF(p^{2e})$ corresponds to the map $x \mapsto x^{p^e}$ on the subfield.

For the reverse implication, if $M(p^{2e})$ is a subgroup of $M(p^{2f})$ then e divides f, and we prove f/e is odd by induction on f/e. □

If f/e is odd, a particularly important copy of $M(p^{2e})$ in $M(p^{2f})$ is formed by all the pairs $[X, \iota_X]$ with $X \in GL(2, p^{2e})$ such that all entries of X lie in the subfield F_{2e} of F; this copy will be called *canonical*. The copy of $PSL(2, p^{2e})$ in $M(p^{2f})$ formed by all the pairs $[X, 0]$ with $X \in SL(2, F_{2e})$ will be called *canonical* as well. We now prove a useful auxiliary result on canonical subgroups.

Proposition 10.5. *Let f/e be an odd integer and let $H \cong M(p^{2e})$ be a subgroup of $G = M(p^{2f})$ such that H contains the canonical copy of $PSL(2, p^{2e})$. Then H is equal to the canonical copy of $M(p^{2e})$ in G.*

Proof (Hriňáková, Širáň). Let H be a copy of $M(p^{2e})$ in G such that $H_0 = H \cap PSL(2, p^{2f})$ is equal to the canonical copy of $PSL(2, p^{2e})$ in G. Obviously, H_0 is a normal subgroup of H of index two. Let $[A, 1]$ be an element of $H \backslash H_0$, where A is the 2×2 matrix with rows (a, b) and (c, d) for some $a, b, c, d \in F$ with $\delta = ad - bc \in N(F)$. We may assume that the entry c in the lower left corner of A is non-zero. Indeed, if $c = 0$ and $b \neq 0$, letting D be an off-diagonal matrix with entries -1 and 1 we may replace $[A, 1]$ with the product $[D, 0][A, 1] \in H \backslash H_0$, and if A is a diagonal matrix we may replace $[A, 1]$ with the product $[D', 0][A, 1] \in H \backslash H_0$ for a matrix D' with rows $(1, 0)$ and $(1, 1)$. Then, since we are working with projective groups, we may assume that $c = 1$, so that $\delta = ad - b$.

By our assumption the group H_0 also contains the element $[C, 0]$ with C having rows $(1, 1)$ and $(0, 1)$. Normality of H_0 in H implies that $[A, 1][C, 0][A, 1]^{-1} = [ACA^{-1}, 0] \in H_0$ and also $[A, 1]^{-1}[C, 0][A, 1] = [(A^\sigma)^{-1} C A^\sigma, 0] \in H_0$. Evaluating the products we obtain

$$\varepsilon ACA^{-1} = \begin{pmatrix} a & b \\ 1 & d \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} d & -b \\ -1 & a \end{pmatrix} = \begin{pmatrix} \delta - a & a^2 \\ -1 & \delta + a \end{pmatrix} \text{ , and}$$

$$\varepsilon' (A^\sigma)^{-1} C A^\sigma = \begin{pmatrix} d^\sigma & -b^\sigma \\ -1 & a^\sigma \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} a^\sigma & b^\sigma \\ 1 & d^\sigma \end{pmatrix} = \begin{pmatrix} \delta^\sigma + d^\sigma & (d^\sigma)^2 \\ -1 & \delta^\sigma - d^\sigma \end{pmatrix}$$

for some $\varepsilon, \varepsilon' \in F^*$. Since H_0 is assumed to be equal to the canonical copy $PSL(2, p^{2e})$ in G, all the remaining entries of the two matrices on the right-hand sides

above must lie in F_{2e}. This readily implies that both $a, \delta, d^\sigma \in F_{2e}$, and since F_{2e} is setwise preserved by σ we also have $d \in F_{2e}$ and so $b = ad - \delta \in F_{2e}$ as well. We conclude that $A \in \mathrm{GL}(2, p^{2e})$ and hence the subgroup H generated by $[A, 1]$ and H_0 is identical with the canonical copy of $M(p^{2e})$ in G. □

As a consequence we prove that all twisted subgroups of G are conjugate. Recall that G_0 denotes the (unique) subgroup of G isomorphic to $PSL(2, p^{2f})$.

Proposition 10.6. *If f/e is an odd integer, then all subgroups of $G = M(p^{2f})$ isomorphic to $M(p^{2e})$ are conjugate in G_0.*

Outline proof. We let $H \cong M(p^{2e})$ be a subgroup of G. Consider $H \cap G_0 \cong PSL(2, p^e)$. By the known classification of the subgroups of $PSL(2, p^{2f})$, we may conclude that all such subgroups are conjugate in G_0, then apply Proposition 10.5. □

We conclude with a sufficient condition for a subgroup of G to be twisted; this result will be of key importance later. In order to state it, we will say that a subgroup H of G *stabilises a point* if, in the natural action of G on the set $F \cup \{\infty\}$ via linear fractional mappings from Section 10.2, there exists a point in $F \cup \{\infty\}$ fixed by all linear fractional mappings corresponding to elements of H. Also, for any positive divisor g of $2f$ let G_g be the canonical copy of $PSL(2, p^g)$ in the group $G = M(p^{2f})$. Moreover, if g is even, we let G_g^* denote the copy of $PGL(2, p^{g/2})$ in G_g formed by (equivalence classes of) non-singular 2×2 matrices over $GF(p^{g/2})$.

Proposition 10.7. *Let H be a subgroup of $G = M(p^{2f})$ not contained in the subgroup $G_0 = PSL(2, p^{2f})$ and let $H_0 = H \cap G_0$. If H_0 does not stabilise a point and is neither dihedral nor isomorphic to A_4, S_4 or A_5, then H is conjugate in $\overline{G}$ to a subgroup isomorphic to $M(p^{2e})$ for some positive divisor e of f such that f/e is odd.*

Outline proof. By a summary [48] of Dickson's classification of subgroups of projective special linear groups over finite fields, subgroups of G_0 comprise point stabilisers, dihedral groups, A_4, S_4, A_5, and $PSL(2, p^g)$ for divisors g of $2f$ together with $PGL(2, p^{g/2})$ for even divisors g of $2f$. Our assumptions imply that H_0 must be isomorphic to one of the last two types of subgroups. We now use arguments from the proof of Proposition 10.5 to conclude that $H \cong M(p^{2e})$ for some e. □

10.4 Representatives of twisted elements

A key element of our strategy is to understand the conjugacy classes of twisted elements. In this section and the next, we develop this understanding. Although many of the detailed calculations were completed by the other authors of our joint paper [30], we include these details here because of their importance to the overall result.

Recalling the notation introduced in Section 10.2, we begin by identifying elements in conjugacy classes of $K\backslash K_0$ that have a particularly simple form. To facilitate the description here and also in the sections that follow, we let $\mathrm{dia}(\alpha,\beta)$ and $\mathrm{off}(\alpha,\beta)$, respectively, denote the 2×2 matrix with diagonal entries α,β (from the top left corner) and zero off-diagonal entries, and the 2×2 matrix with off-diagonal entries α,β (from the top right corner) and zero diagonal entries.

For every element $(A,1)\in K\backslash K_0$ we have $(A,1)^2=(AA^\sigma,0)$. In the study of conjugacy in $K\backslash K_0$ it turns out to be important to understand the behaviour of the products AA^σ. Observe that if $\delta=\det(A)\in N(F)$, then $\det(AA^\sigma)=\delta\delta^\sigma\in N(F_0)$.

Let $(A,1)\in K\backslash K_0$ and let $\{\lambda_1,\lambda_2\}$ be the spectrum of AA^σ in a smallest extension F' of F of degree at most two in which σ may still be assumed to be given by $x\mapsto x^q$. Since $A^\sigma A$ is both a conjugate and also a σ-image of AA^σ, we have $\{\lambda_1,\lambda_2\}^\sigma=\{\lambda_1,\lambda_2\}$. This means that either (1) $\lambda_i^\sigma=\lambda_i$ for $i=1,2$, or (2) $\lambda_1^\sigma=\lambda_2$ and $\lambda_2^\sigma=\lambda_1$. Note that (2) implies $\lambda_1^{q^2}=(\lambda_1^q)^q=\lambda_2^q=\lambda_1$ and, similarly, $\lambda_2^{q^2}=\lambda_2$. We conclude that $F'=F$ in both the situations (1) and (2) and so both λ_1,λ_2 are in F. Observe that $\lambda_1\neq\lambda_2$, as otherwise we would have $\lambda_1=\lambda_1^\sigma\in F_0$ and $\det(AA^\sigma)=\lambda_1^2\in S(F_0)$, a contradiction. Moreover, it follows that in the case (1) we have $\lambda_i\in F_0$ for $i=1,2$ with $\lambda_1\lambda_2\in N(F_0)$, and in the case (2) $\lambda_i\in N(F)\subset F\backslash F_0$ since $\det(AA^\sigma)=\lambda_1\lambda_1^\sigma\in N(F_0)$.

We will now refine our considerations of AA^σ. As before, let $q=p^f$ for some odd prime p and let e be the *smallest* positive divisor of f with f/e odd such that $AA^\sigma=\varepsilon C$ for some $C\in\mathrm{SL}(2,p^{2e})$ and for some $\varepsilon\in F^*$. In other words, we look for the smallest subfield F_{2e} of F, with f/e odd, such that all entries of C lie in F_{2e}; note that we may assume C to have determinant 1 since the determinant of AA^σ is a non-zero square of F. If $\{\mu,\mu^{-1}\}$ is the spectrum of C, we have, without loss of generality, $\lambda_1=\varepsilon\mu$ and $\lambda_2=\varepsilon\mu^{-1}$. Observe that since $\lambda_1,\lambda_2,\varepsilon\in F$, we have $\mu,\mu^{-1}\in F$. Now, μ,μ^{-1} are roots of a quadratic polynomial over F_{2e} and therefore both belong to F_{2e} or to a quadratic extension of F_{2e}. But as f/e is odd, the field F does not contain a quadratic extension of F_{2e}. We conclude that $\mu,\mu^{-1}\in F_{2e}$.

The facts in the previous paragraphs imply that if $(A,1) \in K\backslash K_0$, then the matrix AA^σ is diagonalisable over F and C is diagonalisable over F_{2e}. In particular, there exists a $P \in GL(2,p^{2e})$ such that $P^{-1}CP = D'$ for $D' = \mathrm{dia}(\mu,\mu^{-1})$; multiplying by ε then gives $P^{-1}AA^\sigma P = D$ for $D = \mathrm{dia}(\lambda_1,\lambda_2)$. Here, either $\lambda_1,\lambda_2 \in F_0$ with $\lambda_1\lambda_2 \in N(F_0)$, or $\lambda_1,\lambda_2 \in F\backslash F_0$ and $\lambda_1^\sigma = \lambda_2$. With A, P, D and D' as above, in K we let $(B,1) = (P,0)^{-1}(A,1)(P,0) = (P^{-1}AP^\sigma,1)$. Then,

$$(BB^\sigma,0) = (P,0)^{-1}(A,1)(A,1)(P,0) = (P^{-1}AA^\sigma P,0) = (D,0) = (\varepsilon D',0)$$

and it follows that $BB^\sigma = D = \varepsilon D'$.

We now derive more details about the matrix $B = P^{-1}AP^\sigma$; recall that $P \in GL(2,p^{2e})$. Let u_1,u_2 be linearly independent (column) eigenvectors of C and AA^σ for the eigenvalues μ,μ^{-1} and λ_1,λ_2, respectively; we have $Cu_1 = \mu u_1$, $Cu_2 = \mu^{-1}u_2$, and $AA^\sigma u_i = \lambda_i u_i$ for $i \in \{1,2\}$. Taking the σ-image of the last equation and then multiplying by A from the left we obtain $AA^\sigma(Au_i^\sigma) = \lambda_i^\sigma(Au_i^\sigma)$ for $i = 1,2$. This means that the column vectors Au_i^σ are also eigenvectors of AA^σ for the eigenvalues λ_i^σ, $i = 1,2$. It follows that if $\lambda_i = \lambda_i^\sigma$ for $i = 1,2$, then we must have $Au_i^\sigma = \varepsilon_i u_i$, and if $\lambda_i = \lambda_{3-i}^\sigma$, then $Au_i^\sigma = \varepsilon_{3-i}u_{3-i}$, in both cases for some $\varepsilon_1,\varepsilon_2 \in F$. The last bit we need is the fact that for the matrix P we may take $P = (u_1,u_2)$, i.e., the matrix formed by the columns u_1,u_2, with entries in F_{2e}. Now, for $i = 1,2$, in the case $\lambda_i = \lambda_i^\sigma$ we have $AP^\sigma = (Au_1^\sigma,Au_2^\sigma) = (\varepsilon_1u_1,\varepsilon_2u_2) = P\,\mathrm{dia}(\varepsilon_1,\varepsilon_2)$, and in the case $\lambda_i = \lambda_{3-i}^\sigma$ a similar calculation gives $AP^\sigma = (Au_1^\sigma,Au_2^\sigma) = (\varepsilon_2u_2,\varepsilon_1u_1) = P\,\mathrm{off}(\varepsilon_1,\varepsilon_2)$. This shows that our matrix $B = P^{-1}AP^\sigma$ is equal to $\mathrm{dia}(\varepsilon_1,\varepsilon_2)$ or to $\mathrm{off}(\varepsilon_1,\varepsilon_2)$ for suitable $\varepsilon_1,\varepsilon_2 \in F$, depending on whether λ_1^σ is equal to λ_1 or λ_2. In both cases, of course, $\varepsilon_1\varepsilon_2 \in N(F)$.

Recalling our notation $[A,i]$ for the cosets $(A,i)L = \{(\delta A,i);\ \delta \in F^*\}$, the above calculations lead to the following result.

Proposition 10.8. *Let $G = M(p^{2f})$ for some odd prime p. Then, every element of the form $[A,1] \in G$ is conjugate in $\overline{G}$ to $[B,1]$ with $B = \mathrm{dia}(\lambda,1)$ or $B = \mathrm{off}(\lambda,1)$ for some $\lambda \in N(F)$. If, in addition, $[AA^\sigma,0] = [C,0]$ for some $C \in \mathrm{SL}(2,p^{2e})$ with f/e odd, then $[B,1] = [P,0]^{-1}[A,1][P,0]$ for some $P \in \mathrm{GL}(2,p^{2e})$, and $\lambda\lambda^\sigma \in F_{2e}$ or $\lambda/\lambda^\sigma \in F_{2e}$, depending on whether B is equal to $\mathrm{dia}(\lambda,1)$ or to $\mathrm{off}(\lambda,1)$.*

Proof. We have proven everything except for the last assertion. We have seen that if $[AA^\sigma,0] = [C,0]$ for some $C \in \mathrm{SL}(2,p^{2e})$ with f/e odd, then $BB^\sigma = P^{-1}(AA^\sigma)P = \varepsilon\,\mathrm{dia}(\mu,\mu^{-1})$ for some $\varepsilon \in F^*$, $\mu \in F_{2e}$ and some $P \in \mathrm{GL}(2,p^{2e})$. If $B = \mathrm{dia}(\lambda,1)$, then we have

$\mathrm{dia}(\lambda\lambda^\sigma, 1) = BB^\sigma = \varepsilon C = \varepsilon\,\mathrm{dia}(\mu, \mu^{-1})$, which implies that $\varepsilon = \mu$ and $\lambda\lambda^\sigma = \mu^2 \in F_{2e}$. In the case when $B = \mathrm{off}(\lambda, 1)$ we have $\mathrm{off}(\lambda, \lambda^\sigma) = BB^\sigma = \varepsilon C = \varepsilon\,\mathrm{dia}(\mu, \mu^{-1})$, from which we obtain $\lambda/\lambda^\sigma = \mu^2 \in F_{2e}$. □

Let us have another look at conjugation in the group $\overline{G} = J/L$. Observe that if $(P, i) \in J$, then $(P, i)^{-1} = ((P^{\sigma^i})^{-1}, i)$. Conjugates of $(B, 1) \in K$ by (P, i) have the form $(P, 0)^{-1}(B, 1)(P, 0) = (P^{-1}BP^\sigma, 1)$ if $i = 0$, and $(P, 1)^{-1}(B, 1)(P, 1) = ((P^\sigma)^{-1}B^\sigma P, 1)$ if $i = 1$. It follows that two elements $(B, 1)$ and $(B', 1)$ of K are conjugate in J if and only if $B' = P^{-1}BP^\sigma$ or $B' = (P^\sigma)^{-1}B^\sigma P$ for some $P \in GL(2, F)$. Taking the σ-image in the second case and passing onto $G = K/L$ we have:

Proposition 10.9. *Two elements $[B, 1]$ and $[B', 1]$ of G are conjugate in $\overline{G}$ if and only if $P^{-1}BP^\sigma = \varepsilon B'$ or $P^{-1}BP^\sigma = \varepsilon B'^\sigma$ for some $\varepsilon \in F^*$ and some $P \in GL(2, F)$.*

We will write the two conditions of Proposition 10.9 in the unified form $P^{-1}BP^\sigma = \varepsilon B'^{(\sigma)}$, or, equivalently, $BP^\sigma = \varepsilon P B'^{(\sigma)}$ where $B'^{(\sigma)}$ is equal to B' or B'^σ, depending on whether $i = 0$ or $i = 1$ when using the element $[P, i]$ for conjugation.

10.5 Conjugacy of representatives of twisted elements

We continue with identification of elements of $\overline{G}$ that conjugate a diagonal (or an off-diagonal) element from Proposition 10.8 to another such element. As a by-product we will be able to identify $\overline{G}$-stabilisers of our representatives of twisted elements in G.

We begin with the case when $B = \mathrm{dia}(\lambda, 1)$ and $B' = \mathrm{dia}(\lambda', 1)$; by Proposition 10.9 it is sufficient to find the nonsingular matrices $P \in GL(2, F)$ and $\varepsilon \in F^*$ for which $BP^\sigma = \varepsilon P B'^{(\sigma)}$ in the sense of the notation introduced at the end of Section 10.4. Throughout the computation we will use the symbols $\lambda^{(\sigma)}$ and $\lambda'^{(\sigma)}$ in an analogous way as explained for $B'^{(\sigma)}$. Assuming that P has entries $\alpha, \beta, \gamma, \delta$, the above condition says that

$$\begin{pmatrix} \lambda & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \alpha^\sigma & \beta^\sigma \\ \gamma^\sigma & \delta^\sigma \end{pmatrix} = \varepsilon \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} \begin{pmatrix} \lambda'^{(\sigma)} & 0 \\ 0 & 1 \end{pmatrix} .$$

Evaluating the products we obtain:

$$\lambda\alpha^\sigma = \varepsilon\lambda'^{(\sigma)}\alpha\ ,\quad \lambda\beta^\sigma = \varepsilon\beta\ ,\quad \gamma^\sigma = \varepsilon\lambda'^{(\sigma)}\gamma\ ,\quad \delta^\sigma = \varepsilon\delta\ .$$

By straightforward manipulation this gives the following system of four equations:

$$\alpha(\lambda\lambda^\sigma - \varepsilon\varepsilon^\sigma\lambda'\lambda'^\sigma) = 0\ ,\quad \beta(\lambda\lambda^\sigma - \varepsilon\varepsilon^\sigma) = 0\ ,\quad \gamma(\varepsilon\varepsilon^\sigma\lambda'\lambda'^\sigma - 1) = 0\ ,\quad \delta(\varepsilon\varepsilon^\sigma - 1) = 0\ .$$

From non-singularity of P it follows that $\delta \neq 0$ or $\beta \neq 0$, that is, $\varepsilon\varepsilon^\sigma = 1$ or $\varepsilon\varepsilon^\sigma = \lambda\lambda^\sigma$.

Consider first the case $\varepsilon\varepsilon^\sigma = 1$, that is, $\varepsilon^{q+1} = 1$ for $q = p^f$. Since $\lambda, \lambda' \in N(F)$, we have $\lambda\lambda^\sigma \neq 1 \neq \lambda'\lambda'^\sigma$. Our equations together with $\varepsilon\varepsilon^\sigma = 1$ then imply that $\beta = \gamma = 0$. Hence $\alpha, \delta \neq 0$, by non-singularity of P; in particular, $\lambda\lambda^\sigma = \lambda'\lambda'^\sigma$, or, equivalently, $(\lambda'/\lambda)^{q+1} = 1$. We are interested in conjugation in the group $\overline{G} = J/L$ and so we may assume that $\delta = 1$, which reduces the relations below our matrix equation to $\varepsilon = 1$ and $\alpha^{q-1} = \lambda'^{(\sigma)}/\lambda$. Since $(\lambda'/\lambda)^{q+1} = 1$, the equation $\eta^{q-1} = \lambda'/\lambda$ has $q-1$ solutions $\eta \in F^*$ (note that $|F^*| = q^2 - 1$). If $\lambda'^{(\sigma)} = \lambda'$, then all solutions of the equation $\alpha^{q-1} = \lambda'^{(\sigma)}/\lambda$ have the form $\alpha = \eta$, and if $\lambda'^{(\sigma)} = \lambda'^\sigma = \lambda'\lambda'^{q-1}$, then all solutions of this equation are $\alpha = \eta\lambda'$.

The second case to consider is $\varepsilon\varepsilon^\sigma = \lambda\lambda^\sigma$ $(\neq 1)$, which implies that $\alpha = \delta = 0$, and also $\lambda\lambda^\sigma\lambda'\lambda'^\sigma = 1$ since γ, β now must be non-zero. By the same token as above we may let $\gamma = 1$ without loss of generality. Then, our equations for γ and β in this case reduce to $\varepsilon\lambda'^{(\sigma)} = 1$ and $\lambda\beta^\sigma = \varepsilon\beta$, the latter now being equivalent to $\beta^{q-1} = 1/(\lambda\lambda'^{(\sigma)})$. Since now $(\lambda\lambda')^{q+1} = 1$, there are $q-1$ solutions ζ of the equation $\zeta^{q-1} = 1/(\lambda\lambda')$ in F^*. If $\lambda'^{(\sigma)} = \lambda'$, then we have $\beta = \zeta$, and if $\lambda'^{(\sigma)} = \lambda'^\sigma = \lambda'\lambda'^{q-1}$, we have $\beta = \zeta/\lambda'$. Summing up, we arrive at the following:

Proposition 10.10. *Let $B = \mathrm{dia}(\lambda, 1)$ and $B' = \mathrm{dia}(\lambda', 1)$ for $\lambda, \lambda' \in N(F)$. If an element $[P, i] \in \overline{G}$ conjugates $[B, 1]$ to $[B', 1]$, then, without loss of generality, $P = \mathrm{dia}(\omega, 1)$ or $P = \mathrm{off}(\omega, 1)$ for suitable $\omega \in F^*$. Moreover:*

1. *If $P = \mathrm{dia}(\omega, 1)$, then $\lambda\lambda^\sigma = \lambda'\lambda'^\sigma$, and if this condition is satisfied, then $[B, 1]$ conjugates to $[B', 1]$ in $\overline{G}$ exactly by the $q-1$ elements $[P, 0]$ such that $\omega = \eta$ and the $q-1$ elements $[P, 1]$ with $\omega = \eta\lambda'$, where $\eta \in F^*$ is one of the $q-1$ solutions of the equation $\eta^{q-1} = \lambda'/\lambda$.*

2. *If $P = \mathrm{off}(\omega, 1)$, then $\lambda\lambda^\sigma\lambda'\lambda'^\sigma = 1$, and if this holds, then $[B, 1]$ conjugates to $[B', 1]$ in $\overline{G}$ exactly by the $q-1$ elements $[P, 0]$ with $\omega = \zeta$ and the $q-1$ elements $[P, 1]$ such that $\omega = \zeta/\lambda'$, where $\zeta \in F^*$ is one of the $q-1$ solutions of the equation $\zeta^{q-1} = 1/(\lambda\lambda')$.*

We now repeat this process but now with matrices $B = \mathrm{off}(\lambda, 1)$ and $B' = \mathrm{off}(\lambda', 1)$. Conjugating by $[P, i]$ and assuming that P has entries $\alpha, \beta, \gamma, \delta$, the unified form $BP^\sigma = \varepsilon PB'^{(\sigma)}$ of the condition of Proposition 10.9 now translates into the matrix

equation

$$\begin{pmatrix} 0 & \lambda \\ 1 & 0 \end{pmatrix} \begin{pmatrix} \alpha^\sigma & \beta^\sigma \\ \gamma^\sigma & \delta^\sigma \end{pmatrix} = \varepsilon \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} \begin{pmatrix} 0 & \lambda'^{\sigma^i} \\ 1 & 0 \end{pmatrix} .$$

It follows that

$$\alpha^\sigma = \varepsilon\delta \ , \quad \beta^\sigma = \varepsilon\lambda'^{\sigma^i}\gamma \ , \quad \lambda\gamma^\sigma = \varepsilon\beta \ , \quad \lambda\delta^\sigma = \varepsilon\lambda'^{\sigma^i}\alpha \ ,$$

which, after some manipulation, yield the following two equations:

$$\gamma(\lambda^\sigma - \varepsilon\varepsilon^\sigma\lambda'^{\sigma^i}) = 0 \quad \text{and} \qquad \delta^\sigma(\lambda - \varepsilon\varepsilon^\sigma\lambda'^{\sigma^i}) = 0 \ .$$

This all means that either (a) $\lambda = \varepsilon\varepsilon^\sigma\lambda'^{\sigma^i}$, and then $\beta = \gamma = 0$ and we may assume $\delta = 1$, or else (b) $\lambda^\sigma = \varepsilon\varepsilon^\sigma\lambda'^{\sigma^i}$, and then we have $\alpha = \delta = 0$ and, without loss of generality, $\gamma = 1$. Since $\varepsilon\varepsilon^\sigma, \lambda\lambda^\sigma \in F_0^*$, these conditions are equivalent to $\lambda/\lambda' \in F_0^*$ or $\lambda\lambda' \in F_0^*$, independently of the value of i, but in our analysis below it is still useful to refer to i.

In the case (a), when $\lambda/\lambda'^{\sigma^i} = \varepsilon\varepsilon^\sigma \in F_0^*$, for every $i \in \{0,1\}$ there are $q+1$ $(q+1)^{\text{th}}$ roots $\eta_{(i)}$ of $\lambda/\lambda'^{\sigma^i}$ in F^*. From $\delta = 1$ we have $a^\sigma = \varepsilon$ and $\lambda = \varepsilon\lambda'^{\sigma^i}\alpha$, that is, $\alpha^{q+1} = \lambda/\lambda'^{\sigma^i}$. This implies that $\alpha = \eta_{(i)}$ is one of the $(q+1)^{\text{th}}$ roots of $\lambda/\lambda'^{\sigma^i}$, giving $q+1$ conjugation elements $[P,i]$ such that $P = \mathrm{dia}(\eta_{(i)}, 1)$. In the case (b), $\lambda^\sigma/\lambda'^{\sigma^i} = \varepsilon\varepsilon^\sigma \in F_0^*$ and since also $\lambda\lambda^\sigma \in F_0^*$, we have $\lambda\lambda'^{\sigma^i} \in F_0^*$. It follows that for every $i \in \{0,1\}$ there are $q+1$ $(q+1)^{\text{th}}$ roots $\zeta_{(i)}$ of $\lambda\lambda'^{\sigma^i}$ in F^*. From $\gamma = 1$ we obtain $\lambda = \varepsilon\beta$ and $\beta^\sigma = \varepsilon\lambda'^{\sigma^i}$, which means that $\beta^{q+1} = \lambda\lambda'^{\sigma^i}$. Consequently, $\beta = \zeta_{(i)}$ and we have in this second case $q+1$ conjugation elements $[P,i]$ such that $P = \mathrm{off}(\zeta_{(i)}, 1)$. Realising that the condition (a) for $i = 0$ is equivalent to (b) for $i = 1$ (and equivalent to $\lambda/\lambda' \in F_0^*$) and, similarly, the condition (a) for $i = 1$ is equivalent to (b) for $i = 0$ (and equivalent to $\lambda\lambda' \in F_0^*$), we conclude that:

Proposition 10.11. *Let $B = \mathrm{off}(\lambda, 1)$ and $B' = \mathrm{off}(\lambda', 1)$ for $\lambda, \lambda' \in N(F)$. Further, for $i \in \{0,1\}$, let $\eta_{(i)}, \zeta_{(i)} \in F^*$ be any of the $q+1$ roots of the equation $\eta_{(i)}^{q+1} = \lambda/\lambda'^{\sigma^i}$ and $\zeta_{(i)}^{q+1} = \lambda\lambda'^{\sigma^i}$, respectively. Then, an element $[P,i] \in \overline{G}$ conjugates $[B,1]$ to $[B',1]$ if and only if $\lambda/\lambda'^{\sigma^i}$ or $\lambda\lambda'^{\sigma^i}$ are elements of F_0^*. In an equivalent form, $[B,1]$ is conjugate to $[B',1]$ if and only if either*

1. *$\lambda/\lambda' \in F_0^*$, in which case the conjugation is realised by exactly $q+1$ elements $[P,0]$ with $P = \mathrm{dia}(\eta_{(0)}, 1)$ and exactly $q+1$ elements $[P,1]$ with $P = \mathrm{off}(\zeta_{(1)}, 1)$, or*

2. *$\lambda\lambda' \in F_0^*$, by the conjugation realised by exactly $q+1$ elements $[P,0]$ with $P = \mathrm{dia}(\eta_{(1)}, 1)$ and exactly $q+1$ elements $[P,1]$ with $P = \mathrm{off}(\zeta_{(0)}, 1)$.*

10.6 Conjugacy classes of twisted elements

With the help of the calculations in the previous section we can now prove a useful result about identification of suitable representatives of conjugacy classes (in the group $\overline{G}$) of elements of $G\backslash G_0$.

Theorem 10.12. *Let ξ be a primitive element of F and let $[A,1]$ be an element of G. Then, exactly one of the following two cases occur:*

1. *There exists exactly one odd $i \in \{1,2,\ldots,(q-1)/2\}$ such that $[A,1]$ is conjugate in $\overline{G}$ to $[B,1]$ with $B = \mathrm{dia}(\xi^i,1)$; the order of $[A,1]$ in G is then $2(q-1)/\gcd\{q-1,i\}$.*
2. *There exists exactly one odd $i \in \{1,2,\ldots,(q+1)/2\}$ such that $[A,1]$ is conjugate in $\overline{G}$ to $[B,1]$ with $B = \mathrm{off}(\xi^i,1)$, and the order of $[A,1]$ in G is $2(q+1)/\gcd\{q+1,i\}$.*

Furthermore, we have:

3. *The stabiliser of $[B,1]$ for $B = \mathrm{dia}(\lambda,1)$, $\lambda \in N(F)$, in the group $\overline{G}$ is isomorphic to the cyclic group $\mathbb{Z}_{2(q-1)}$ generated by (conjugation by) $[P,1]$ for $P = \mathrm{dia}(\mu\lambda,1)$ with a suitable $(q-1)^{\text{th}}$ root of unity μ, except when λ is a $(q+1)^{\text{th}}$ root of -1 and $q \equiv -1 \bmod 4$, in which case the stabiliser is isomorphic to $\mathbb{Z}_{2(q-1)} \cdot \mathbb{Z}_2$.*
4. *The stabiliser of $[B,1]$ for $B = \mathrm{off}(\lambda,1)$, $\lambda \in N(F)$, in the group $\overline{G}$ is isomorphic to the cyclic group $\mathbb{Z}_{2(q+1)}$ generated by (conjugation by) $[P,1]$ for $P = \mathrm{off}(\mu\lambda,1)$, where μ is a suitable $(q+1)^{\text{th}}$ root of unity, except when λ is a $(q-1)^{\text{th}}$ root of -1 and $q \equiv 1 \bmod 4$, when the stabiliser is isomorphic to $\mathbb{Z}_{2(q+1)} \cdot \mathbb{Z}_2$.*

Outline proof. By Proposition 10.8, each element $[A,1] \in G$ is conjugate in G to $[B,1]$ with $B = \mathrm{dia}(\lambda,1)$ or $B = \mathrm{off}(\lambda,1)$ for some $\lambda \in N(F)$. Parts 1 and 3 follow from some detailed combination of this with Proposition 10.10, and parts 2 and 4 similarly using Proposition 10.11. □

Let us remark that the exceptional cases in the items 3 and 4 above correspond precisely to elements $[B,1]$ of order 4 in G. By inspecting possible orders of $[B,1]$ we also have:

Corollary 10.13. *Every element of $G\backslash G_0$ has order divisible by 4.*

10.7 Non-singular pairs and twisted subgroups

Our aim in this and the following two sections is to determine representatives of selected conjugacy classes $\{(x,y)^g;\ g \in \overline{G} = G\langle\sigma\rangle\}$ of elements $x, y \in G$ satisfying $(xy)^2 = 1$, and make important conclusions about subgroups the corresponding pairs (x,y) generate. We note that the action of σ need not be considered separately, because $[I,1][A,i][I,1] = [A^\sigma, i]$ for $i \in \{0,1\}$, which means that the action of σ is equivalent to conjugation by the element $[I,1] \in \overline{G}$.

Since we want xy to have order 2, both x and y as above must lie in $G\backslash G_0$ because, by Corollary 10.13, there are no involutions in $G\backslash G_0$. By Theorem 10.12 we may assume that $y = [B,1]$ for $B = \mathrm{dia}(\lambda,1)$ or $B = \mathrm{off}(\lambda,1)$ for a suitable $\lambda \in N(F)$. Letting $x = [A,1]$, the pair x, y may in general generate a proper subgroup of $G = M(p^{2f})$; such cases will still be of interest for our intended classification of orientably-regular maps as long as the subgroup $\langle x,y\rangle$ is twisted, that is, isomorphic to $M(p^{2e})$ for a suitable divisor e of f.

We now identify conditions on A implied by the requirement that $([A,1][B,1])^2$ be the identity in G and begin with the case when $B = \mathrm{dia}(\lambda,1)$. Let $A \in GL(2,F)$ have rows (a,b) and (c,d), with determinant $ad - bc \in N(F)$. Then, $[A,1][B,1] = [AB^\sigma, 0]$, where

$$AB^\sigma = \begin{pmatrix} a\lambda^\sigma & b \\ c\lambda^\sigma & d \end{pmatrix} .$$

Since AB^σ lies in $PSL(2,F)$, it has order 2 if and only if its trace $a\lambda^\sigma + d$ is equal to zero. If one of a, d was equal to zero, both would have to be zero and then $[A,1]$ and $[B,1]$ would clearly not generate a twisted subgroup of G. Therefore both a, d are non-zero and we may assume without loss of generality that $a = -1$ and $d = \lambda^\sigma$. We will thus consider only elements $[A,1] \in G$ of the form

$$A = \begin{pmatrix} -1 & b \\ c & \lambda^\sigma \end{pmatrix} , \quad u = bc + \lambda^\sigma \in N(F) . \tag{10.1}$$

Next, consider the case when $B = \mathrm{off}(\lambda,1)$. For a matrix $A \in GL(2,F)$ with rows (a,b) and (c,d) such that $ad - bc \in N(F)$ we now have $[A,1][B,1] = [AB^\sigma, 0]$, where

$$AB^\sigma = \begin{pmatrix} b & a\lambda^\sigma \\ d & c\lambda^\sigma \end{pmatrix} .$$

Again, $AB^\sigma \in PSL(2,F)$ has order 2 if and only if its trace $b + c\lambda^\sigma$ is equal to zero. If one of b, c was equal to zero, we would have $b = c = 0$, but then $[A, 1]$ and $[B, 1]$ would again not generate a twisted subgroup of G. Therefore both b and c are non-zero and we may assume that $c = -1$ and $b = \lambda^\sigma$. It follows that, without loss of generality, we only need to consider elements $[A, 1] \in G$ such that

$$A = \begin{pmatrix} a & \lambda^\sigma \\ -1 & d \end{pmatrix} \ , \quad u = ad + \lambda^\sigma \in N(F) \ . \tag{10.2}$$

With A and B as above we can still identify obvious instances when $[A, 1]$ and $[B, 1]$ do not generate a twisted subgroup of G. This is certainly the case if

(i) both $[A, 1]$ and $[B, 1]$ have order 4, as then the two elements generate a solvable group, cf. [24], or

(ii) $B = \mathrm{dia}(\lambda, 1)$ and A is an upper- or a lower-triangular matrix, as then $[A, 1]$ and $[B, 1]$) generate a triangular subgroup of G, or else

(iii) $B = \mathrm{off}(\lambda, 1)$ and A is an off-diagonal matrix, as then $[A, 1]$ and $[B, 1]$ clearly do not generate a twisted subgroup of G.

For $B = \mathrm{dia}(\lambda, 1)$ and A given by (10.1) and for $B = \mathrm{off}(\lambda, 1)$ and A given by (10.2), an ordered pair $([A, 1], [B, 1])$ not satisfying any of (i), (ii) and (iii) will be called *non-singular*.

We are now in position to classify the subgroups of $G = M(p^{2f})$ generated by non-singular pairs. To do so we will again use knowledge of the situation in the subgroup $G_0 \simeq PSL(2, p^{2f})$ of G. Recall that a subgroup H of $G = M(p^{2f})$ was said to stabilise a point if there exists an element in $F \cup \{\infty\}$ fixed by all linear fractional mappings corresponding to elements of H; also, G_0 denotes the (unique) copy of $PSL(2, p^{2f})$ in G.

Proposition 10.14. *Let H be a subgroup of G generated by a non-singular pair $([A, 1], [B, 1])$. Then H is isomorphic to $M(p^{2e})$ for some positive divisor e of f such that f/e is odd.*

Outline proof. Let $H_0 = H \cap G_0$. The classification of [48] tells us that subgroups of G_0 fall into the following categories: point stabilisers, dihedral groups, A_4, S_4, A_5, and $\mathrm{PSL}(2, p^g)$ for divisors g of $2f$ together with $PGL(2, p^{g/2})$ for even divisors g of $2f$. For our subgroup H_0 we rule out all but the penultimate case.

We then apply Proposition 10.7 to the subgroup H_0 to conclude that H is conjugate in $\overline{G}$ to a subgroup isomorphic to $M(p^{2e})$ for some positive divisor e of f such that f/e is odd, completing the proof. □

It follows that a pair $([A,1],[B,1])$ of elements of G as above generates a twisted subgroup of G if and only if the pair is non-singular.

10.8 Orbits of non-singular pairs: The diagonal case

We will identify representatives of $\overline{G}$-orbits of non-singular pairs $([A,1],[B,1])$, dealing with $B = \mathrm{dia}(\lambda,1)$ and A as in (10.1) here and deferring the case $B = \mathrm{off}(\lambda,1)$ to the next section.

Instead of working with matrices, the form of A in (10.1) suggests to look at the corresponding quadruples (λ,b,c,u), also called *non-singular*, under the induced action of the stabiliser of $[B,1]$ in $\overline{G}$. We recall that the values of λ and identification of the stabiliser are in items 1 and 3 of Theorem 10.12. To simplify the notation in what follows, for any $\omega \in F$ we will use the symbol $\sqrt[r]{\omega}$ to denote the set of all r^{th} roots of ω in $F = GF(q^2)$, $q = p^f$. The analysis in the third part of the proof of Theorem 10.12 tells us that the stabiliser of $[B,1]$ in $\overline{G}$ consists exactly of the following elements of $\overline{G}$:

$[P_1(\eta),0]$, where $P_1 = \mathrm{dia}(\eta,1)$ and $\eta \in F_0^*$;
$[P_2(\eta),1]$, where $P_2 = \mathrm{dia}(\eta\lambda,1)$ and $\eta \in F_0^*$;
$[P_3(\zeta),0]$, where $P_3 = \mathrm{off}(\zeta,1)$ if $\lambda \in \sqrt[q+1]{-1}$ and $\zeta \in \sqrt[q-1]{\lambda^{-2}}$;
$[P_4(\zeta),1]$, where $P_4 = \mathrm{off}(\zeta/\lambda,1)$ if $\lambda \in \sqrt[q+1]{-1}$ and $\zeta \in \sqrt[q-1]{\lambda^{-2}}$.

To find the corresponding orbit of $[A,1]$ we first evaluate the products $[P_j(\eta),0]^{-1}[A,1][P_j(\eta),0]$ for A as in (10.1) and $j \in \{1,2,3,4\}$:

$$[P_1(\eta),0]^{-1}[A,1][P_1(\eta),0] = [C_1,1], \quad \text{where} \quad C_1 = \begin{pmatrix} -1 & b\eta^{-1} \\ c\eta & \lambda^\sigma \end{pmatrix};$$

$$[P_2(\eta),1]^{-1}[A,1][P_2(\eta),1] = [C_2,1], \quad \text{where} \quad C_2 = \begin{pmatrix} -1, & b^\sigma(\eta\lambda)^{-1} \\ c^\sigma\eta\lambda^\sigma & \lambda^\sigma \end{pmatrix};$$

$$[P_3(\zeta),0]^{-1}[A,1][P_3(\zeta),0] = [C_3,1], \quad \text{where} \quad C_3 = \begin{pmatrix} -1 & c\zeta/\lambda \\ b\lambda/\zeta & \lambda^\sigma \end{pmatrix}; \quad \text{and}$$

$$[P_4(\zeta),1]^{-1}[A,1][P_4(\zeta),1] = [C_4,1], \quad \text{where} \quad C_4 = \begin{pmatrix} -1 & -c^\sigma\zeta/\lambda^2 \\ b^\sigma/\zeta & \lambda^\sigma \end{pmatrix}.$$

Let $\lambda = \xi^i$ for a fixed primitive element $\xi \in F$ and some odd i such that

$1 \le i \le (q-1)/2$; note that here $\lambda \in \sqrt[q+1]{-1}$ if and only if $i = (q-1)/2$. It follows that we have either $(q-1)/4$ such odd values of i if $q \equiv 1 \bmod 4$ and all are smaller than $(q-1)/2$, or else $(q-3)/4$ such odd $i < (q-1)/2$ together with $i = (q-1)/2$ if $q \equiv -1 \bmod 4$.

The strategy now is to count the number of $\overline{G}$-orbits for each such i, by bringing together the information found so far. For each such $i < (n-1)/2$ we obtain, after some manipulation:

$$n_1 = (q+1)\left\lfloor \frac{q-1}{4} \right\rfloor \frac{q^2-3}{4} \ . \tag{10.3}$$

If $q \equiv -1 \pmod 4$, then we also need to consider $i = (q-1)/2$ and in that case the count turns out to be:

$$\begin{aligned} n_2 &= \frac{(q^2-1)(q-1)^2/4 + ((q^2-3)/2 - (q^2-1)/4)(q^2-1)}{4(q-1)} \\ &= \frac{1}{8}\left((q+1)(q^2-3) - (q^2-1)\right) \ . \end{aligned} \tag{10.4}$$

10.9 Orbits of non-singular pairs: The off-diagonal case

The counting of orbits in the off-diagonal case proceeds in an analogous manner. In this case the result for $i < (q+1)/2$ is:

$$n_3 = (q-1)\left\lfloor \frac{q+1}{4} \right\rfloor \frac{q^2+1}{4} \ . \tag{10.5}$$

If $q \equiv 1 \pmod 4$ then we must also consider $i = (q+1)/2$ and in that case the count turns out to be:

$$n_4 = \frac{(q^2-1)(q+1)(q-3)/4 + ((q^2-3)/2 - (q^2-1)/4)(q^2-1) + 2(q^2-1)}{4(q+1)} \ ,$$

which simplifies to

$$n_4 = \frac{1}{8}\left((q-1)(q^2+1) - (q^2-1)\right) \ . \tag{10.6}$$

10.10 Enumeration of orientably-regular maps on $M(q^2)$

We have seen in Section 10.7 that a pair $([A,1],[B,1])$ of elements of G, with diagonal B and A given by (10.1) or with off-diagonal B and A given by (10.2), and with

product of order two, generates a twisted subgroup of G if and only if the pair is non-singular. In the previous two sections we have counted orbits of non-singular pairs in G under conjugation in $\overline{G}$, with no regard to subgroups the pairs generate. The number of these orbits turns out to be $n_1 + n_3 + n_4$ if $q \equiv 1 \bmod 4$ and $n_1 + n_2 + n_3$ if $q \equiv -1 \bmod 4$; in both cases the sum is equal to $(q^2-1)(q^2-2)/8$. We state this as a separate result.

Proposition 10.15. *The number of $\overline{G}$-orbits of non-singular pairs in $G = M(q^2)$ is equal to $(q^2-1)(q^2-2)/8$.* □

We will now refine our considerations and take into account subgroups generated by non-singular pairs. For our group $G = G_{2f} = M(p^{2f})$ and for any positive divisor e of f such that f/e is odd we let G_{2e} denote the canonical copy of $M(p^{2e})$ in G. In Lemma 10.4 we saw that the automorphism σ of $F = F_{2f} = GF(p^{2f})$ of order two restricts to an automorphism σ_{2e} of order two of the subfield $F_{2e} = GF(p^{2e})$ of F. We recall that $\overline{G} = \overline{G}_{2f} = G_{2f}\langle\sigma\rangle$ and we similarly introduce $\overline{G}_{2e}$ for every e as above by letting $\overline{G}_{2e} = G_{2e}\langle\sigma_{2e}\rangle$.

Let $\mathrm{orb}_f(e)$ denote the number of $\overline{G}_{2f}$-orbits of non-singular pairs $([A,1],[B,1])$ of G that generate a subgroup of G isomorphic to $M(p^{2e})$. At the same time, let $\mathrm{orb}(e)$ be the number of orbits of non-singular pairs of G_{2e} which generate G_{2e}. The two quantities, are, in fact, equal, which is fundamental for our enumeration.

Proposition 10.16. *For each positive divisor e of f with f/e odd, we have* $\mathrm{orb}_f(e) = \mathrm{orb}(e)$.

Outline proof. It is clear that every $\overline{G}_{2e}$-orbit of a non-singular pair in the canonical copy $G_{2e} \cong M(p^{2e})$ in G_{2f} is contained in a $\overline{G}_{2f}$-orbit of the same pair. In the reverse direction, let a non-singular pair in G generate a subgroup isomorphic to $M(p^{2e})$. Since, by the important Proposition 10.6, all such subgroups are $\overline{G}_{2f}$-conjugate in G_{2f}, we may assume that the non-singular pair is contained in G_{2e}. But then the $\overline{G}_{2f}$-orbit of this pair obviously contains a $\overline{G}_{2e}$-orbit of the same pair. The proof now proceeds by establishing the following fact:

Let $([A,1],[B,1])$ and $([A',1],B',1])$ be two non-singular pairs of G_{2e} both generating G_{2e} and lying in the same $\overline{G}_{2f}$-orbit of G_{2f}. Then the two pairs are contained in the same $\overline{G}_{2e}$-orbit of G_{2e}. □

For positive integers x let us define a function h by $h(x) = (p^{2x}-1)(p^{2x}-2)/8$. In terms of h and the numbers $\mathrm{orb}_f(e)$, Proposition 10.15 simply says that

$\sum_e \mathrm{orb}_f(e) = h(f)$, where summation is taken over all positive divisors e of f such that f/e is odd. By Proposition 10.16 we may replace $\mathrm{orb}_f(e)$ with $\mathrm{orb}(e)$ and obtain $\sum_e \mathrm{orb}(e) = h(f)$, with the same summation convention. This miniature but important detail enables us to make a substantial advance in the enumeration.

Let $f = 2^\alpha o$ where o is an odd integer and let e be a divisor of f such that f/e is odd; equivalently, $e = 2^\alpha d$ where d is a positive (and necessarily odd) divisor of o. Taking the above notes into account, Proposition 10.15 may then be restated as follows:

$$\sum_{d|o} \mathrm{orb}(2^\alpha d) = h(2^\alpha o) \ . \tag{10.7}$$

Using the Möbius inversion we obtain $\mathrm{orb}(f) = \mathrm{orb}(2^\alpha o) = \sum_{d|o} \mu(o/d) h(2^\alpha d)$, where μ is the classical number-theoretic Möbius function μ on positive integers. We thus arrive at our first main result.

Theorem 10.17. *Let $q = p^f$ for an odd prime p, let $G = M(q^2)$, and let $f = 2^\alpha o$ with o odd. The number of $\overline{G}$-orbits of non-singular generating pairs of G is equal to*

$$\sum_{d|o} \mu(o/d) h(2^\alpha d) \ , \quad \text{where} \quad h(x) = (p^{2x} - 1)(p^{2x} - 2)/8 \ .$$

The last step is to study conjugacy of non-singular pairs of $M(q^2)$ under the action of the group $\mathrm{Aut}(M(q^2))$ which, as we know by Proposition 10.3, is isomorphic to $P\Gamma L(2, q^2)$. Since for $q = p^f$ we have $P\Gamma L(2, q^2) \cong PGL(2, q^2) \rtimes \mathbb{Z}_{2f} \cong \overline{G} \rtimes \mathbb{Z}_f$, it is sufficient to investigate the induced action of the Galois automorphisms $\sigma_j : \ z \mapsto z^{p^j}$ for $z \in F = GF(p^{2f})$ and $1 \le j \le f - 1$ on the $\overline{G}$-orbits of our non-singular pairs $([A, 1], [B, 1])$. We will use the natural notation O^{σ_j} for the σ_j-image of a $\overline{G}$-orbit O of a pair $([A, 1], [B, 1])$ of elements of G. Note that $\sigma_f = \sigma$, and we also have $O^{\sigma_f} = O$, by the remark made at the beginning of Section 10.7. Clearly, if $O^{\sigma_j} \cap O \neq \emptyset$, then $O^{\sigma_j} = O$.

Proposition 10.18. *Let O be the orbit of a non-singular pair $([A, 1], [B, 1])$ of elements of G under conjugation in $\overline{G}$ and let j be the smallest positive integer for which $O^{\sigma_j} = O$. If $[A, 1]$ and $[B, 1]$ generate G, then $j = f$.*

Proof (Hriňáková, Širáň). We may assume that $f \ge 2$, otherwise the result is trivial. Suppose that the pair $([A, 1], [B, 1])^{\sigma_j} = ([A, 1]^{\sigma_j}, [B, 1]^{\sigma_j})$ is $\overline{G}$-conjugate to the pair $([A, 1], [B, 1])$, that is, there exists some $C \in GL(2, q^2)$ and $i \in Z_2$ such that $[A, 1]^{\sigma_j} = [C, i]^{-1}[A, 1][C, i]$ and $[B, 1]^{\sigma_j} = [C, i]^{-1}[B, 1][C, i]$. It follows that for every $[X, 1] \in \langle [A, 1], [B, 1] \rangle$ we have $[X, 1]^{\sigma_j} = [C, i]^{-1}[X, 1][C, i]$. Using our assumption

that $\langle [A,1],[B,1]\rangle = G$, we conclude that the above is valid also for $X = \mathrm{dia}(\xi,1)$, where ξ is a primitive element of $F = GF(p^{2f})$. Letting C have elements $\alpha,\beta,\gamma,\delta$ in the usual order, the equivalent form $[C,i][\mathrm{dia}(\xi^{p^j},1),1] = [\mathrm{dia}(\xi,1),1][C,i]$ of the above equation yields

$$\begin{pmatrix}\alpha & \beta\\ \gamma & \delta\end{pmatrix}\begin{pmatrix}(\xi^{p^j})^{(\sigma)} & 0\\ 0 & 1\end{pmatrix} = \varepsilon\begin{pmatrix}\xi & 0\\ 0 & 1\end{pmatrix}\begin{pmatrix}\alpha^\sigma & \beta^\sigma\\ \gamma^\sigma & \delta^\sigma\end{pmatrix}$$

for some $\varepsilon \in F^*$; here we used the (σ)-convention introduced at the end of section 10.4. This gives the system of equations

$$\alpha(\xi^{p^j})^{(\sigma)} = \varepsilon\xi\alpha^\sigma \ , \quad \beta = \varepsilon\xi\beta^\sigma \ , \quad \gamma(\xi^{p^j})^{(\sigma)} = \varepsilon\gamma^\sigma \ , \delta = \varepsilon\delta^\sigma \ .$$

Consider first the case $\delta \neq 0$; without loss of generality we then may assume $\delta = 1$. Then $\varepsilon = 1$, and the equation for α gives $(\xi^{p^j})^{(\sigma)}\xi^{-1} = \alpha^\sigma\alpha^{-1}$, or, equivalently, $\xi^{p^{j+if}-1} = \alpha^{p^f-1}$. It follows that $p^f - 1$ is a divisor of $p^{j+if} - 1$, which implies that f divides $j + if$ and hence f divides j, which, since $j \leq f$, shows that $j = f$. If $\delta = 0$ then, without loss of generality, $\beta = 1$ and so $\varepsilon = \xi^{-1}$. The equation for γ now implies $\xi^{p^{j+if}+1} = \gamma^{p^f-1}$. It follows that $p^f - 1$ divides $p^{j+if} + 1$ and hence also $p^{2(j+if)} - 1$ and so f must divide $2(j + if)$. Thus, f is a divisor of $2j$ and as $j \leq f$, we have either $j = f$ or $j = f/2$ (assuming f is even). But the last case is easily seen to be impossible since $p^f - 1$ is not a divisor of $p^{f/2} + 1$ or $p^{3f/2} + 1$. This completes the proof. □

Proposition 10.18 tells us that if a non-singular pair $([A,1],[B,1])$ of elements of G actually generates G and gives rise to an orbit O under conjugation in $\overline{G}$, then the action of the group $\mathrm{Aut}(M(q^2))$ fuses the f orbits O^{σ_j} for $j \in \{0,1,\ldots,f-1\}$ into a single orbit. Recalling the one-to-one correspondence between isomorphism classes of orientably-regular maps supported by the group $G = M(q^2)$ and orbits of (necessarily non-singular) generating pairs of G under conjugation by $\mathrm{Aut}(G)$, Theorem 10.17 then immediately implies our second main result.

Theorem 10.19. *Let $q = p^f$ for an odd prime p and let $f = 2^\alpha o$ with o odd. The number of orbits of non-singular generating pairs of $M(q^2)$ under the action of the group $\mathrm{Aut}(M(q^2))$, and hence the number of isomorphism classes of orientably-regular maps $\mathcal{M}$ with $\mathrm{Aut}^+(\mathcal{M}) \cong M(q^2)$, is equal to*

$$\frac{1}{f}\sum_{d|o}\mu(o/d)h(2^\alpha d) \ , \quad \text{where} \quad h(x) = (p^{2x}-1)(p^{2x}-2)/8 \ .$$

10.11 Enumeration of reflexible maps

Recall that a map is called *reflexible* if it admits an automorphism reversing the orientation of the surface. For orientably-regular maps represented by triples (G, x, y) as indicated in Section 10.1, reflexibility is equivalent with the existence of an automorphism θ of the group G such that $\theta(x) = x^{-1}$ and $\theta(y) = y^{-1}$. Note that if such a θ exists, then $\theta^2 = id$.

In the specific situation considered in this chapter, namely, when $G = M(q^2)$ for $q = p^f$, we established in Proposition 10.3 that $\mathrm{Aut}(G) \cong P\Gamma L(2, q^2)$. Moreover, it is well known that every automorphism in $P\Gamma L(2, q^2) \cong PGL(2, q^2) \rtimes \mathbb{Z}_{2f}$ is a composition of a conjugation by some element of $PGL(2, q^2)$ and a power of the Frobenius automorphism $z \mapsto z^p$ of the Galois field $F = GF(p^{2f})$. It follows that an *involutory* automorphism θ of $G = M(q^2)$ is a composition of a conjugation as above with σ^i for $i \in \{0, 1\}$, where σ is the automorphism of F sending z to z^q. By the remark at the beginning of Section 10.7, however, the action of σ is equivalent to conjugation in $\overline{G} = G\langle\sigma\rangle$ by the element $[I, 1]$. Consequently, an orientably-regular map on the group $G = M(q^2)$ generated by a pair of elements $x = [A, 1]$ and $y = [B, 1]$ is reflexible if and only if the ordered pairs (x, y) and (x^{-1}, y^{-1}) are conjugate by an *involutory* element of $\overline{G}$.

In this section we will count the number of reflexible orientably-regular maps on $M(q^2)$. In particular, we will see that not all orientably-regular maps with automorphism group $M(q^2)$ are reflexible, in contrast with the position for $PGL(2, q^2)$, see e.g. [23]. We will use techniques similar to the main enumeration in previous sections and structure our explanations accordingly.

10.11.1 Conjugating involutions

By writing out the conjugating equations, it follows after some tedious but elementary algebraic manipulation that the possible conjugating involutions as described above can take only certain well-defined forms.

In the case $B = \mathrm{dia}(\lambda, 1)$ the possible conjugating elements can have the following forms.

$$[C, 0]; \quad C = \begin{pmatrix} 0 & \beta \\ 1 & 0 \end{pmatrix}; \quad \beta^{q-1} = \lambda^{q-1} \tag{10.8}$$

$$[C,1]; \quad C = \begin{pmatrix} 0 & \beta \\ 1 & 0 \end{pmatrix}; \quad \beta \in F'^{*} \tag{10.9}$$

In the case $B = \mathrm{off}(\lambda, 1)$ the possible conjugating elements can have the following forms.

$$[C,0]; \quad C = \begin{pmatrix} 0 & \beta \\ 1 & 0 \end{pmatrix}; \quad \beta^{q+1} = \lambda^{q+1} \tag{10.10}$$

$$[C,1]; \quad C = \begin{pmatrix} 1 & 0 \\ 0 & \delta \end{pmatrix}; \quad \delta \in \sqrt[q+1]{1} \tag{10.11}$$

10.11.2 Enumeration

We now proceed to the actual enumeration and follow the same strategy as we used for general maps, namely, counting orbits in the diagonal and off-diagonal cases for B as in Sections 10.8 and 10.9 and then deriving the final enumeration result using the Möbius inversion formula as in Section 10.10.

10.11.2.1 Counting orbits: The case $B = \mathrm{dia}(\lambda, 1)$

The first possibility is that the conjugating element has the form $[C, 0]$ with $C = \begin{pmatrix} 0 & \beta \\ 1 & 0 \end{pmatrix}$ for some β such that $\beta^{q-1} = \lambda^{q-1}$. Following some detailed computations, the number of orbits, summed over all $i \leq (q-1)/2$, turns out to be:

$$r_1 = \frac{(q^2-1)(q-2)}{8} \tag{10.12}$$

The second possibility is that the conjugating element has the form $[C, 1]$ with $C = \begin{pmatrix} 0 & \beta \\ 1 & 0 \end{pmatrix}$ for some $\beta \in F'^{*}$. The number of orbits in this case is:

$$r_2 = \frac{(q^2-1)(q-1)}{16} \tag{10.13}$$

10.11.2.2 Counting orbits: The case $B = \mathrm{off}(\lambda, 1)$

The first possibility is that the conjugating element has the form $[C, 0]$ with $C = \begin{pmatrix} 0 & \beta \\ 1 & 0 \end{pmatrix}$ for some β such that $\beta^{q+1} = \lambda^{q+1}$. The number of orbits, summed over

all $i \leq (q+1)/2$, is:

$$r_3 = \frac{q(q^2-1)}{8} \tag{10.14}$$

The remaining possibility is that the conjugating element has the form $[C, 1]$ with $C = \begin{pmatrix} 1 & 0 \\ 0 & \delta \end{pmatrix}$ for some δ such that $\delta^{q+1} = 1$. In this case the number of orbits is:

$$r_4 = \frac{(q+1)(q^2-1)}{16} \tag{10.15}$$

10.11.2.3 Summary of counting orbits

For $B = \mathrm{dia}(\lambda, 1)$ the total number of $\overline{G}$-orbits of generating pairs for reflexible maps is

$$R_1 = r_1 + r_2 = \frac{(q^2-1)(3q-5)}{16} \tag{10.16}$$

For $B = \mathrm{off}(\lambda, 1)$ the total number of such orbits is

$$R_2 = r_3 + r_4 = \frac{(q^2-1)(3q+1)}{16} \tag{10.17}$$

The total number of orbits is therefore:

$$R = R_1 + R_2 = \frac{(q^2-1)(3q-2)}{8} \tag{10.18}$$

10.11.3 Counting reflexible maps

We may enumerate the orientably-regular reflexible maps on $M(q^2)$ by using the above calculations in conjunction with the logic of Section 10.10. Since details of this process are exactly as in Section 10.10 except for using the input on counting orbits from Subsection 10.11.2.3 we present just the final result.

Theorem 10.20. *Let $q = p^f$ be an odd prime power, with $f = 2^\alpha o$ where o is odd. The number of orientably-regular reflexible maps $\mathcal{M}$ with $\mathrm{Aut}^+(\mathcal{M}) \cong M(q^2)$ is, up to isomorphism, equal to*

$$\frac{1}{f} \sum_{d|o} \mu(o/d) \tilde{h}(2^\alpha d) \ ,$$

where $\tilde{h}(x) = (p^{2x}-1)(3p^x-2)/8$ and μ is the Möbius function.

10.12 Remarks

As stated in Section 10.1, orientably-regular maps have been enumerated for a few classes of non-trivial groups, including the linear fractional groups $PSL(2,q)$ and $PGL(2,q)$ [23, 64]. It should be noted, however, that the available results for $PSL(2,q)$ and $PGL(2,q)$ are more detailed by giving 'closed formulae' for the number of orientably-regular maps of every given type, whereas our main results in Theorems 10.19 and 10.20 contain formulae for the total number of such maps.

In order to obtain a refined version of our enumeration of orientably-regular maps with automorphism group isomorphic to a twisted linear fractional group $G = M(q^2)$ one could follow [39], which requires setting up both a character table for G and the Möbius function for the lattice of subgroups of G. The number of orientably-regular maps on the group G is then obtained as a combination of a character-theoretic formula for counting solutions of the equation $xyz = 1$ for x, y, z in given conjugacy classes of G (a special case of a general formula of Frobenius [33]) combined with Möbius inversion, which is a forthcoming project of the authors. Whether the project will return a 'nice' formula, however, is not clear due to another significant difference between the family of orientably-regular maps on $M(q^2)$ compared to those on $PGL(2,q)$. Namely, in the case of $PGL(2,q)$, for any even $k, \ell \geq 4$ not both equal to 4 there is an orientably-regular map for infinitely many values of q, cf. [23]. Our next result shows that this fails to hold in the case of $M(q^2)$.

Proposition 10.21. *If $k, \ell \equiv 0 \pmod 8$ and $k \not\equiv \ell \pmod{16}$ then there is no orientably-regular map of type (k, ℓ) on $M(q^2)$ for any q.*

Proof. By Theorem 10.12, orders of elements in $G \backslash G_0$ are $o_i = 2(q-1)/\gcd\{q-1, i\}$ and $o'_i = 2(q+1)/\gcd\{q+1, i\}$ for odd i such that $1 \leq i \leq (q-1)/2$ and $1 \leq i \leq (q+1)/2$, respectively. Note that if $o_i \equiv 0 \mod 8$ then $o'_i \equiv 4 \mod 8$ and vice versa. Further, if $o_i \equiv 8 \mod 16$ then $q - 1 \equiv 4 \mod 8$ since i is odd, and if $o_i \equiv 0 \mod 16$ then $q - 1 \equiv 0 \mod 8$. It follows that for a given q we cannot have a non-singular generating pair of orders $o_i \equiv 0 \mod 16$ and $o_j \equiv 8 \mod 16$. The argument for orders of the form o'_i is similar. □

Besides reflexibility, another frequently studied property of orientably-regular maps is self-duality. In general, an oriented map is *positively self-dual* if it is isomorphic to its dual with the same orientation, and *negatively self-dual* if it is isomorphic to its oppositely oriented dual map. In terms of orientably-regular maps represented by triples (G, x, y), positive and negative self-duality is equivalent to the existence of an

	$B = \mathrm{dia}(\lambda, 1)$				$B = \mathrm{off}(\lambda, 1)$			
q	$\#(k{=}\ell)$	+ self-dual	− self-dual	both	$\#(k{=}\ell)$	+ self-dual	− self-dual	both
3	0	0	0	0	3	3	3	3
5	15	15	5	5	10	10	6	6
7	28	28	8	8	78	42	14	14
9	95	45	9	9	68	36	10	10
11	276	132	24	24	265	165	33	33
13	469	273	39	39	666	234	42	42
17	2556	612	68	68	1312	544	72	72
19	1960	760	80	80	2799	855	95	95

Table 10.1: Numbers of self-dual maps on $M(q^2)$

(involutory) automorphism of G sending the ordered pair (x, y) onto (y, x) and (y^{-1}, x^{-1}), respectively. By the same arguments as in the second paragraph of Section 10.11 one concludes that for our group $G = M(q^2)$, an orientably-regular map defined by a generating pair $[A, 1], [B, 1]$ will be positively self-dual if and only if there exists an involution $[C, i] \in \overline{G}$ conjugating the two generators, and the map will be negatively self-dual if there is such an involution conjugating $[A, 1]$ to $[B, 1]^{-1}$.

Setting up the corresponding matrix equations for such conjugations, however, lead to enormously complicated formulae from which we were not able to extract 'nice' closed formulae. Clearly in a self-dual map we have $k = \ell$ so that the orders of $[A, 1]$ and $[B, 1]$ must be equal. We used GAP[35] to construct all the regular maps on $M(q^2)$, for small values of q, with the generators of equal order, and then tested for self-duality by determining if a conjugating element $[C, i]$ as above exists. The results of this computation are given in Table 10.1, showing the numbers $\#(k{=}\ell)$ of maps that have generators of equal orders, those which are positively or negatively self-dual, and those which are both.

Note that the computational evidence suggests that a negatively self-dual map on $M(q^2)$ is also positively self-dual.

Chapter 11

Conclusion

Our final chapter begins by summarising the asymptotic results in the degree-diameter problem, highlighting those areas where we have been able to improve the entries in the table. We then conclude with some discussion and ideas for future research.

11.1 Revised table of asymptotic results

Chapters 3, 4, 5 and 8 all presented new asymptotic results in the undirected version of the degree-diameter problem. We present here an updated and expanded version of Table 2.1 summarising the new position. New results are highlighted in blue with a reference to the corresponding result from earlier chapters.

It is worth noting that our focus on highly symmetric graphs, in particular Cayley graphs, results in most progress being made on those sections of the table. Typically, apart from Cayley graphs of abelian groups, dihedral groups and the like, no better upper bound than 1 is known for the value of $L^+(k)$. Thus there are still many areas where the gap between the upper and lower asymptotic bounds is large.

Type		Diam 2	Diam 3	Diam 4	Diam 5	Diam k
General graphs						
All graphs	L^-	1.00000	0.29629	0.18750	0.08192	$1/2^k$
	L^+	1.00000	1.00000	0.25000	1.00000	$1/1.6^k$
Vertex-transitive	L^-	0.68762[a]	0.25000[b]	0.09600[c]	0.05859[c]	$k/3^k$
	L^+	1.00000	1.00000	0.09600[c]	0.05859[c]	$1/2^k$
Arc-transitive	L^-	—	—	—	—	—
	L^+	0.25000[d]	0.03703[d]	0.00390[d]	0.00032[d]	$1/k^k$[d]
Cayley graphs						
All groups	L^-	0.68762[a]	0.25000[b]	0.09600[c]	0.05859[c]	$k/3^k$
	L^+	1.00000	1.00000	0.09600[c]	0.05859[c]	$k/3^k$
Circulant	L^-	0.36111[e]	0.05600[f]	0.00815[m]	0.00081[f]	$1.20431^k/k^k$[g]
	L^+	0.36111	0.05700[f]	0.00815[m]	0.00081[f]	$1.20431^k/k^k$[g]
General abelian	L^-	0.39062	0.07031	0.00815[m]	0.00081[f]	$1.20431^k/k^k$[g]
	L^+	0.44444	0.07031	0.00815[m]	0.00081[f]	$1.20431^k/k^k$[g]
Elementary abelian 2-groups	L^-	—	—	—	—	—
	L^+	0.28444[h]	0.06250[l]	—	—	—
Dihedral	L^-	0.50000†[i]	0.31059‡[j]	0.10983‡[k]	—	—
	L^+	0.50000†[i]	0.31059‡[j]	0.10983‡[k]	—	—

a Theorem 5.8 b Theorem 5.4 c Theorem 5.6 d Theorem 9.1 e Theorem 3.7
f Theorem 3.11 g Theorem 3.20 h Theorem 8.24 i Theorem 4.5 j Theorem 4.8
k Theorem 4.9 l Theorem 8.27 m Corollary 3.17 † exact value ‡ upper bound

Table 11.1: Revised asymptotic lower bounds on orders of undirected graphs

11.2 Concluding remarks and future research

We have seen that graphs with a high degree of symmetry are a fruitful area of research in the degree-diameter problem, as well as in related areas such as the girth problem and the group-theoretical study of product-free sets. We conclude with some final remarks and possible future research, structured according to the topics of the preceding chapters.

In Chapter 3 we proved new asymptotic lower bounds on the orders of circulant graphs of diameters greater than 2. The diameter 2 case seems to be resistant to complete understanding, and the current best bound of 13/36 is some way away from the theoretical limit of 1/2. A future research project could usefully try to extend the ideas of this chapter to general abelian groups, for diameters 3 and above.

In Chapter 4 we completely settled the asymptotic position for diameter 2 Cayley graphs of dihedral groups. However, it would be interesting to attempt a construction for, say, dicyclic groups where we proved an upper bound of 1/2, to try to arrive at a similar closed position. For diameters 3 and 4 we have an upper asymptotic bound, but our diameter 2 construction cannot readily be extended to larger diameters to provide a lower bound. Given our progress in Chapter 3 on cyclic groups at diameters 3 and 4, it would be interesting to try to extend similar techniques to the dihedral case.

Chapter 5 introduced a new generalised construction providing new asymptotic bounds for Cayley graphs of diameters 3 and above. This relies heavily on computer search for feasible solutions, and it would be preferable to prove some general results about existence of solutions. A current project is investigating the existence of solutions for Cayley graphs of groups of the form $H^k \rtimes K$, in the case where the right hand side group K is either cyclic or dihedral.

In the mixed graph problem, Chapter 6 introduced a correction to the published formula for the Moore bound, and ruled out the possibility of existence of mixed Cayley Moore graphs in a number of open cases. There are still infinitely many open cases, and the topic is an active area of current research. While some limited further progress may be made with computational investigations, it is likely that new combinatorial ideas will be required to understand the problem fully. One direction in which the Cayley graph search technique might be extended is the problem of mixed *almost Moore* graphs, that is with order one less than the Moore bound. Although again there are infinitely many open cases, this problem seems just as intractable as in the Moore case.

In the girth problem, Chapter 7 derived new information about the graphs of Lazebnik, Ustimenko and Woldar. We have an open conjecture on the automorphism groups of these graphs, and it would be profitable to pursue this, perhaps by studying the stabilisers of 3-arcs in the graphs. Since these are still the best graphs known in an asymptotic sense, any new information might lead to useful research avenues.

The study of filled groups in Chapter 8 made substantial progress towards a classification of groups with this property. Filled groups, and the more general question of product-free sets in groups, have links to other areas of combinatorics including the degree-diameter problem, as we saw in the chapter. Partitions of a group into symmetric product-free sets (equivalent to the decomposition of a complete graph into edge-disjoint Cayley graphs) is an interesting topic in its own right, and has links to mainstream research areas such as Ramsey problems.

From Chapter 9, a promising line of research would be to try to improve the current position for arc-transitive graphs in the diameter problem. This is a new area, and it seems likely that constructions along similar lines to ones we have used for other related problems might improve our initial bounds.

In Chapter 10 we took a slight detour into the area of embeddings of maps on surfaces, although the underlying theme of studying graphs with a high degree of symmetry is constant. The groups $M(q^2)$ of this chapter are less studied in the

literature than their better-known cousins $PSL(2, q)$ and $PGL(2, q)$. One project will be to completely document the conjugacy classes and character tables of these groups. In addition, our understanding of the numbers of positively and negatively self-dual maps is currently based only on computational evidence, and seems resistant to theoretical attack using the same methods we used for the main enumeration. It would be interesting to try to resolve this.

Bibliography

[1] Marcel Abas, *Cayley graphs of diameter two and any degree with order half of the Moore bound*, Discrete Applied Mathematics **173** (2014), 1–7.

[2] ———, *Cayley graphs of diameter two with order greater than 0.684 of the Moore bound for any degree*, European Journal of Combinatorics **57** (2016), 109–120.

[3] Chimere Anabanti, Grahame Erskine, and Sarah Hart, *Groups whose locally maximal product-free sets are complete*, arXiv preprint arXiv:1609.09662.

[4] Chimere S. Anabanti and Sarah B. Hart, *On a conjecture of Street and Whitehead on locally maximal product-free sets*, Australas. J. Combin. **63** (2015), 385–398. MR 3414072

[5] Martin Bachratý, Jana Šiagiová, and Jozef Širáň, *Asymptotically approaching the Moore bound for diameter three by Cayley graphs*, preprint.

[6] Roger C Baker and Glyn Harman, *The difference between consecutive primes*, Proceedings of the London Mathematical Society **3** (1996), no. 2, 261–280.

[7] Roger C Baker, Glyn Harman, and János Pintz, *The difference between consecutive primes, II*, Proceedings of the London Mathematical Society **83** (2001), no. 3, 532–562.

[8] Eiichi Bannai and Tatsuro Ito, *On finite Moore graphs*, J. Fac. Sci. Tokyo Univ. **20** (1973), 191–208.

[9] Yakov Berkovich and Zvonimir Janko, *Groups of prime power order, vol. 1*, Volume 46 of de Gruyter Expositions in Mathematics, 2008.

[10] David Bevan, *Large "De Bruijn" Cayley graphs and digraphs*, arXiv preprint arXiv:1507.08926, July 2015.

[11] David Bevan, Grahame Erskine, and Robert Lewis, *Large circulant graphs of fixed diameter and arbitrary degree*, Ars Mathematica Contemporanea **13** (2017), no. 2, 275–291.

[12] Norman Biggs, *Algebraic graph theory*, Cambridge University Press, Cambridge, 1993.

[13] ———, *Constructions for cubic graphs with large girth*, Electron. J. Combin. **5** (1998), Article 1, 25 pp. (electronic). MR 1661181 (99j:05097)

[14] Béla Bollobás, *Modern graph theory*, vol. 184, Springer Science & Business Media, 2013.

[15] Juraj Bosák, *Partially directed Moore graphs*, Mathematica Slovaca **29** (1979), no. 2, 181–196.

[16] William G Bridges and Sam Toueg, *On the impossibility of directed Moore graphs*, Journal of Combinatorial theory, series B **29** (1980), no. 3, 339–341.

[17] William G Brown, *On graphs that do not contain a Thomsen graph*, Canad. Math. Bull **9** (1966), no. 2, 1–2.

[18] Josep M Brunat, Margarida Espona, Miguel Angel Fiol, and Oriol Serra, *On Cayley line digraphs*, Discrete mathematics **138** (1995), no. 1, 147–159.

[19] Dominique Buset, Mourad El Amiri, Grahame Erskine, Mirka Miller, and Hebert Pérez-Rosés, *A revised Moore bound for mixed graphs*, Discrete Mathematics **339** (2016), no. 8, 2066–2069.

[20] Peter J Cameron, *Portrait of a typical sum-free set*, Surveys in combinatorics **123** (1987), 13–42.

[21] Eduardo A Canale and José Gómez, *Asymptotically large (δ, d)-graphs*, Discrete applied mathematics **152** (2005), no. 1, 89–108.

[22] Marston Conder, Peter Dobscanyi, Brendan McKay, and Gordon Royle, *Cubic symmetric graphs (the Foster census)*, `http://staffhome.ecm.uwa.edu.au/ 00013890/remote/foster/`, Accessed 27/03/17.

[23] Marston Conder, Primož Potočnik, and Jozef Širáň, *Regular hypermaps over projective linear groups*, Journal of the Australian Mathematical Society **85** (2008), no. 02, 155–175.

[24] H. S. M. Coxeter and W. O. J. Moser, *Generators and relations for discrete groups.*, Springer-Verlag, Berlin, 1980.

[25] John Cullinan and Farshid Hajir, *Primes of prescribed congruence class in short intervals*, Integers **12** (2012), A56.

[26] Charles Delorme, *Grands graphes de degré et diametre donnés*, European Journal of Combinatorics **6** (1985), no. 4, 291–302.

[27] ———, *Examples of products giving large graphs with given degree and diameter*, Discrete Applied Mathematics **37** (1992), 157–167.

[28] Randall Dougherty and Vance Faber, *The degree-diameter problem for several varieties of Cayley graphs I: The abelian case*, SIAM Journal on Discrete Mathematics **17** (2004), no. 3, 478–519.

[29] Grahame Erskine, *Diameter 2 Cayley graphs of dihedral groups*, Discrete Mathematics **338** (2015), no. 6, 1022–1024.

[30] Grahame Erskine, Katarína Hriňáková, and Jozef Širáň, *Orientably-regular maps on twisted linear fractional groups*, arXiv preprint arXiv:1701.05781.

[31] Geoffrey Exoo and Robert Jajcay, *Dynamic cage survey*, Electronic Journal of Combinatorics **DS16v3** (2013), 55pp.

[32] Ramiro Feria-Puron, Hebert Perez-Roses, and Joe Ryan, *Searching for large circulant graphs*, arXiv preprint arXiv:1503.07357, 2015.

[33] Ferdinand Georg Frobenius, *Über Gruppencharaktere*, (1896).

[34] Ernst M Gabidulin, Alexander A Davydov, and Leonid M Tombak, *Linear codes with covering radius 2 and other new covering codes*, IEEE Transactions on Information Theory **37** (1991), no. 1, 219–224.

[35] The GAP Group, *GAP – Groups, Algorithms, and Programming, Version 4.8.6*, 2016.

[36] Michael Giudici, *Maximal subgroups of almost simple groups with socle $PSL(2, q)$*, arXiv preprint math/0703685 (2007).

[37] Michael Giudici and Sarah Hart, *Small maximal sum-free sets*, Electronic Journal Of Combinatorics **16** (2009), no. 1, R59.

[38] Alan J Hoffman and Robert R Singleton, *On Moore graphs with diameters 2 and 3*, IBM Journal of Research and Development **4** (1960), no. 5, 497–504.

[39] G. A. Jones, *Combinatorial categories and permutation groups*, Ars Math. Contemp. **10** (2016), no. 2, 237–254.

[40] Gareth A Jones and David Singerman, *Theory of maps on orientable surfaces*, Proc. London Math. Soc.(3), vol. 37, 1978, pp. 273–307.

[41] Leif K. Jørgensen, *New mixed Moore graphs and directed strongly regular graphs*, Discrete Mathematics **338** (2015), no. 6, 1011 – 1016.

[42] W. H. Kautz, *Bounds on directed (d, k) graphs*, Theory of cellular logic networks and machines, AFCRL-68-0668, SRI Project 7258, Final Report (1968), 20–28.

[43] F. Lazebnik, V. A. Ustimenko, and A. J. Woldar, *A characterization of the components of the graphs $D(k, q)$*, Discrete Math. **157** (1996), 271–283. MR 1417299 (97h:05150)

[44] Felix Lazebnik and Shuying Sun, *Some families of graphs, hypergraphs and digraphs defined by systems of equations: a survey*, preprint available at `http://www.math.udel.edu/ lazebnik/papers/adgsurvey2016aa.pdf`, 2016.

[45] Felix Lazebnik and Vasiliy A Ustimenko, *Explicit construction of graphs with an arbitrary large girth and of large size*, Discrete Applied Mathematics **60** (1995), no. 1, 275–284.

[46] Felix Lazebnik, Vasiliy A Ustimenko, and Andrew J Woldar, *A new series of dense graphs of high girth*, Bulletin of the American Mathematical Society **32** (1995), no. 1, 73–79.

[47] ———, *New upper bounds on the order of cages*, Electronic Journal of Combinatorics **4** (1997), no. 2, R13.

[48] Dimitri Leemans and Julie de Saedeleer, *On the rank two geometries of the groups psl (2, q): part i*, Ars Mathematica Contemporanea **3** (2010), no. 2.

[49] Robert R. Lewis, *Improved upper bounds for the order of some classes of Abelian Cayley and circulant graphs of diameter two*, arXiv preprint arXiv:1506.02844, June 2015.

[50] Nacho López, Josep M Miret, and Cèsar Fernández, *Non existence of some mixed Moore graphs of diameter 2 using SAT*, Discrete Mathematics **339** (2016), no. 2, 589–596.

[51] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak, *Ramanujan graphs*, Combinatorica **8** (1988), no. 3, 261–277.

[52] AM Macbeath, *Generators of the linear fractional groups*, Proc. Symp. Pure Math, vol. 12, 1969, pp. 14–32.

[53] Heather Macbeth, Jana Šiagiová, and Jozef Širáň, *Cayley graphs of given degree and diameter for cyclic, Abelian, and metacyclic groups*, Discrete Mathematics **312** (2012), no. 1, 94–99.

[54] Heather Macbeth, Jana Šiagiová, Jozef Širáň, and Tomáš Vetrík, *Large Cayley graphs and vertex-transitive non-Cayley graphs of given degree and diameter*, Journal of Graph Theory **64** (2010), no. 2, 87–98.

[55] Brendan McKay, *nauty user's guide (version 1.5), Technical report TR-CS-90-02*, Australian National University, Computer Science Department, 1990.

[56] Brendan D McKay, Mirka Miller, and Jozef Širáň, *A note on large graphs of diameter two and given maximum degree*, Journal of Combinatorial Theory, Series B **74** (1998), no. 1, 110–118.

[57] Mirka Miller and Jozef Širáň, *Moore graphs and beyond: A survey of the degree/diameter problem*, Electronic Journal of Combinatorics **DS14v2** (2013), 92pp.

[58] H. L. Montgomery and R. C. Vaughan, *The large sieve*, Mathematika **20** (1973), no. 02, 119.

[59] Minh Hoang Nguyen and Mirka Miller, *Moore bound for mixed networks*, Discrete Mathematics **308** (2008), no. 23, 5499 – 5503.

[60] Minh Hoang Nguyen, Mirka Miller, and Joan Gimbert, *On mixed Moore graphs*, Discrete Mathematics **307** (2007), no. 7, 964–970.

[61] J Plesnık and Š Znám, *Strongly geodetic directed graphs*, Acta FRN Univ. Comen.-Mathematica **29** (1974), 29–34.

[62] Alexander Pott and Yue Zhou, *Cayley graphs of diameter two from difference sets*, arXiv preprint arXiv:1506.05780, June 2015.

[63] Olivier Ramaré and Robert Rumely, *Primes in arithmetic progressions*, Math. Comp. **65** (1996), no. 213, 397–425. MR 1320898 (97a:11144)

[64] Chih-han Sah, *Groups related to compact Riemann surfaces*, Acta Mathematica **123** (1969), no. 1, 13–42.

[65] Jana Šiagiová and Jozef Širáň, *Approaching the Moore bound for diameter two by Cayley graphs*, Journal of Combinatorial Theory, Series B **102** (2012), no. 2, 470–473.

[66] Jana Šiagiová, Jozef Širáň, and Mária Ždímalová, *Large graphs of diameter two and given degree*, International Workshop on Optimal Network Topologies, Iniciativa Digital Politècnica, 2011.

[67] Jozef Širáň, *How symmetric can maps on surfaces be?*, Surveys in Combinatorics 2013 **409** (2013), 161.

[68] N. J. A. Sloane (editor), *The on-line encyclopedia of integer sequences*, `http://www/oeis.org/`, accessed 2016-11-08.

[69] Leonard Soicher, *The GRAPE package for GAP, Version 4.7*, 2016.

[70] Anne Penfold Street and Earl Glen Whitehead, *Group Ramsey theory*, Journal of Combinatorial Theory, Series A **17** (1974), no. 2, 219–226.

[71] James Tuite, *Diameter 6 construction*, Private communication.

[72] Tomáš Vetrík, *Large Cayley digraphs of given degree and diameter*, Discrete Mathematics **312** (2012), no. 2, 472–475.

[73] ———, *Cayley graphs of given degree and diameter 3, 4 and 5*, Discrete Mathematics **313** (2013), 213–216.

[74] ———, *Abelian Cayley graphs of given degree and diameter 2 and 3*, Graphs and Combinatorics **30** (2014), no. 6, 1587–1591.

[75] Hans Zassenhaus, *Über endliche Fastkörper*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg, vol. 11, Springer, 1935, pp. 187–220.

[76] Sanming Zhou, *Unitary graphs*, Journal of Graph Theory **75** (2014), no. 1, 37–47.